


COSYNE

MAIN CONFERENCE

**Denver, CO:
27 feb - 1 mar**

2020

cosyne.org



SIMONS
FOUNDATION

BURROUGHS
WELLCOME
FUND 

Program Summary

Thursday, 27 February

- 4:00p **Registration opens**
- 4:45p **Welcome reception**
- 6:15p **Opening remarks**
- 6:30p **Session 1: The ABCs—Artificial and Biological Circuits**
Invited speaker: Matthew Botvinick; 2 accepted talks
- 8:00p **Poster Session I**

Friday, 28 February

- 8:30a **Session 2: Latent space**
Invited speaker: Mehrdad Jazayeri; 3 accepted talks
- 10:30a **Session 3: Learning to decide**
Invited speaker: Linda Wilbrecht; 3 accepted talks
- 12:00n **Lunch break**
- 2:00p **Session 4: From theory to practice**
Invited speaker: John Cunningham; 3 accepted talks
- 4:15p **Session 5:**
Invited speaker: Marta Zlatic; 2 accepted talks
- 5:30p **Dinner break**
- 8:00p **Poster Session II**

Saturday, 29 February

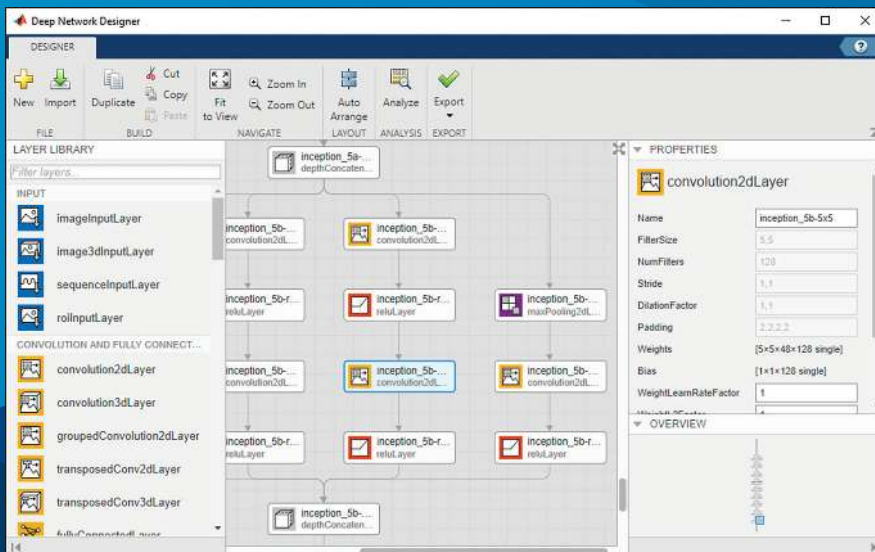
- 8:30a **Session 6: Visual hierarchy**
Invited speaker: Hendrikje Nienborg; 3 accepted talks
- 10:30a **Session 7: Motor learning**
Invited speaker: Megan Carey; 3 accepted talks
- 12:00n **Lunch break**
- 2:00p **Session 8: Planning and exploration**
Invited speakers: Sam Gershman, Wei Ji Ma
- 4:15p **Session 9: The social brain**
Invited speaker: Gul Dolen; 2 accepted talks
- 5:30p **Dinner break**
- 8:00p **Poster Session III**

Sunday, 01 March


- 8:30a **Session 10: Where am I?**
 Invited speaker: Lisa Giocomo; 3 accepted talks
- 10:30a **Session 11: What and where is that smell?**
 Invited speaker: Rainer Friedrich; 3 accepted talks
- 12:00n **Lunch break**
- 2:00p **Session 12: Can you be flexible?**
 Invited speaker: Christopher Harvey; 3 accepted talks

Deep Learning for Novices and Pros


Use MATLAB® to import, design, train, evaluate, and deploy deep networks



Labeled Ground Truth



Network Prediction



Learn more at mathworks.com/deep-learning


MATLAB®

 MathWorks®



Neurotechnology Research Systems

DataLogger - Wireless Neural Recording System



The New DataLogger system is an untethered, battery-powered neural recording system that allows wideband continuous recording at 40 kHz sampling rate. The DataLogger system consists of the Base unit, a 32 channel headstage, an infrared emitter, and the BaseControl software.

OmniPlex - Neural Data Acquisition System

Plexon's flagship neural data acquisition system for program selectable amplification, filtering and spike sorting. Available in 16 to 512 channel configurations, with either the DigiAmp™ for use with analog headstages or our newest Digital Headstage Processor (DHP) subsystem utilizing digital headstages.



N-FORM® U.S. Patent Nos. 9,095,267, 10,086,192 Patents Pending

N-Form Array - Customizable 3D Neural Probe

Plexon Inc is excited to announce a partnership with Modular Bionics™ Inc. to offer the N-Form® for research applications. The N-Form® array is a customizable, chronically implanted, 3D micro electrode array for use in long-term recording experiments.

Linear Probes



S Probe



V Probe



U Probe

To contact a Plexon Sales Engineer for more information visit us at plexon.com or find us on Facebook @Plexoninc or on Twitter @PlexonNeurosci



NeuroNexus

512 Channels

Finally, something worth watching

Meet the new SmartBox Pro




More info at www.NeuroNexus.com

**A high-performance, multi-functional instrument
used for signal acquisition, experiment control,
and probe diagnostics**

5200 S State St, Suite 200, Ann Arbor, MI 48108

Phone: +1.734.915.8858

Email: sales@neuronexus.com



grapevine

neural interface system

Hardware and Software for Advanced Electrophysiology

Signal Processing



- Supports up to 512 channels of stimulation and recording
- Optimized for complex experimental designs
- Compact, portable form-factor



Equipment Interfacing



Recording & Stimulating



RZ10X

A turnkey solution for calcium imaging and optogenetic stimulation.

Superior simplicity - one processor integrates all your fiber photometry needs.



Fiber Photometry & Optogenetics

- TDT's RZ10X is fully configured for recording various fluorophore responses, including GCaMP and dLight.
- With Synapse software, simplify your experiment design, testing, and data collection.
- View real-time GCaMP responses in raw and $\Delta F/F$ plots for easy data visualization.

Come by our Booth to see our latest industry leading recording system today.





3i Intelligent
Imaging
Innovations

ADVANCED MULTIMODAL MICROSCOPY SYSTEMS

Lattice LightSheet™

Spatiotemporal Superresolution Microscope

Marianas LightSheet™

Versatile Live Cell diSPIM Microscope

Marianas™

Inverted Microscope Platform for Live Cell Research

Cleared Tissue LightSheet™

High speed high resolution imaging of cleared tissue and whole organs

VIVO™


Intravital Microscope for Neuroscience and Cardiovascular Research

VIVO Multiphoton™

Modular Multiphoton Microscope Series

SlideBook 2020

SlideBook digital microscopy software advances research microscopy through the entire experimental process. By managing everything from instrument control to image processing and data analysis, SlideBook allows scientists to focus on investigation rather than instrumentation.



VIVO Multiphoton OPEN configuration for all-optical electrophysiology. **Phasor** computer-generated holography, combined with **Temporal Focusing** uses "sculpted light" to provide selective, precise optogenetic stimulation of many individual neurons simultaneously. Stimulation is seamlessly integrated with the dual galvanometric and resonant scanning acquisition supported by the **Vector RS+ Hybrid Scanner**. Three Photon excitation and 2pFLIM options are also available.

About Cosyne

The annual Cosyne meeting provides an inclusive forum for the exchange of experimental and theoretical/computational approaches to problems in systems neuroscience.

To encourage interdisciplinary interactions, the main meeting is arranged in a single track. A set of invited talks are selected by the Executive Committee and Organizing Committee, and additional talks and posters are selected by the Program Committee, based on submitted abstracts.

Cosyne topics include (but are not limited to): neural basis of behavior, sensory and motor systems, circuitry, learning, neural coding, natural scene statistics, dendritic computation, neural basis of persistent activity, nonlinear receptive field mapping, representations of time and sequence, reward systems, decision-making, synaptic plasticity, map formation and plasticity, population coding, attention, and computation with spiking networks. Participants include pure experimentalists, pure theorists, and everything in between.

Cosyne 2020 Leadership

Organizing Committee

General Chairs

Eugenia Chiappe (Champalimaud) and Christian Machens (Champalimaud)

Program Chairs

Anne-Marie Oswald (University of Pittsburgh) and Srdjan Ostojic (Ecole Normale Supérieure Paris)

Workshop Chairs

Catherine Hartley (New York University) and Blake Richards (McGill University)

Undergraduate Travel Chairs

Angela Langdon (Princeton University) and Robert Wilson (University of Arizona)

Publicity Chair

Adam Calhoun (Princeton University)

Diversity Chairs

Eva Dyer (Georgia Tech, Emory University) and Eric Shea-Brown (University of Washington)

Communications Chair

Xaq Pitkow (Rice University)

Tutorial Chair

Il Memming Park (Stony Brook University)

Development Chair

Michael Long (New York University)

Executive Committee

Stephanie Palmer (University of Chicago)

Zachary Mainen (Champalimaud)

Alexandre Pouget (University of Geneva)

Anthony Zador (Cold Spring Harbor Laboratory)

Program Committee

Anne-Marie Oswald (University of Pittsburgh), co-chair
Srdjan Ostojic (Ecole Normale Supérieure Paris), co-chair
Athena Akrami (University College London)
Demba Ba (Harvard University)
Omri Barak (Technion)
Brice Bathellier (Paris)
Laura Busse (LMU Munich)
Yoram Burak (Hebrew University)
Steve Chase (Carnegie Mellon University)
Anne Collins (University of California, Berkeley)
Christine Constantinople (New York University)
Saskia De Vries (Allen Institute for Brain Science)
Victor de Lafuente (National Autonomous University of Mexico)
Jan Drugowitsch (Harvard University)
Tatiana Engel (Cold Spring Harbor Laboratory)
Sean Escola (Columbia University)
Annegret Falkner (Princeton University)
Kevin Franks (Duke University)
Julijana Gjorgjieva (Max Planck Institute, Frankfurt)
Tim Hanks (University of California, Davis)
Ann Hermundstad (Janelia Farm Research Campus)
Andrew Leifer (Princeton University)
Mattieu Louis (University of California, Santa Barbara)
Jakob Macke (Technical University of Munich)
Mackenzie Mathis (Ecole polytechnique fédérale de Lausanne)
Cris Niell (University of Oregon)
Adrien Peyrache (McGill University)
Xaq Pitkow (Rice University)
Christina Savin (New York University)
Marshall Shuler (Johns Hopkins University)
Robert Rosenbaum (University of Notre Dame)
Daniela Vallentin (Max Planck Institute, Munich)
Joel Zylberberg (York University)

Cosyne 2020 reviewers

Luigi Acerbi, Alireza Alemi, Cristiano Alessandro, Mikio Aoi, Yeka Aponte, Hiroki Asari, Asli Ayaz, Shahab Bakhtiari, Joao Barbosa, Tania Barkat, Andrea Barreiro, Eleanor Batty, Jeffrey Beck, Philipp Berens, Gordon Berman, Edgar Bermudez, Michael Beyeler, Johannes Bill, Yazan Billeh, William Bishop, Kevin Bolding, Kathryn Bonnen, Hannah Bos, Yves Boubenec, Aine Byrne, Denise Cai, Adam Calhoun, Dario Campagner, Ioana Carcea, N Alex Cayco Gajic, Sourish Chakravarty, Chandramouli Chandrasekaran, Adam Charles, Rishidev Chaudhuri, Xiaomo Chen, Hannah Choi, SueYeon Chung, Kelly Clancy, Adam Claridge-Chang, Ruben Coen-Cagli, Julia Cox, Matt Creamer, Frédéric Crevecoeur, Ran Darshan, Ian Davison, Jaime de la Rocha, Kristen Delevich, Daniel Denman, Ashesh Dhawale, Long Ding, Mario Dipoppa, Benjamin Dunn, Eva Dyer, Rodrigo Echeveste, Alexander S Ecker, Felix Effenberger, Robert Egger, Seth Egger, Roian Egnor, Valerie Ego-Stengel, Sami El-Boustani, Ahmed El-Hady, Ben Engelhard, Vahid Esmaeili, Rose Faghih, Shiva Farashahi, Katie Ferguson, Andrew Fink, Walter Fischler, Alexander Fleischmann, Nick Franklin, Izumi Fukunaga, Marina Garrett, Jeff Gavornik, Gabrielle Girardeau, Chad Giusti, Joshua Glaser, Matt Golub, Pedro Goncalves, Andrew Gordus, Yi Gu, Ranier Gutierrez, Hannah Haberkern, Ralf Haefner, Steffen Hage, Arif Hamid, Adam Hantman, Kiah Hardcastle, Guillaume Hennequin, Matthias Hennig, Yu Hu, Chengcheng Huang, Simon Jacob, Monika Jadi, Xiaoxuan Jia, Maximilian Jösch, Kresimir Josic, Vijay Mohan K Nambodiri, Jonathan Kadmon, Naama Kadmon Harpaz, Raphael Kaplan, Nikolaos Karalis, Matthias Kaschube, Loukia Katsouri, Alfred Kaye, Hokto Kazama, Alex Kell, Ann Kennedy, Sepiedeh Keshavarzi, Zachary Kilpatrick, Christopher Kim, Sung Soo Kim, Ugne Klibaite, Sue Ann Koay, Mihaly Kollo, Charles Kopec, Nataliya Kraynyukova, Jens Kremkow, Kishore Kuchibhotla, Guillaume Lajoie, Angela Langdon, Richard Lange, Benjamin Lansdell, Kenneth Latimer, Yaara Lefler, Luis Lemus, Anna Levina, Sam Lewallen, Liang Liang, Sukbin Lim, Scott Linderman, Grace Lindsay, Ashok Litwin-Kumar, Yoav Livneh, Laureline Logiaco, Matthew Lovett-Barron, Jan-Matthis Lueckmann, Emily Mackevicius, David Margolis, Olivier Marre, Jesse Marshall, Carlotta Martelli, Tanya Marton, Paul Masset, Francesca Mastrogiuseppe, Alexander Mathis, Luca Mazzucato, Matthew McGinley, Jenna McHenry, Jorge Mejias, Raoul-Martin Memmesheimer, Jorge Menendez, Leenoy Meshulam, Florent Meyniel, Daniel Millman, Aaron Milstein, Andrew Miri, Gal Mishne, Wiktor Mlynarski, Joseph Monaco, Ruben Moreno Bote, James Murray, Malavika Murugan, Farzaneh Najafi, Miho Nakajima, Devika Narain, Eva Naumann, Kristina Nielsen, Cian O'Donnell, Daniel O'Connor, Gabriel Ocker, Michael Okun, Amy Orsborn, Yusuf Osmanlioglu, David Oswald, Marino Pagan, Agostina Palmigiano, Chethan Pandarinath, Cengiz Pehlevan, Abigail Person, Megan Peters, Warren Pettine, Alex Piet, Lucas Pinto, Eric Plourde, Cindy Poo, Andrew Pruszynski, Maximilian Puelma Touzel, Srikanth Ramaswamy, Pavan Ramdya, Erin Rich, Lisa Roux, Alon Rubin, Mikail Rubinov, Pavel Rueda Orozco, Douglas Ruff, Caroline Runyan, Chaitanya Ryali, João Sacramento, Shreya Saxena, Ian Schmitt, David Schneider, Monika Scholz, Carl Schoonover, Manuel Schottdorf, Tilo Schwaller, Benjamin Scott, Kelly Seagraves, Madineh Sedigh-Sarvestani, Nidhi Seethapathi, Hideaki Shimazaki, Marshall Shuler, Joshua Siegle, Julie Simpson, Fabian H Sinz, Anne Smith, Spencer Smith, Hansem Sohn, Kimberly Stachenfeld, Natalie A Steinemann, Merav Stern, Weinan Sun, Vladislav Susoy, David Sussillo, Sina Tafazoli, Corinne Teeter, Dougal Tervo, Kay Thurley, Balazs Ujfalussy, Anne Urai, Jonathan Vacher, Timo van Kerkoerle, Julia Veit, Vivek Venkatachalam, Martin Vinck, Siwei Wang, Michael Waskom, Xue-Xin Wei, Bob Wilson, Marina Wosniack, Guangyu Robert Yang, Ilker Yildirim, Noga Zaslavsky, Fleur Zeldenrust, Friedemann Zenke, Wenhao Zhang, and Petr Znamenskiy.

Special thanks to Titipat Achakulvisut, Daniel Acuna, and Konrad Kording for writing and managing the automated software for reviewer abstract assignment.

About Cosyne

Conference Support

Administrative Support, Registration, Hotels

Leslie Weekes, Cosyne

Social media policy

Cosyne encourages the use of social media before, during, and after the conference, so long as it falls within the following rules:

- Do not capture or share details of any unpublished data presented at the meeting.
- If you are unsure whether data is unpublished, check with the presenter before sharing the information.
- Respect presenters' wishes if they indicate the information presented is not to be shared.

Stay up to date with Cosyne 2020 #cosyne20

Cosyne Talks Channel

Selected Cosyne talks are available on Cosyne YouTube Channel. *Cosyne is very grateful to Carlos Stein Brito, Robin Cao, and Charline Tessereau for recording and broadcasting the conference presentations in 2020 and previous years.*

Travel Grants

The Cosyne community is committed to bringing talented scientists together at our annual meeting, regardless of their ability to afford travel. Thus, a number of travel grants are awarded to students, postdocs, and PIs for travel to the Cosyne meeting. Each award covers at least \$500 towards travel and meeting attendance costs. Five award granting programs were available for Cosyne 2020.

The generosity of our sponsors helps make these travel grant programs possible. Cosyne Travel Grant Programs are supported entirely by the following corporations and foundations:

- The Gatsby Charitable Foundation
- Burroughs Wellcome Fund
- Simons Foundation
- DeepMind



Cosyne Presenters Travel Grant Program

These grants support early career scientists with highly scored abstracts to enable them to present their work at the meeting.

The 2020 recipients are

Margarida Agrochao, Pinglei Bao, Manuel Beiran, Matthias Christenson, Alexandra Ferguson, Jonathan Gill, Evren Gokcen, Ni Ji, Yul Kang, Jan Hendrik Kirchner, Darby Losey, Hirofumi Nakayama, Dina Obeid, Ioannis Pisokas, Mahdi Ramadan, Nimrod Shaham, Sergey Shuvaev, Iris Stone, Xulu Sun, and Avner Wallach.

Cosyne New Attendees Travel Grant Program

These grants help bring scientists that have not previously attended Cosyne to the meeting for exchange of ideas with the community.

The 2020 recipients are

Dhruba Banerjee, Aanchal Bhatia, Ho Yin Chau, Prannoy Chaudhuri-Vayalambone, Weiwei Chen, Jessica Dafflon, Subhadra Mokashe, Katharina Dobs, Aleksander Domanski, Pablo Franco, Sven Goedeke, Siddharth Jayakumar, Jessica Kohn, Marino Krstulovic, Divyansh Mittal, Samia Mohinta, Philipp Norton, Matthijs oude Lohuis, Fabian Pallasdies, Youngjin Park, Amol Pasarkar, Krishna Pusuluri, Drew Schreiner, Patricia Stan, Kyra Swanson, Danil Tyulmankov, Praveen Venkatesh, Shenghao Wu, Sichao Yang, and Meytar Zemer.

Cosyne Mentorship Travel Grant Program

These grants provide support for early-career scientists of underrepresented minority groups to attend the meeting. A Cosyne PI must act as a mentor for these trainees and the program also is meant to recognize these PIs ("Cosyne Mentors").

The 2020 Cosyne Mentors and mentees are:

Dori Derdikman and Dorgham Khatib, Il Memming Park and Diego Arribas, Henning Sprekeler and Hymavathy Balasubramanian.

Cosyne Undergraduate Travel Grant Program

These grants help bring promising undergraduate students with strong interest in neuroscience to the meeting.

The 2020 recipients are

Robert Dabagia, Lavonna Mark, Habiba Noamany, Diana Polhac, Yongsoo Ra, Noushin Quazi, Michael Xie, and Peng Xu.

Cosyne Childcare Travel Grant Program

Cosyne Childcare Grants help cover childcare expenses incurred by participation in Cosyne 2020.

The 2020 recipients are

Nadav Ben Shushan, Alexis Dubreuil, Jonathan Kadmon, Julija Krupic, Aneesha Sahu, Sacha Sokoloski, Hidenori Tanaka, and Sudha Yellapantula.

Code of Conduct at Cosyne

Cosyne is adopting the practice of defining a meeting code of conduct and providing information for participants on our ethics policies as well as resources for filing complaints. This code of conduct is based on standards and language set at other meetings, whose organizing boards convened special working groups of scientific and legal experts to set their policies. We follow, in particular, those guidelines established for the Gordon Research Conferences and the Society for Neuroscience Annual Meeting.

The following code of conduct has been adapted from:

<https://www.grc.org/about/grc-policies-and-legal-disclaimers>

<https://www.sfn.org/Membership/Professional-Conduct/Code-of-Conduct-at-SfN-Events>

other online resources:

<http://changingourcampus.org>

<https://www.sfn.org/Membership/Professional-Conduct/SfN-Ethics-Policy>

At Cosyne, we strive for open and honest intellectual debate as part of a welcoming and inclusive atmosphere. Cosyne asks each session chair and organizing and reviewing committee member to promote rigorous analysis of all science presented for or at the meeting in a manner respectful to all attendees. To help maintain an open and respectful community of scientists, Cosyne does not tolerate illegal or inappropriate behavior at the main meeting or the workshops, including violations of applicable laws pertaining to sale or consumption of alcohol, destruction of property, or harassment of any kind, including sexual harassment.

Cosyne condemns inappropriate or suggestive acts or comments that demean another person by reason of their gender, gender identity or expression, race, religion, ethnicity, age or disability or that are unwelcome or offensive to other members of the community or

their guests. Cosyne will review allegations of any such behavior on a case-by-case basis, and violations may result in the prohibition on future attendance by particular individuals.

Cosyne believes home institutions and employers of our event attendees, including research institutions, companies, and other organizations, are best equipped to investigate allegations or violations by their faculty members, trainees, employees or affiliated individuals, and to evaluate and determine appropriate actions consistent with their employment and academic obligations. In the event of an allegation of harassment at a Cosyne-organized event, Cosyne will document and will generally refer allegations to an attendee's organization for its review. To the extent possible, documentation will be held in confidence by Cosyne's executive staff.

If you believe you have been subjected to or have otherwise experienced behavior at a Cosyne event that violates Cosyne's Code of Conduct, please act promptly to report the issue so that steps may be taken to address the situation immediately. You may notify the chair of the session you are attending, or one of the meeting organizers, each of whom we have asked to address the matter promptly or refer it back to the Cosyne main office. If you are reluctant to speak with an organizer or session chair for any reason, you may notify Leslie Weekes, Cosyne's Executive Administrator, at the Cosyne main office by email to LESLIE.WEEKES@COSYNE.ORG.

Reports requesting anonymity will be respected, although Cosyne reserves the right to notify appropriate law enforcement if the allegations are deemed serious enough to warrant such notice. It should be noted that Cosyne's ability to investigate or address anonymous reports may be limited or otherwise affected by the need to balance concerns over privacy and fairness to all concerned. Following completion of Cosyne's inquiry, any action to be taken by Cosyne against the person accused of acting inappropriately will be determined by Cosyne in its sole discretion and may include discharge from the conference or restrictions on their future attendance.

Program

Note: Institutions listed in the program are the primary affiliation of the first author. For the complete list, please consult the abstracts.

Thursday, 27 February

- 8:00a **Neurodata without borders Tutorial**
- 12:00n **Normative approaches to understanding neural coding and behavior. *Cosyne 2020 Tutorial session sponsored by the Simons Foundation***
- 4:00p **Registration opens**
- 4:45p **Welcome reception**
- 6:15p **Opening remarks**

Session 1: The ABCs—Artificial and Biological Circuits

(Chair: Anne-Marie Oswald, Srdjan Ostojic)

- 6:30p [Deep reinforcement learning and its neuroscientific implications](#)
Matthew Botvinick, DeepMind (**invited**) 29
- 7:30p [Flexible neuronal dynamics during categorization task switching in artificial and biological neural networks](#)
K. Mohan, O. Zhu, D. Freedman, The University of Chicago 34
- 7:45p [Dissecting modularity and redundancy in multi-regional circuits by population recording and modeling](#)
B. Kang*, G. Chen*, N. Li**, S. Druckmann**, Stanford University 34
- 8:00p **Poster Session I**

Friday, 28 February

Session 2: Latent space

(Chair: David Sussillo)

- 8:30a [Computation-through-dynamics for the temporal control of behavior](#)
Mehrdad Jazayeri, Massachusetts Institute of Technology (**invited**) 29
- 9:15a [Flexible neural control of motor units revealed via latent factor models](#)
N. Marshall, J. Glaser, S. Perkins, L. Abbott, J. Cunningham, M. Churchland, Columbia University 35
- 9:30a [Evidence of a memory trace in motor cortex after short-term learning](#)
D. Losey, J. Hennig, E. Oby, M. Golub, P. Sadtler, K. Quick, S. Ryu, E. Tyler-Kabara, A. Batista, B. Yu, S. Chase, Carnegie Mellon University 36
- 9:45p [Separation of preparatory neural states when learning multiple arm-movement dynamics](#)
X. Sun, D. O’Shea, M. Golub, E. Trautmann, S. Ryu, K. Shenoy, Stanford University 36
- 10:00a **Coffee break**

Program

Session 3: Learning to decide

(Chair: Tatiana Engel)

10:30a	The role of the striatal indirect pathway/D2R+ spiny projection neurons in choice and rejection Linda Wilbrecht, University of California, Berkeley (invited)	30
11:15a	Long-range inhibition mediates decision making in the superior colliculus J. Essig, J. Hunt, G. Felsen, University of Colorado	37
11:30a	Spontaneous activity and multi-sensory integration in the developing higher-order cortex M. Wosniack, J. H. Kirchner, P. Maldonado, N. Zabouri, C. Lohmann, J. Gjorgjieva, Max Planck Institute for Brain Research	38
11:45a	Learning rate adjustments in a volatile environment by mouse medial prefrontal cortex H. Atilgan, C. Murphy, H. Wang, A. Kwan, Yale University	38
12:00n	Lunch break	
12:30p	Workshop <i>Growing up in science</i> Wei Ji Ma	

Session 4: From theory to practice

(Chair: Christina Savin)

2:00p	Automating parameter inference in theoretical neuroscience John Cunningham, Columbia University (invited)	30
2:45p	Sequential Component Analysis (SCA) V. Rutten, A. Bernacchia, G. Hennequin, Gatsby Computational Neuroscience Unit, University College London	39
3:00p	Nonlinear computations in semi-balanced networks C. Baker, V. Zhu, R. Rosenbaum, University of Notre Dame	39
3:15p	Gradient-based learning with Hebbian plasticity in structured and deep neural networks D. Obeid, C. Pehlevan, Harvard University	40
3:30p	Coffee break	

Session 5:

(Chair: Ashok Litwin-Kumar)

4:15p	Circuits for learning, predicting, and deciding in <i>Drosophila</i> Marta Zlatic, HHMI Janelia Research Campus (invited)	30
5:00p	Neural mechanism for illusory motion perception from stationary patterns M. Agrochao, R. Tanaka, D. Clark, E. Salazar-Gatzimas, Yale University	40
5:15p	Sensory control of neuromodulation enables flexible selection of behavioral states N. Ji, G. Madan, G. Fabre, A. Dayan, C. Baker, I. Nwabudike, S. Flavell, Massachusetts Institute of Technology	41
5:30p	Dinner break	
8:00p	Poster Session II	

Saturday, 29 February

Session 6: Visual hierarchy

(Chair: Xaq Pitkow)

8:30a	Evidence for task agnostic feedback in visual cortex Hendrikje Nienborg, NEI/NIH (invited)	31
9:15a	Dissecting feedforward and feedback interactions between populations of neurons E. Gokcen, J. Semedo, A. Zandvakili, C. Machens, A. Kohn, B. Yu, Carnegie Mellon University	41
9:30a	Tracking information flow in subnetworks of mouse visual cortex X. Jia, J. Siegle, S. Olsen, Allen Institute for Brain Science	42
9:45p	Assembly structure expands the dimension of shared variability in cortical networks D. Rager, S. Khanna, M. Smith, B. Doiron, Carnegie Mellon University	43
10:00a	Coffee break	

Session 7: Motor learning

(Chair: Daniela Vallentin)

10:30a	The complex behaviors of the 'simple' cerebellar circuit Megan Carey, Champalimaud Centre for the Unknown (invited)	31
11:15a	Isolating a locus for reach adaptation in the cerebellar cortex D. Calame, M. Becker, A. Person, University of Colorado	43
11:30a	Identical encoding of flexible and automatic motor sequences in the dorsal lateral striatum K. Mizes, B. Olveczky, Harvard University	44
11:45a	Exploring the hidden state of Coma S. Mofakham, A. Fry, J. Nassar, J. Adachi, C. Mikell, Stony Brook University	44
12:00n	Lunch break	
12:00n	Diversity and inclusion lunch	

Session 8: Planning and exploration

(Chair: Jan Drugowitsch)

2:00p	The architecture of exploration Sam Gershman, Harvard University (invited)	32
2:45p	Human planning in large state spaces Wei Ji Ma, New York University (invited)	32
3:30p	Coffee break	

Session 9: The social brain

(Chair: Annegret Falkner)

4:15p	Parallel social information processing circuits are differentially impacted in autism Gul Dolen, Johns Hopkins University (invited)	32
5:00p	Deep neural network modeling of visuomotor transformations during social interaction A. Calhoun, B. Cowley, J. Pillow, M. Murthy, Princeton University	45
5:15p	Multi-region network models of brain-wide interactions during adaptive and maladaptive state transitions M. Perich, A. Andalman, E. Carter, K. Rajan, K. Deisseroth, Icahn School of Medicine at Mt. Sinai	45

Program

5:30p **Dinner break**
8:00p **Poster Session III**

Sunday, 01 March

Session 10: Where am I?

(Chair: Yoram Burak)

- 8:30a [Multiple maps for navigation](#)
Lisa Giocomo, Stanford University (**invited**) 33
- 9:15a [Accurate angular integration with only a handful of neurons](#)
M. Noorman, V. Jayaraman, S. Romani, A. Hermundstad, HHMI Janelia Research Campus 46
- 9:30a [Preexisting hippocampal network dynamics constrain optogenetically induced place fields](#)
S. McKenzie, G. Buzski, R. Huszar, D. English, K. Kim, E. Yoon, New York University . . . 47
- 9:45p [A map of object space in primate inferotemporal cortex](#)
P. Bao, L. She, M. McGill, D. Tsao, California Institute of Technology 47
- 10:00a **Coffee break**

Session 11: What and where is that smell?

(Chair: Anne-Marie Oswald)

- 10:30a [Connectivity and computation in olfaction](#)
Rainer Friedrich, FMI Basel (**invited**) 33
- 11:15a [A map for odors and place in posterior piriform cortex](#)
C. Poo, G. Agarwal, N. Bonnachi, Z. Mainen, Champalimaud Centre for the Unknown . . . 48
- 11:30a [Rapid representational drift in primary olfactory cortex](#)
C. Schoonover, A. Fink, R. Axel, Columbia University 48
- 11:45a [Manipulating synthetic optogenetic odors reveals the coding logic of olfactory perception](#)
E. Chong, M. Moroni, C. Wilson, S. Shoham, S. Panzeri, D. Rinberg, New York University 49
- 12:00n **Lunch break**

Session 12: Can you be flexible?

(Chair: Kanaka Rajan)

- 2:00p [A neuro-economic model for rats' willingness to work for water](#)
P. Reinagel, University of California, San Diego 49
- 2:15p [The role of frontal cortex in multisensory decision-making](#)
P. Coen, M. Wells, P. Zatka-Haas, M. Carandini, K. Harris, University College London . . . 50
- 2:30p [Task representation in evoked and spontaneous activity in the visual cortex](#)
M. Hajnal, G. Orban, P. Golshani, D. T. Tran, M. C. Einstein, P. Polack, HAS Wigner
Research Centre for Physics 51
- 2:45p [Cortical networks for flexible decisions during spatial navigation](#)
Christopher Harvey, Harvard University (**invited**) 33
- 3:30p **Closing remarks**

Poster Session I

8:00p Thursday 27 February

I-1. Chimera states as epileptic seizure predictors Nuttida Rungratsameetaweemana, Claudia Lainscsek, Sydney S Cash, Terrence Sejnowski, University of California, San Diego	51
I-2. Measuring and modeling a fine spatial scale clustering of orientation tuning in L2/3 of mouse V1 Peijia Yu, Yuhan Yang, Daisong Pan, Jiang Lan Fan, Chengcheng Huang, Mario Dipoppa, Kenneth Miller, Na Ji, Brent Doiron, Carnegie Mellon University	52
I-3. Computational cognitive requirements of random decision problems Pablo Franco, Karlo Doroc, Nitin Yadav, Peter Bossaerts, Carsten Murawski, University of Melbourne	53
I-4. Mechanisms of top-down attentional control in thalamic reticular circuits Qinglong Gu, Norman Lam, John D Murray, Yale School of Medicine	53
I-5. Time-dependent simultaneous configural and elemental odorant mixture processing Tingkai Liu, Chung-Heng Yeh, Aurel A Lazar, Columbia University	54
I-6. Inferring low-dimensional latent descriptions of animal vocalizations Jack Goffinet, John Pearson, Richard Mooney, Duke University	54
I-7. Model metamers reveal that deep neural network invariances differ from human perceptual invariances Jenelle Feather, Alex Durango, Ray Gonzalez, Josh McDermott, Massachusetts Institute of Technology	55
I-8. Prospective representation of navigational goals in orbitofrontal cortex Raunak Basu, Hiroshi Ito, Max Planck Institute for Brain Research	55
I-9. Different types of plasticity and interneurons shape circuit dynamics after sensory deprivation Leonidas Richter, Julijana Gjorgjieva, Max Planck Institute for Brain Research	56
I-10. STP controls non-linear responses and contrast-invariant tuning in spatially structured networks Laura Bernaez Timon, Tatjana Tchumatchenko, Max Planck Institute for Brain Research	57
I-11. Cognitive context alters cortical involvement in identical perceptual decisions Charlotte Arlt, Roberto Barroso-Luque, Ningjing Xia, Christopher Harvey, Harvard University	57
I-12. Hierarchical neural network models that more closely match primary visual cortex tend to better explain higher level visual cortical responses Tiago Marques, Martin Schrimpf, James J DiCarlo, Massachusetts Institute of Technology	58
I-13. Algorithmically engineered synaptic plasticity rules Basile Confavreux, Friedemann Zenke, Everton J Agnes, Tim Vogels, University of Oxford	58
I-14. Sensorimotor processing for tracking action outcomes in the zebrafish serotonergic system Takashi Kawashima, Ziqiang Wei, Sujatha Narayan, Liam Paninski, Misha Ahrens, Weizmann Institute of Science	59
I-15. Landmark stability needed for formation but not expression of spatial maps in retrosplenial cortex Dhruba Banerjee, Zaneta Navratilova, Bruce McNaughton, Sunil Gandhi, University of California, Irvine	60
I-16. Balancing excitation and inhibition in spiking networks with input dependent inhibitory plasticity (idip) Klara Kaleb, Victor Pedrosa, Claudia Clopath, Imperial College London	60
I-17. Multiscale associative memory recall by modulation of inhibitory circuits Tatsuya Haga, Tomoki Fukai, RIKEN Center for Brain Science	61
I-18. How do time and uncertainty motivate information seeking? Ethan Bromberg-Martin, Takaya Ogasawara, Yang-Yang Feng, J Kael White, Kaining Zhang, Ilya Monosov, Washington University in St Louis	61
I-19. Insights from a deep convolutional neural network into mid-level representation in visual cortex Wyeth Bair, Dean Pospisil, Anthony Bigelow, Dina Popovkina, Taekjun Kim, University of Washington	62

I-20. Probabilistic successor representations allow for flexible behaviour
 Jesse Geerts, Kimberly Stachenfeld, Neil Burgess, University College London 62

I-21. Multivariate phase coupling networks across brain regions using torus graphs
 Josue Orellana, Natalie Klein, Scott Brincat, Earl Miller, Robert Kass, Carnegie Mellon University 63

I-22. Inferring random network parameters from continuous-time trajectories
 Alexander van Meegen, Tobias Kuhn, Moritz Helias, Research Centre Jülich 64

I-23. Local synaptic balancing for network robustness
 Christopher Stock, Sarah Harvey, Samuel A Ocko, Surya Ganguli, Stanford University 64

I-24. Spectral regularization in biological and artificial neural networks
 Josue Nassar, Piotr A Sokol, SueYeon Chung, Kenneth Harris, Il Memming Park, Stony Brook University 65

I-25. IronClust: Scalable and drift-resistant spike sorting for long-duration, high-channel count recordings
 James Jun, Jeremy Magland, Catalin Mitelut, Alex Barnett, flatiron Institute 65

I-26. A common model explaining flexible decision making, grid fields, and cognitive control
 Payam Piray, Nathaniel D Daw, Princeton University 66

I-27. Recurrent neural network models of multi-area computation underlying decision-making
 Michael Kleinman, Chandramouli Chandrasekaran, Jonathan Kao, University of California, Los Angeles 67

I-28. Neuronal representations in mice barrel cortex and ANNs performing a whisker discrimination task
 Ramon Nogueira, Chris Rodgers, Randy Bruno, Stefano Fusi, Columbia University 67

I-29. Simultaneous dimensionality reduction and deconvolution for calcium imaging recordings
 Tze Hui Koh, William Bishop, Takashi Kawashima, Misha Ahrens, Steven Chase, Byron Yu, Carnegie Mellon University 68

I-30. Proprioceptive tuning emerges in a task-driven hierarchical deep neural network model
 Kai Sandbrink, Pranav Mamidanna, Claudio Michaelis, Mackenzie Mathis, Matthias Bethge, Alexander Mathis, Harvard University 68

I-31. Inter-joint coupling dynamics of walking *Drosophila* from complete 3D leg kinematics
 Pierre Karashchuk, Evyn Dickinson, Katie Rupp, Bingni Brunton, John Tuthill, University of Washington 69

I-32. Unsupervised learning of odor sequences in primary olfactory cortex
 Andrew Fink, Carl Schoonover, Richard Axel, Columbia University 69

I-33. The hourglass organization of the *C. elegans* connectome
 Constantine Dovrolis, Kaeser Sabrin, Yongbin Wei, Martijn van den Heuvel, Georgia Institute of Technology 70

I-34. Experimental demonstration of single neuron specificity during underactuated neurocontrol
 Samuel Brown, Jason Ritt, Boston University 70

I-35. Predicting functional similarity of neural networks from architectural features
 Adam Haber, Elad Schneidman, Weizmann Institute of Science 71

I-36. To what extent can the olfactory cortex be modeled by random connectivity?
 William Dorrell, Julien Grimaud, Venkatesh Murthy, Cengiz Pehlevan, Harvard University 72

I-37. Rule adherence warps choice representations and increases decision-making efficiency
 Becket Ebitz, Cindy Tu, Benjamin Hayden, University of Minnesota 72

I-38. A comprehensive analysis of local and global inputs to motion detection pathways
 Luis Giordano Ramos Traslosheros Lopez, Marion Silies, International Max Planck Research School Neurosciences 73

I-39. L2/3 and L5 pyramidal neuron somata and apical dendrites exhibit distinct responses to unexpected violations of visual flow
 Colleen Gillon, Joel Zylberberg, Blake Richards, Jerome Lecoq, Yoshua Bengio, Timothy P Lillicrap, Jed Perkins, Jason Pina, Sam Seid, Carol Thompson, Ryan Valenza, University of Toronto 74

I-40. Efficient high-dimensional receptive field inference using a flexible spline basis
 Ziwei Huang, Philipp Berens, University of Tübingen 74

I-41. Inhibitory plasticity rules create a contextual switch via complementary receptive fields Everton J Agnes, Andrea I Luppi, Tim Vogels, University of Oxford	75
I-42. Hippocampal CA3 neurons are organized into correlated cell assemblies supporting stable place codes Liron Sheintuch, Nitzan Geva, Alon Rubin, Yaniv Ziv, Weizmann Institute of Science	76
I-43. Dynamical learning of dynamics Christian Klos, Yaroslav Felipe Kalle Kossio, Sven Goedeke, Aditya Gilra, Raoul-Martin Memmesheimer, University of Bonn	76
I-44. Using unsupervised machine learning to understand the role of mPFC in social dominance in mice Kanha Batra, Nancy Padilla-Coreano, Sebastien Hausmann, Laurel Keyes, Ila fiete, Kay Tye, University of California, San Diego	77
I-45. A theory of visual search for foveated visual systems R Calen Walshe, Jared Abrams, Wilson S Geisler, The University of Texas	77
I-46. Modeling the separate functions of feedforward and feedback pathways in the visual system Grace Lindsay, Tom Mrcsic-flogel, Maneesh Sahani, Gatsby Computational Neuroscience Unit, University College London	78
I-47. Naive artificial intelligence Tomer Barak, Yonatan Loewenstein, Hebrew University of Jerusalem	78
I-48. Ignoring correlations causes a failure of retinal population codes under rod-mediated light levels Kiersten Ruda, Joel Zylberberg, Greg field, Duke University	79
I-49. Synaptic dynamics as a substrate for temporal learning Alessandro Barri, Martin Wiechert, David DiGregorio, Institut Pasteur	79
I-50. A two time-scale normative model of <i>C. elegans</i> sensory adaptation and behavior Hamilton White, Sruti Mallik, Baranidharan Raman, ShiNung Ching, Dirk Albrecht, Worcester Polytechnic Institute	80
I-51. Removing uncorrelated noise in system neuroscience using deep interpolation Jerome Lecoq, Michael Oliver, Natalia Orlova, Joshua Siegle, Allen Institute for Brain Science	81
I-52. Hearing-impaired deep neural networks replicate behavioral deficits of hearing-impaired humans Mark Saddler, Jenelle Feather, Andrew Francl, Josh McDermott, Massachusetts Institute of Technology	81
I-53. Steering neural population activity with adaptive closed-loop stimulation system Sina Tafazoli, Timothy Buschman, Camden MacDowell, Daniel Che, Kate Letai, Princeton University	82
I-54. Quantifying the role of divisive normalization in contextual modulation of neuronal variability Dylan Festa, Oren Weiss, Amir Aschner, Adam Kohn, Ruben Coen-Cagli, Albert Einstein College of Medicine	82
I-55. Dopamine firing activity codes reward expectation and drives motivation in a working memory task Nestor Parga, Stefania Sarno, Manuel Beiran, Jose Vergara, Rossi-Pool Roman, Ranulfo Romo, Universidad Autonoma Madrid	83
I-56. Top-down and bottom-up factors regulate cortical dynamics through Vip-Sst mutual inhibition Daniel Millman, Gabriel Ocker, Michael Oliver, Peter Ledochowitsch, Jerome Lecoq, Clay Reid, Michael Buice, Saskia de Vries, Allen Institute for Brain Science	84
I-57. The anterior dorsomedial frontal cortex is causally involved in regulating the time constant of evidence accumulation Thomas Luo, Carlos D Brody, Adrian Bondy, Brian DePasquale, Princeton University	84
I-58. Non-linear matrix factorization methods for extracting calcium traces in moving <i>C. elegans</i> videos Conor McGrory, Amin Nejatbakhsh, Erdem Varol, Columbia University	85
I-59. Emergence of opposite neurons in a circuit for multisensory processing Ho Yin Chau, Xueyan Niu, Wenhao Zhang, Tai Sing Lee, University of California, Berkeley	85

I-60. [Model selection approach for identifying distinct rhythmic entrainment profiles of CA1 interneurons](#)
 Pamela Riviere, Lara M Rangel, University of California, San Diego 86

I-61. [An olfactory pattern generator for functional analysis of neural circuit in Drosophila larva](#)
 Guangwei Si, Jacob Baron, Yu Feng, Aravinthan Samuel, Harvard University 86

I-62. [Slow manifolds and output attractors as the substrate for working memory in recurrent networks](#)
 Elham Ghazizadeh, Lawrence Snyder, ShiNung Ching, Washington University in St Louis 87

I-63. [Probabilistic knowledge of event timing to optimize reward seeking behaviors](#)
 Eizaburo Doi, Joshua Dudman, HHMI Janelia Research Campus 88

I-64. [A framework for unifying and generalizing models of neural dynamics during decision-making](#)
 David Zoltowski, Jacob Yates, Jonathan Pillow, Scott Linderman, Princeton University 88

I-65. [Learning recurrent dynamic patterns in spiking neural networks](#)
 Christopher Kim, Carson Chow, National Institutes of Health 89

I-66. [Burst initiation in the preBotzinger Complex as synchronization spreading on a random graph](#)
 Valentin Slepukhin, Sufyan Ashhad, Jack Feldman, Alex Levine, University of California, Los Angeles 89

I-67. [Evidence for multifaceted uncertainty-driven human exploration in a novel multi-armed bandit task](#)
 Angela Yu, Dalin Guo, florent Meyniel, University of California, San Diego 90

I-68. [The dynamic repertoire of neural activity across the forebrain: the view from the single neuron](#)
 Daniel Levenstein, John Rinzel, Gyorgy Buzsaki, New York University 90

I-69. [Deep reinforcement learning of extra-synaptic control in C. elegans chemosensory response](#)
 Jimin Kim, Eli Shlizerman, University of Washington 91

I-70. [Inverse rational control in continuous problems](#)
 Minhae Kwon, Saurabh Daptardar, Paul Schrater, Xaq Pitkow, Rice University 91

I-71. [VIP interneuron contributions to state-dependent sensory processing](#)
 Katie Ferguson, Calvin Fang, Jessica Cardin, Yale School of Medicine 92

I-72. [Automatically detecting single cell ion channel events of neural activity using generative adversarial and recurrent neural networks](#)
 Numan Celik, Fiona O'Brien, Yalin Zheng, Frans Coenen, Richard Barrett-Jolley, University of Liverpool 93

I-73. [Temporal compression as a mechanism for switching network operation modes](#)
 David Tingley, Gyorgy Buzsaki, New York University 93

I-74. [Local changes, global effects: how localised is the effect of deformed boundaries on grid cells?](#)
 Marino Krstulovic, Prannoy Chaudhuri-Vayalambone*, Marius Bauza, Michael Rule, Pauline Kerekes, Timothy O'Leary, Julija Krupic, University of Cambridge 94

I-75. [flexible recruitment of memory-based choice representations by human medial-frontal cortex](#)
 Juri Minxha, Ralph Adolphs, Stefano Fusi, Adam Mamelak, Ueli Rutishauser, Cedars Sinai Medical Center 94

I-76. [Closing the loop: recurrent neural networks that interact with their environment through reward-modulated learning](#)
 Robert Rosenbaum, Ryan Pyle, University of Notre Dame 95

I-77. [Which comes first: correlated pre- and post-synaptic activity or associative plasticity?](#)
 Aaron Milstein, Katie Bittner, Yiding Li, Christine Grienberger, Ivan Soltesz, Jeffrey Magee, Sandro Romani, Stanford University 96

I-78. [Angular and linear velocity coding in the parahippocampal circuit](#)
 Davide Spalla, Alessandro Treves, Charlotte Boccara, SISSA– Scuola Internazionale Superiore di Studi Avanzati – 96

I-79. [Neural networks for sorting neurons](#)
 JinHyung Lee, Catalin Mitelut, Ian Kinsella, Shenghao Wu, Eleanor Batty, Liam Paninski, Hooshmand Shokri, Ari Pakman, Yueqi Wang, Nishchal Dethe, Kevin Li, Eduardo Blancas Reyes, Alexandra Tikidji-Hamburyan, Georges Goetz, E.J. Chichilnisky, David Carlson, Columbia University 97

I-80. Strongly correlated spatiotemporal encoding and simple decoding in the prefrontal cortex Ehud Karpas, Roozbeh Kiani, Ori Maoz, Elad Schneidman, Weizmann Institute of Science	98
I-81. Measuring human probabilistic segmentation maps Jonathan Vacher, Pascal Mamassian, Ruben Coen-Cagli, Albert Einstein College of Medicine	98
I-82. Task engagement selectively enhances population discrimination of behavior-relevant categories in primary auditory cortex Charles Heller, Daniela Saderi, Mateo Lopez Espejo, Stephen David, Oregon Health and Science University	99
I-83. Non-normal amplification in the barrel cortex connectome Gadi Naveh, Haim Sompolinsky, Hebrew University of Jerusalem	99
I-84. Discovering interpretable models of population dynamics from neural activity recordings Mikhail Genkin, Tatiana Engel, Cold Spring Harbor Laboratory	100
I-85. Heterogeneity of timescales in recurrent networks with clustered connectivity Nicolae Istrate, Merav Stern, Luca Mazzucato, University of Oregon	100
I-86. Premotor output to striatum carries action history information to support ongoing decision-making Drew Schreiner, Christina Gremel, University of California, San Diego	101
I-87. Localized and sequential encoding of linguistic hierarchy in the human auditory cortex Menoua Keshishian, Jose Herrero, Nima Mesgarani, Bahar Khalighinejad, Hassan Akbari, Ashesh Mehta, Columbia University	102
I-88. Inferring neural population spiking rate from wide-field calcium imaging Merav Stern, Eric Shea-Brown, Daniela Witten, University of Washington	102
I-89. Pattern completion and separation in the interplay between olfactory bulb and cortex Gaia Tavoni, David Kersen, Vijay Balasubramanian, University of Pennsylvania	103
I-90. Nonlinear network activity regimes as a result of nonlinear neuronal dynamics Pierre Ekelmans, Nataliya Kravnyukova, Tatjana Tchumatchenko, Max Planck Institute for Brain Research	103
I-91. Mice detect increments of V1 spiking but not equivalent decrements Jackson Cone, Morgan Bade, Nicolas Masse, Elizabeth Page, David Freedman, John Maunsell, The University of Chicago	104
I-92. Disambiguating memory from contrast in monkey inferotemporal cortex Vahid Mehrpour, Travis Meyer, Eero Simoncelli, Nicole Rust, University of Pennsylvania	104
I-93. VTA stimulation induces remapping in multiple hippocampal regions Dorgham Khatib, Gilad Tocker, Jonathan Gross, Genela Morris, Dori Derdikman, Technion	105
I-94. Sex-modulated norepinephrine function in mediating exploration-exploitation tradeoff Cathy Chen, Becket Ebitz, Sylvia Bindas, Evan Knep, Nicola Grissom, University of Minnesota	106
I-95. Oxytocin enables social transmission of maternal behavior Ioana Carcea, Robert Froemke, Rutgers University	106
I-96. Solutions to the assignment problem balance tradeoffs between local and catastrophic errors W Jeffrey Johnston, David Freedman, The University of Chicago	107
I-97. Non-linear inhibitory plasticity can homeostatically regulate excitatory weight dynamics Christoph Miehl, Julijana Gjorgjieva, Max Planck Institute for Brain Research	107
I-98. Hierarchical inference guides perceptual decision-making in a dynamic environment Julie Charlton, Wiktor Mlynarski, Yoon Bai, Ann Hermundstad, Robbe Goris, University of Texas at Austin	108
I-99. Unsupervised depth estimation via visual-motor integration in recurrent neural networks Elliott Abe, Philip Parker, Cristopher Niell, Yashar Ahmadian, University of Oregon	108
I-100. Signatures of criticality observed in efficient coding networks Anna Levina, Shervin Safavi, Nikos Logothetis, Matthew Chalk, University of Tubingen	109
I-101. Learning structure from the environment using a recurrent spiking network Amadeus Maes, Mauricio Barahona, Claudia Clopath, Imperial College London	109

I-102. [Structure-function relationships in neural networks from geometric features of error landscapes](#)
Tirthabir Biswas, James fitzgerald, HHMI Janelia Research Campus 110

I-103. [Complementary encoding pathways build a memory hierarchy in a model of hippocampus](#)
Louis Kang, Taro Toyozumi, University of California, Berkeley 110

I-104. [A GLM approach for characterizing orbitofrontal cortical responses to multiple reward attributes](#)
David Hocker, Cristina Savin, Christine Constantinople, New York University 111

I-105. [The liminal state: Neural signatures of perceptual awareness](#)
Meenakshi Asokan, Kenneth Hancock, Ross Williamson, Daniel Polley, Harvard University 111

I-106. [Using task-optimized neural networks to understand why brains have specialized processing for faces](#)
Katharina Dobs, Alex Kell, Julio Martinez, Michael Cohen, Nancy Kanwisher, Massachusetts Institute of Technology 112

I-107. [A model for oscillatory gating of information flow between neural circuits as a function of local recurrence](#)
Mikhail Khona, Qianli Xu, Ila Fiete, Massachusetts Institute of Technology 113

I-108. [Unraveling the neural circuitry of predictive coding in mouse visual cortex](#)
Hannah Choi, Marina Garrett, Rylan Larsen, Nicholas Cain, University of Washington 113

I-109. [Action context influences a sensorimotor transformation across cortical areas](#)
Yi-Ting Chang, Eric A finkle, Jae Hun Lee, Emily E Lubin, Daniel O'Connor, Johns Hopkins University . . 114

I-110. [Transient working memory is controlled by slow synaptic dynamics](#)
Sophia Becker, Tatjana Tchumatchenko, Andreas Nold, Max Planck Institute for Brain Research 114

I-111. [fine-scale organization of the functional diversity in mouse superior colliculus](#)
Hymavathy Balasubramanian, Jeremie Sibille, Jens Kremkow, Bernstein Centre for Computational Neuroscience 115

I-112. [Hippocampal population representations in reinforced continual learning](#)
Samia Mohinta, Stephane Ciocchi, Rui Ponte Costa, University of Bristol 116

I-113. [Where can a place cell put its fields? Let us count the ways](#)
Thibaud Tallefumier, Man Yi Yim, Lorenzo Sadun, Ila fiete, University of Texas at Austin 116

I-114. [Deep neuroethology of a virtual rodent](#)
Josh Merel, Diego Aldarondo, Jesse Marshall, Yuval Tassa, Greg Wayne, Bence Olveczky, DeepMind . . 117

I-115. [Emergence of functional and structural properties of head direction system by optimization of RNNs](#)
Christopher Cueva, Peter Wang, Matthew Chin, Xue-Xin Wei, Columbia University 117

I-116. [State space oscillator models to identify and parameterize oscillatory signals in EEG](#)
Amanda Beck, Hugo Soulat, Emily Stephen, Patrick Purdon, Massachusetts Institute of Technology . . . 118

I-117. [Dissection of low-dimensional models of V1 circuitry with multiple inhibitory types](#)
Kenneth Miller, Agostina Palmigiano, Nataliya Krainyukova, Daniel Mossing, Hillel Adesnik, Tatjana Tchumatchenko, Columbia University 118

I-118. [Predicting optogenetic responses through full-dimensional models of interneuronal circuitry](#)
Agostina Palmigiano, Daniel Mossing, Hillel Adesnik, Kenneth Miller, Francesco Fumarola, Columbia University 119

I-119. [Recurrent neural circuits overcome partial inactivation by compensation and re-learning](#)
Colin Bredenber, Roozbeh Kiani, Cristina Savin, New York University 120

I-120. [Multiscale low-dimensional neural dynamics explain naturalistic 3D movements](#)
Hamidreza Abbaspourazad, Mahdi Choudhary, Yan Wong, Bijan Pesaran, Maryam Shanechi, University of Southern California 120

I-121. [The PHATE of learning networks](#)
Scott Gigante, Adam Charles, Smita Krishnaswamy, Gal Mishne, Yale University 121

I-122. [Learning and stabilizing manifold attractors in heterogeneous networks](#)
Alexander Rivkind, Ran Darshan, The Weizmann Institute of Science 121

I-123. [Spatiotemporal neural correlations and network dynamics](#)
Yanliang Shi, Roxana Zeraati, Anna Levina, Tatiana Engel, Cold Spring Harbor Laboratory 122

I-124. [Neural distribution alignment via hierarchical optimal transport](#)
Max Dabagia, Eva Dyer, Georgia Institute of Technology 122

II-1. Optimal population response for robust linear readout Blake Bordelon, Cengiz Pehlevan, Harvard University	123
II-2. Slow coordinated activity across hemispheres in PFC predicts pupil size Akash Umakantha, Ryan Williamson, Byron Yu, Matthew Smith, Carnegie Mellon University	123
II-3. Mixed representations in retrosplenial cortex mediate flexible navigation decisions in mice Shinichiro Kira, Giuseppe Pica, Stefano Panzeri, Christopher Harvey, Harvard University	124
II-4. Blockage of cerebellar outflow changes the spatiotemporal organization of muscle activity Nirvik Sinha, Abdurraheem Nashef, Sharon Israely, Yifat Prut, Northwestern University	124
II-5. Inferring function from structure with connectome and task constrained neural networks Janne Lappalainen, Sridhama Prakhya, Mason McGill, Fabian Tschopp, Srini Turaga, HHMI Janelia Research Campus	125
II-6. Representation of uncertainty in macaque visual cortex Zoe Boundy-Singer, Olivier Henaff, Kristof Meding, Corey Ziemba, Robbe Goris, University of Texas at Austin	126
II-7. Theory of gating in recurrent neural networks Kamesh Krishnamurthy, Tankut Can, David Schwab, Princeton University	126
II-8. Dynamical mechanisms of flexible timing by temporal scaling Manuel Beiran, Mehrdad Jazayeri, Srdjan Ostojic, Ecole Normale Supérieure	127
II-9. A learning-forgetting tradeoff from the perspective of control: a Drosophila case study James Bennett, Luca Manneschi, Eleni Vasilaki, Thomas Nowotny, University of Sussex	127
II-10. Sequential and efficient neural-population coding of complex task information Sue Ann Koay, Stephan Thiberge, Carlos D Brody, David Tank, Princeton University	128
II-11. Convolutional dictionary learning of stimulus from spiking data Andrew Song, Bahareh Tolooshams, Simona Temereanca, Demba Ba, Massachusetts Institute of Technology	128
II-12. Reinforcement learning models reveal quantitative reward prediction errors in ventral pallidum David Ottenheimer, Bilal Bari, Elissa Sutlief, Kurt Fraser, Tabitha Kim, Jocelyn Richard, Jeremiah Cohen, Patricia Janak, Johns Hopkins University	129
II-13. Discrete latent states underlie sensory decision-making behavior in mice Zoe Ashwood, Anne Urai, Valeria Aguilon, International Brain Laboratory, Anne Churchland, Alexandre Pouget, Jonathan Pillow, Princeton University	130
II-14. Human visual motion perception shows hallmarks of Bayesian structural inference Sichao Yang, Johannes Bill, Jan Drugowitsch, Samuel J Gershman, University of Wisconsin-Madison	130
II-15. Mechanistic model of a neural circuit for ZCA-type whitening in olfaction Nikolai Chapochnikov, Cengiz Pehlevan, Dmitri Chklovskii, flatiron Institute	131
II-16. Meta-learning biologically plausible semi-supervised learning rules Niru Maheswaranathan, Keren Gu-Lamberg, Sam Greydanus, Luke Metz, Jascha Sohl-Dickstein, Google Brain	132
II-17. Revealing human brain dynamics underlying ECoG power via Directed Information in frequency Sudha Yellapantula, Kiefer Forseth, Nitin Tandon, Behnaam Aazhang, Rice University	132
II-18. Probing song motor circuit architecture and function in Drosophila melanogaster Frederic Roemschied, Diego Pacheco Pinedo, Xiping Li, Elise Ireland, Mala Murthy, Princeton University	133
II-19. Attractor dynamics gate cortical information flow during decision-making Lorenzo Fontolan, Arseny finkelstein, Michael Economo, Nuo Li, Sandro Romani, Karel Svoboda, HHMI Janelia Research Campus	133

[II-20. Representing relational knowledge in cognitive maps for generalization](#)
 Mona Garvert, Nicolas Schuck, Christian Doeller, Max Planck Institute for Human Cognitive and Brain Sciences 134

[II-21. Neural basis of active odour trail tracking in rodents](#)
 Siddharth Jayakumar, Alexander Mathis, Ventkatesh Murthy, Harvard University 134

[II-22. EASE: EM-Assisted Source Extraction from calcium imaging data](#)
 Pengcheng Zhou, Jacob Reimer, Ding Zhou, Amol Pasarkar, Ian Kinsella, Emmanouil Froudarakis, Dimitri Yatsenko, Paul Fahey, Agnes Bodor, JoAnn Buchanan, Dan Bumbarger, Gayathri Mahalingam, Russel Torres, Sven Dorkenwald, Dodam Ih, Kisuk Lee, Ran Lu, Thomas Macrina, Jingpeng Wu, Nuno Costa, Clay Reid, Andreas Tolias, Liam Paninski, Columbia University 135

[II-23. Inferring network motifs from neural activity using analytic input-output relation of LIF neurons](#)
 Safura Rashid Shomali, Majid Nili Ahmadabadi, Seyyed Nader Rasuli, Hideaki Shimazaki, Institute for Research in Fundamental Sciences (IPM) 136

[II-24. Dynamics of natural odor plumes structure the activity of neural populations at the first stage of odor processing in mice](#)
 Suzanne Lewis, Lai Xu, Mohammad Tariq, Venkatesh Gopal, Agnese Seminara, Merav Stern, David Gire, University of Washington 137

[II-25. Activation of distinct cortical circuitries by expected and unexpected stimuli](#)
 Farzaneh Najafi, Natalia Orlova, Dmitri Tsybouski, Sam Seid, India Kato, Sara Kivikas, Marina Garrett, Peter Groblewski, Rylan Larsen, Doug Ollerenshaw, David Sullivan, Wayne Wakeman, Quinn L'Heureux, Shawn Olsen, Jerome Lecoq, Allen Institute for Brain Science 138

[II-26. Precise excitation-inhibition balance - computational implications for feedforward circuits](#)
 Aanchal Bhatia, Sahil Moza, Upinder Bhalla, National Centre for Biological Sciences 138

[II-27. Reward expectancy restructures and enhances spatial encoding in the hippocampus](#)
 Seetha Krishnan, Chery Cherian, Mark Sheffield, University of Chicago 139

[II-28. Contribution of sensory-motor gain to signal and noise in motor control](#)
 Seth Egger, Stephen Lisberger, Duke University School of Medicine 140

[II-29. Hippocampal replay slows down with experience, as greater detail is incorporated](#)
 Alice Berners-Lee, Ting Feng, Delia Silva, Xiaojing Wu, Ellen Ambrose, Brad Pfeiffer, David Foster, Johns Hopkins University 140

[II-30. Cellular and synaptic mechanisms of familiarity-evoked theta oscillations in the visual cortex](#)
 Alexander Chubykin, Mang Gao, Purdue University 141

[II-31. The dentate gyrus classifies cortical representations of learned stimuli](#)
 Fabio Stefanini, Nicholas Woods, Daniel Apodaca-Montano, Jeremy S Biane, Mazen Kheirbek, Columbia University 141

[II-32. Accurate extraction of membrane potential for in vivo voltage imaging](#)
 Michael Xie, Yoav Adam, Linlin Fan, Urs Boehm, Ian Kinsella, Ding Zhou, Liam Paninski, Adam Cohen, Harvard University 142

[II-33. Anterior and posterior striatum play distinct roles in evidence accumulation](#)
 Adrian Bondy, Diksha Gupta, Carlos D Brody, Thomas Luo, Princeton University 142

[II-34. Hebbian STDP protects against assembly fusion in stochastic networks](#)
 Xinrui Yang, Brent Doiron, University of Pittsburgh 143

[II-35. The role of hippocampal sequences in deliberative control: a computational and algorithmic theory](#)
 David Theurel, Matthew Wilson, Zhe Chen, Massachusetts Institute of Technology 143

[II-36. Adaptive integration behavior from static integration windows](#)
 Richard McWalter, Josh McDermott, Massachusetts Institute of Technology 144

[II-37. Resonating neurons stabilize heterogeneous grid-cell networks](#)
 Divyansh Mittal, Rishikesh Narayanan, Indian Institute of Science, Bangalore 145

Posters II

II-38. Non-perturbative renormalization group analysis of strongly-coupled spiking networks	
Braden Brinkman, Stony Brook University	145
II-39. Interneuron diversity emerges in spiking networks optimized for a compartment-specific excitation-inhibition balance	
Joram Keijser, Sprekeler Henning, TU Berlin	146
II-40. The influence of noise on the efficient coding of natural scenes	
Na Young Jun, Greg field, John Pearson, Duke University	146
II-41. Visual pattern analysis by motor neurons	
Antimo Buonocore, Matthias Baumann, Ziad Hafed, University of Tuebingen	147
II-42. 3-factor Hebbian learning rules in deep networks: an information bottleneck approach	
Roman Pogodin, Peter Latham, Gatsby Computational Neuroscience Unit, University College London	147
II-43. Temporal basis decomposition reveals a linear relationship between neural activity and kinematics	
David Sabatini, Matthew Kaufman, The University of Chicago	148
II-44. Experience dependent context discrimination in the hippocampus	
Mark Plitt, Lisa Giocomo, Stanford University	148
II-45. Parallel emergence of mixed tuning to task variables and cell assembly coordination in mPFC during learning	
Aleksander Domanski, Daniel Durstewitz, Matt Jones, University of Bristol	149
II-46. The interplay between randomness and structure during learning in RNNs	
Friedrich Schuessler, Alexis Dubreuil, Francesca Mastrogiuseppe, Srdjan Ostojic, Omri Barak, Technion	149
II-47. Meta-learning Hebbian plasticity for continual familiarity detection	
Danil Tjulmankov, Guangyu Robert Yang, Larry Abbott, Columbia University	150
II-48. Uncovering the organization of neural circuits with generalized phase-locking analysis	
Shervin Safavi, Fanis Panagiotaropoulos, Vishal Kapoor, Juan F Ramirez-Villegas, Nikos Logothetis, Michel Besserve, Max Planck Institute for Biological Cybernetics	150
II-49. Impact of single neurons in cortical circuits with dense and strong connectivity	
Juan Luis Riquelme, Mike Hemberger, Gilles Laurent, Julijana Gjorgjieva, Max Planck Institute for Brain Research	151
II-50. Biologically plausible supervised learning via similarity matching	
Shanshan Qin, Cengiz Pehlevan, Harvard University	152
II-51. Disentangling the roles of dimensionality and cell classes in neural computations	
Alexis Dubreuil, Adrian Valente, Francesca Mastrogiuseppe, Srdjan Ostojic, Ecole Normale Supérieure	152
II-52. Neuromechanical modeling of Hydra behaviors	
Hengji Wang, Adrienne Fairhall, University of Washington	153
II-53. Serotonin neuron firing rates track expected uncertainty during dynamic decision making	
Cooper Grossman, Bilal Bari, Jeremiah Cohen, Johns Hopkins University	153
II-54. SMA-ACC network guides flexible strategy selection in rats	
Mikhail Proskurin, Maxim Manakov, Maria Rysakova, Elena Kuleshova, Andy Lustig, Reza Behnam, Dougal Tervo, Alla Karpova, HHMI Janelia Research Campus	154
II-55. Hierarchical Markov modeling of social behavior from 3D mouse pose	
Caleb Weinreb, Sandeep Robert Datta, Alia Newman-Boulle, Scott W Linderman, Harvard University	154
II-56. Predictive coding explains the link between complex cells and the topographic organization of orientation maps in V1	
Angelo Franciosini, Victor Boutin, Frederic Chavane, Laurent Perrinet, Aix Marseille Université	155
II-57. Optogenetically-evoked modulation of noise correlations in the visual cortex	
Alan Akil, Ariana Andrei, Robert Rosenbaum, Sorin Pojoga, Roger Janz, Valentin Dragoi, Kresimir Josic, University of Houston	156

[II-58. Continual weight updates and convolutional architectures for equilibrium propagation](#)
 Maxence Ernout, Yoshua Bengio, Julie Grollier, Damien Querlioz, Benjamin Scellier, Universite Paris Sud 156

[II-59. Long-range horizontal connectivity in mammalian visual cortex for balancing performance and cost](#)
 Youngjin Park, Seungdae Baek, Se-Bum Paik, KAIST 157

[II-60. Efficient and sufficient: generalized optimal codes](#)
 Luke Rast, Jan Drugowitsch, Harvard University 158

[II-61. Spatio-temporal structure of awake hippocampal replay](#)
 Emma Krause, Jan Drugowitsch, Harvard University 158

[II-62. Emergent evolutionary dynamics over recurrent neural computations](#)
 Czege Daniel, Hamza Giaffar, Marton Csillag, Eors Szathmary, Eotvos University 159

[II-63. Fast and deep neuromorphic learning with first-spike coding](#)
 Julian Goltz, Andreas Baumbach, Sebastian Billadelle, Oliver Breitwieser, Laura Kriener, Akos Ferenc Kungl, Karlheinz Meier, Johannes Schemmel, Mihai A Petrovici, Heidelberg University 159

[II-64. Goal-directed state space models of mouse behavior](#)
 Blue Sheffer, Sandeep Robert Datta, Scott Linderman, Stanford University 160

[II-65. Differential covariance of fMRI predicts structural connectivity and behavior](#)
 Yusi Chen, Qasim Bukhari, Tiger Wuton Lin, Terrence Sejnowski, Salk Institute for Biological Studies . . . 160

[II-66. Evolving the olfactory systems](#)
 Guangyu Robert Yang, Peter Wang, Richard Axel, Larry Abbott, Yi Sun, Ashok Litwin-Kumar, Columbia University 161

[II-67. Time continuity of early visual experience drives the development of ventral stream representations](#)
 Eis Annavini, Mattia D'Andola, Alessio Ansuini, Davide Zoccolan, SISSA– Scuola Internazionale Superiore di Studi Avanzati – 161

[II-68. A central source for fixational eye drifts](#)
 Nadav Ben-Shushan, Nimrod Shaham, Mati Joshua, Yoram Burak, The Hebrew University of Jerusalem . 162

[II-69. Inhibitory output specificity and the speed-accuracy tradeoff in perceptual decision-making](#)
 James Roach, Anne Churchland, Tatiana Engel, Cold Spring Harbor Laboratory 162

[II-70. Identifying and mitigating statistical biases in neural models of tuning and functional coupling](#)
 Pratik Sachdeva, Jesse Livezey, Maximilian Dougherty, Sharmodeep Bhattacharyya, Kristofer Bouchard, University of California, Berkeley 163

[II-71. Multi-dimensional perceptual decision making under dynamic belief states](#)
 Cheng Xue, Lily Kramer, Marlene Cohen, University of Pittsburgh 163

[II-72. Precise coupling of the thalamic head-direction system to hippocampal ripples](#)
 Guillaume Viejo, Adrien Peyrache, McGill University 164

[II-73. Kura-Net: Exploring systems of coupled oscillators with deep learning](#)
 Matthew Ricci, Aneri Soni, Thomas Serre, Yuwei Zhang, Minju Jung, Brown University 164

[II-74. Synergistic population coding through correlated variability in V1](#)
 Daniel Denman, Clay Reid, University of Colorado 165

[II-75. Variational Autoencoder account of the early visual hierarchy](#)
 Gergo Orban, Mihaly Banyai, David Nagy, Wigner RCP 165

[II-76. For bursts with plateau potentials, GCaMP6-signal reflects plateau duration rather than burst size](#)
 Peter Ledochowitsch, Lawrence Huang, Gabriel Ocker, Michael Oliver, Jack Waters, Hongkui Zeng, Saskia de Vries, Michael Buice, Allen Institute for Brain Science 166

[II-77. Reorganization of phase space structure during network learning](#)
 Rainer Engelken, Larry Abbott, Columbia University 167

[II-78. Meta-learning biologically plausible learning algorithms with feedback and local plasticity](#)
 Jack Lindsey, Ashok Litwin-Kumar, Columbia University 167

Posters II

II-79. Biophysics-based frequency domain spatiotemporal pattern separation of the multichannel lfp signals	
Weiwei Chen, Anton Sirota, Ludwig-Maximilians University Munich	168
II-80. Combining MRI connectomics and reservoir computing to quantify the coding properties of large-scale brain networks	
Estefany Suarez, Guillaume Lajoie, Bratislav Misis, McGill University	168
II-81. Experimentally observed features of sequential activity with temporally asymmetric Hebbian learning	
Maxwell Gillett, Nicolas Brunel, Duke University	169
II-82. Single neurons in rat orbitofrontal cortex predict confidence-guided time investment and choice updating	
Torben Ott, Paul Masset, Armin Lak, Junya Hirokawa, Adam Kepecs, Cold Spring Harbor Laboratory . . .	169
II-83. Modeling behaviorally relevant neural dynamics with a novel preferential subspace identification (PSID)	
Omid Sani, Bijan Pesaran, Maryam Shanechi, University of Southern California	170
II-84. Normalized reinforcement learning predicts asymmetric and adaptive choice behavior	
Kenway Louie, New York University	171
II-85. Modeling the population-level statistics and structure of hippocampal place cell responses	
Man Yi Yim, Jae Sung Lee, Albert Lee, Thibaud Taillefumier, Ila fiete, University of Texas at Austin	171
II-86. Predictive coding via optimal sampling regulates processing of odor dynamics	
Morgan Brown, Yashar Ahmadian, Roma Shusterman, University of Oregon	172
II-87. Minimax dynamics of optimal excitatory-inhibitory balanced networks	
Qianyi Li, Cengiz Pehlevan, Harvard University	172
II-88. Compressed interpretable models of neurons in cortical area V4	
Reza Abbasi-Asl, Bin Yu, University of California, San Francisco	173
II-89. Balanced spiking networks with conditionally Poisson neurons	
Camille Rullan, Jonathan Pillow, New York University	173
II-90. Motor signals promote rapid basis set formation in model cerebellar granule cell layer	
Jesse Gilmer, Ioannis Delis, Abigail Person, University of Colorado	174
II-91. Functional relevance of sensory- and task-related representations in posterior parietal cortex	
Matthijs oude Lohuis, Pietro Marchesi, Umberto Olcese, Cyriel Pennartz, University of Amsterdam	174
II-92. Controlling processing speed via cell-type specific perturbations in spiking neural circuits	
David Wyrick, Luca Mazzucato, University of Oregon	175
II-93. A novel computational framework for the role of nucleus accumbens dopamine in information processing	
Munir Kutlu, Jennifer Zachry, Patrick Melugin, Erin Calipari, Vanderbilt University School of Medicine . . .	175
II-94. Task-selective connectivity improves rapid task switching in an interpretable model of SC	
Sean Bittner, Carlos D Brody, John Cunningham, Alex Piet, Chunyu Duan, Columbia University	176
II-95. An anatomically and functionally constrained model of direction selectivity in Drosophila	
Jacob Portes, Jessica Kohn, Larry Abbott, Rudy Behnia, Columbia University	176
II-96. LOOPER: A tool for the semi-supervised extraction of behaviorally relevant dynamics from observations of neural data	
Connor Brennan, Alex Proekt, University of Pennsylvania	177
II-97. Normative evidence-accumulation and decision-making in foraging	
Zachary Kilpatrick, Jacob Davidson, Ahmed El Hady, University of Colorado, Boulder	177
II-98. Learning sequences of correlated patterns in recurrent networks	
Subhadra Mokashe, Nicolas Brunel, Duke University	178

[II-99. Dynamic long-range coordination patterns in macaque motor cortex](#)
David Dahmen, Moritz Layer, Lukas Deutz, Paulina Anna Dąbrowska, Nicole Voges, Michael von Papen, Sonja Grun, Markus Diesmann, Moritz Helias, Research Centre Julich 179

[II-100. Higher visual areas contribute to visually guided detection behavior in the mouse](#)
Hannah Goldbach, Bradley Akitake, Caitlin Leedy, Mark H Histed, National Inst. of Mental Health 179

[II-101. Neural dynamics across brains of socially interacting animals](#)
Lyle Kingsbury, Shan Huang, Jun Wang, Ken Gu, Peyman Golshani, Emily Wu, Weizhe Hong, University of California, Los Angeles 180

[II-102. Transient chaotic dimensionality expansion by recurrent networks](#)
Christian Keup, Tobias Kuhn, David Dahmen, Moritz Helias, Research Centre Julich 180

[II-103. Foveal action for the control of extra-foveal vision](#)
Ziad Hafed, Xiaoguang Tian, Tong Zhang, University of Tuebingen 181

[II-104. Using noise to probe network structure and prune synapses](#)
Rishidev Chaudhuri, University of California, Davis 181

[II-105. Automatic fitting of spiking network models to neuronal activity reveals limits of model flexibility](#)
Shenghao Wu, Chengcheng Huang, Adam Snyder, Matthew Smith, Brent Doiron, Byron Yu, Carnegie Mellon University 182

[II-106. Development of negative prediction-error neurons in the presence of three inhibitory interneuron types](#)
Loreen Hertaeg, Sprekeler Henning, TU Berlin 183

[II-107. Hexadirectional coding of decision trajectories through abstract and discrete spaces](#)
Seongmin Park, Douglas Miller, Erie Boorman, University of California, Davis 183

[II-108. Diverse behavioral strategies and neural coding during a visual change detection task](#)
Alex Piet, Nick Ponvert, Marina Garrett, Anton Arkhipov, Yazan Billeh, Allen Institute for Brain Science . . . 184

[II-109. An inference perspective on urgency in decision-making: A drunkard's walk case study](#)
Maximilian Puelma Touzel, Raymond Chua, Yoshua Bengio, Paul Cisek, Guillaume Lajoie, Universite de Montreal 184

[II-110. Excitatory and inhibitory tone shapes decision regimes in hierarchical neural networks to maximize reward across environments](#)
Warren Pettine, Kenway Louie, John D Murray, Xiao-Jing Wang, New York University 185

[II-111. A cognitive map in the primate entorhinal cortex for mental navigation](#)
Sujaya Neupane, Mehrdad Jazayeri, Massachusetts Institute of Technology 185

[II-112. Economic choices with simultaneous or sequential offers rely on the same neural circuit mechanism](#)
Weikang Shi, Sebastien Ballesta, Camillo Padoa-Schioppa, Washington University in St Louis 186

[II-113. Suboptimal evidence integration reveals adaptive, nonlinear mechanisms in spatial vision](#)
Takahiro Doi, Johannes Burge, University of Pennsylvania 186

[II-114. Time-dependent decision urgency in dynamic reward and evidence-quality conditions](#)
Nicholas Barendregt, Kresimir Josic, Zachary Kilpatrick, University of Colorado, Boulder 187

[II-115. From single neurons to behavior in the jellyfish *Aurelia aurita*](#)
Fabian Pallasdies, Sven Goedeke, Wilhelm Braun, Raoul-Martin Memmesheimer, University of Bonn . . . 188

[II-116. Mixture of poisson models for neural correlations and coding](#)
Sacha Sokoloski, Amir Aschner, Ruben Coen-Cagli, Albert Einstein College of Medicine 188

[II-117. Understanding encoding and redundancy in grid cells using partial information decomposition](#)
Praveen Venkatesh, Pulkit Grover, Carnegie Mellon University 189

[II-118. Isolating attention components in superior colliculus: perceptual sensitivity, criterion and motor bias](#)
Supriya Ghosh, John Maunsell, University of Chicago 189

Posters II

II-119. Formalizing a perceptual-mnemonic theory of the medial temporal lobe	
Tyler Bonnen, Dan Yamins, Anthony Wagner, Stanford University	190
II-120. The effect of boundaries on grid cell patterns	
Mauro Miguel Monsalve Mercado, Christian Leibold, Columbia University	190
II-121. Optimal dendritic processing of spatiotemporal codes	
Brendan Bicknell, Michael Hausser, University College London	191
II-122. Dimensionality control in the critical regime of balanced networks	
Eric Shea-Brown, Stefano Recanatesi, David Dahmen, Moritz Helias, University of Washington	191
II-123. Encoding object transformations across distinct areas of rodent visual cortex	
Juliana Rhee, Cesar Echavarria, Javier Masis, Davidf Cox, Harvard University	192
II-124. The sensorimotor strategies and neuronal representations of whisker-based shape recognition in mice	
Chris Rodgers, Ramon Nogueira, Esther Greeman, Stefano Fusi, Randy Bruno, Columbia University . . .	192

Poster Session III

8:00p Saturday 29 February

III-1. Motor cortical representation and decoding of attempted handwriting in a person with tetraplegia Frank Willett, Donald Avansino, Leigh Hochberg, Jaimie Henderson, Krishna Shenoy, Stanford University	193
III-2. Reinforcement regulates context-dependent timing variability in thalamus Jing Wang, Eghbal Hosseini, Nicolas Meirhaeghe, Adam Akkad, Mehrdad Jazayeri, Massachusetts Institute of Technology	194
III-3. Dynamic recombination of temporally coincident presynaptic modules underlies responses in cortex Arjun Bharioke, Stuart Trenholm, Adrian Wertz, Daniel Hillier, Emilie Mace, Keisuke Yonehara, Balazs Rozsa, Karl-Klaus Conzelmann, Rava Azeredo da Silveira, Botond Roska, Institute of Ophthalmology Basel	194
III-4. Linking structure to function in a model of early color processing Matthias Christenson, Sarah Heath, Maia Weisenhaus, Larry Abbott, Rudy Behnia, Columbia University	195
III-5. Internal models for active sensing in freely swimming electric fish Avner Wallach, Nate Sawtell, Columbia University	196
III-6. Multi-animal pose tracking and predictive modeling of social interactions Talmo Pereira, Mala Murthy, Nathaniel Tabris, Shruthi Ravindranath, Junyu Li, Joshua Shaevitz, Princeton University	196
III-7. The functional architecture of glutamatergic release onto CA1 neurons Jason Climer, Michael Adoff, Heydar Davoudi, Daniel Dombeck, Northwestern University	197
III-8. Tongue kinematics reveal cortex-dependent corrections as the mouse tongue reaches for, and misses, targets Teja Bollu, Brendan Ito, Sam Whitehead, Brian Kardon, Mei Liu, James Redd, Jesse Goldberg, Cornell University	197
III-9. Precise optical probing of perceptual detection in olfactory circuits Jonathan Gill, Gilad Lerman, Hetince Zhao, Benjamin Stetler, Shy Shoham, Dmitry Rinberg, New York University	198
III-10. Neural mechanisms of selection in visual working memory Matthew Panichello, Timothy Buschman, Princeton University	198
III-11. Multiple overlapping hypothalamus-brainstem circuits drive rapid threat avoidance Matthew Lovett-Barron, Ritchie Chen, Susanna Bradbury, Aaron S Andalman, Mahendra Wagle, Su Guo, Karl Deisseroth, Stanford University	199
III-12. Action selection in continuous time and space in the basal ganglia Sean Escola, James Murray, Columbia University	199
III-13. The role of TRN for feedback-enhanced surround suppression in mouse dLGN Gregory Born, Sinem Erisken, Felix Schneider, Agne Klein, Martin Spacek, Milad Mobarhan, Davide Crombie, Chu Lan Lao, Gaute Einevoll, Laura Busse, Ludwig-Maximilians-University Munich	200
III-14. A theory of learning with a constrained weight distribution Ben Sorscher, Weishun Zhong, Daniel D Lee, Haim Sompolinsky, Stanford University	200
III-15. Latent-state models reveal a state-dependent contribution of the striatum to decision-making Iris Stone, Scott Bolkan, Ilana Witten, Jonathan Pillow, Princeton University	201
III-16. Remembrance of things practiced: A two-pathway circuit for sequential learning James Murray, Sean Escola, Columbia University	202
III-17. Local and global organization of synaptic inputs on cortical dendrites Jan Hendrik Kirchner, Julijana Gjorgjieva, Max Planck Institute for Brain Research	202
III-18. Probing the neural substrates of movement across the rodent behavioral repertoire Jesse Marshall, Diego Aldarondo, William Wang, Tim Dunn, Gordon Berman, Bence Olveczky, Harvard University	203

Posters III

III-19. flygenectors: large-scale dynamics of internal and behavioral states in a small animal Evan Schaffer, Neeli Mishra, Wenze Li, Matthew Whiteway, Jason Freedman, Kripa Patel, Venkatakaushik Voleti, Liam Paninski, Larry Abbott, Elizabeth Hillman, Richard Axel, Columbia University	203
III-20. Studying neural codes of odor perception using a novel behavioral paradigm in mice Hirofumi Nakayama, Dmitry Rinberg, New York University Neuroscience Institute	204
III-21. Developmental and evolutionary principles of olfactory circuit designs Naoki Hiratani, Peter Latham, Gatsby Computational Neuroscience Unit, University College London	205
III-22. Natural gradient learning for spiking neurons Elena Kreuzer, Mihai A Petrovici, Walter Martin Senn, University of Bern	205
III-23. Algorithmic basis of foraging decisions Sergey Shuvaev, Sarah Starosta, Duda Kvitsiani, Adam Kepecs, Alexei Koulakov, Cold Spring Harbor Laboratory	206
III-24. Efficient routing of information to high-dimensional neural representations Samuel Muscinelli, Ashok Litwin-Kumar, Columbia University	206
III-25. Expressivity of expand-and-sparsify representations Sanjoy Dasgupta, University of California, San Diego	207
III-26. Functional organization of posterior cortex during flexible navigation decisions Shih-Yi Tseng, Selmaan Chettih, Charlotte Arlt, Roberto Barroso-Luque, Christopher Harvey, Harvard University	207
III-27. Perceptual straightening of natural videos arises from a cascaded computation Robbe Goris, Yoon Bai, Olivier Henaff, Corey Ziemba, University of Texas at Austin	208
III-28. Strong disinhibition establishes long neuronal timescales critical for working memory Robert Kim, Terrence Sejnowski, University of California, San Diego	208
III-29. Network mechanisms of long-term spatial memory consolidation Andres Grosmark, Fraser Sparks, Matt Davis, Attila Losonczy, Columbia University	209
III-30. Structural and functional differences in the head direction circuit of two insect species Ioannis Pisokas, Stanley Heinze, Barbara Webb, University of Edinburgh	209
III-31. Neural and computational basis for organizing episodic experience into event units Chen Sun, Kelsey Allen, Susumu Tonegawa, Massachusetts Institute of Technology	210
III-32. Geometric representation of abstract learned knowledge by neural manifolds in hippocampus Edward H Nieh, Manuel Schottdorf, Nicolas W Freeman, Sue Ann Koay, Lucas Pinto, Jeffrey L Gauthier, Carlos D Brody, David Tank, Princeton University	211
III-33. Navigational uncertainty provides a unifying account of human navigational behavior and rodent grid field deformations Yul Kang, Daniel Wolpert, Mate Lengyel, University of Cambridge	211
III-34. Auditory activity is diverse and widespread throughout the central brain of drosophila Diego Pacheco Pinedo, Eftychios Pnevmatikakis, Stephan Thiberge, Mala Murthy, Princeton University	212
III-35. Representation of 3D orienting movements in primary visual cortex Grigori Guitchounts, Javier Masis, Steffen Wolff, David Cox, Harvard University	212
III-36. Understanding the emergence of hexagonal grid cells in networks trained to path integrate Gabriel Mel, Surya Ganguli, Ben Sorscher, Samuel A Ocko, Stanford University	213
III-37. Comparing physical inferences in humans, monkeys and task-optimized recurrent neural network models Rishi Rajalingham, Aida Piccato, Mehrdad Jazayeri, Massachusetts Institute of Technology	213
III-38. Modeling large-scale brain network dynamics in response to electrical stimulation Yuxiao Yang, Shaoyu Qiao, Omid Sani, Isaac Sedillo, Breonna Ferrentino, Bijan Pesaran, Maryam Shanechi, University of Southern California	214

III-39. Adaptive stimulus selection improves deep neural network models of macaque V4 Benjamin Cowley, Patricia Stan, Matthew Smith, Jonathan Pillow, Princeton University	215
III-40. Why do neurons spike spontaneously? Chaitanya Chintaluri, Tim Vogels, University of Oxford	215
III-41. A dynamical model with E/I balance explains robustness to optogenetic stimulation in motor cortex Lea Duncker, Daniel O'Shea, Krishna Shenoy, Maneesh Sahani, Gatsby Computational Neuroscience Unit, University College London	216
III-42. System identification in the brain: clustering dynamical rules from sensory data Tiberiu Tesileanu, Samaneh Nasiri, Anirvan Sengupta, Dmitri Chklovskii, flatiron Institute	216
III-43. Synaptic plasticity in balanced networks Kresimir Josic, Alan Akil, Robert Rosenbaum, University of Houston	217
III-44. Neural population dynamics in motor cortex during grasp Aneesha K Suresh, James Goodman, Elizaveta V Okorokova, Matthew Kaufman, Nicholas G Hatsopoulos, Sliman Bensmaia, The University of Chicago	217
III-45. Temporal context invariance reveals a timescale hierarchy in human auditory cortex Sam Norman-Haignere, Laura Long, Orrin Devinsky, Werner Doyle, Guy McKhann, Catherine Schevon, Adeen flinker, Nima Mesgarani, Columbia University	218
III-46. Deep neural networks reveal adaptedness of human sound localization to real-world environments Andrew Francl, Josh McDermott, Massachusetts Institute of Technology	218
III-47. Working memory for temporal estimates via slow neural manifolds Alexandra Ferguson, Manuel Beiran, T Vincenza Parks, Srdjan Ostojic, Mehrdad Jazayeri, Massachusetts Institute of Technology	219
III-48. Stimulus representation and inference by recurrently coupled neuronal populations Wenhao Zhang, Si Wu, Kresimir Josic, Brent Doiron, University of Pittsburgh	220
III-49. Continual learning, replay and consolidation in a forgetful memory network model Nimrod Shaham, Jay Chandra, Kreiman Gabriel, Haim Sompolinsky, Harvard University	220
III-50. A normative hippocampus model: robustly encoding variables on smooth manifolds using spiking neurons Paxon Frady, Friedrich Sommer, University of California, Berkeley	221
III-51. Phase transitions in vocal development are driven by energy-information balance Thiago Tarrat Varella, Yisi Zhang, Daniel Takahashi, Asif Ghazanfar, Princeton University	221
III-52. Laminar-specific sensorimotor integration of context in cortical dynamics Hansem Sohn, Mehrdad Jazayeri, Massachusetts Institute of Technology	222
III-53. Adjudicating between deep neural network models of biological vision with controversial stimuli Tal Golan, Prashant C Raju, Nikolaus Kriegeskorte, Columbia University	223
III-54. Surprise learning: fast adaptation in volatile environments with spiking neural networks Martin Louis Lucien Remy Barry, Carlos Stein N Brito, Wulfram Gerstner, EPfl	223
III-55. Efficient decoding of large-scale neural population responses with Gaussian-process multiclass regression C Daniel Greenidge, Jacob Yates, Benjamin Scholl, Jonathan Pillow, Princeton University	224
III-56. Recurrent neural network dynamics are dependent on learning rule Brandon McMahan, Michael Kleinman, Jonathan Kao, University of California, Los Angeles	224
III-57. Systematic tradeoffs between complexity and accuracy in human decision-making Alexandre filipowicz, Eugenio Piasini, Kamesh Krishnamurthy, Joseph Kable, Joshua Gold, University of Pennsylvania	225
III-58. Cell-type specific inhibitory plasticity enhances within assembly cooperation and between assembly competition in spiking networks Fereshteh Lagzi, Martha Canto Bustos, Anne-Marie Oswald, Brent Doiron, University of Pittsburgh	226

III-59. [Constraints on the time course of neural population activity](#)
 Alan Degenhart, Erinn Grigsby, Nicole McClain, Patrick Marino, Asma Motiwala, Emily Oby, Aaron Batista, Byron Yu, Carnegie Mellon University 226

III-60. [Metastable attractors in secondary motor cortex underlie self-initiated behavior](#)
 Stefano Recanatesi, Ulises Pereira, Masayoshi Murakami, Zachary Mainen, Luca Mazzucato, University of Washington 227

III-61. [Neural circuitry supporting approximate inference captures structure of correlations in hippocampus](#)
 Francesca Mastrogiuseppe, Naoki Hiratani, Peter Latham, Gatsby Computational Neuroscience Unit, University College London 227

III-62. [From deep learning to mechanistic understanding in neuroscience:revealing computational mechanisms of retinal prediction via model reduction](#)
 Hidenori Tanaka, Aran Nayebi, Niru Maheswaranathan, Lane McIntosh, Stephen Baccus, Surya Ganguli, Stanford University 228

III-63. [Learning a texture model for representing cortical area V2](#)
 Nikhil Parthasarathy, Eero Simoncelli, New York University 229

III-64. [Evidence that recurrent pathways between the prefrontal and inferior temporal cortex are critical during core object recognition](#)
 Kohitij Kar, James DiCarlo, Massachusetts Institute of Technology 229

III-65. [Parametric Copula-GP model reveals tuning of neuronal and behavioral relationships to visual stimuli](#)
 Nina Kudryashova, Theoklitos Amvrosiadis, Nathalie Dupuy, Nathalie Rochefort, Arno Onken, University of Edinburgh 230

III-66. [Cooling of medial septum reveals theta phase lag coordination of hippocampal cell assemblies](#)
 Peter Petersen, Gyorgy Buzsaki, New York University 230

III-67. [Identifying neural sequences in continuous time through point-process seqNMF](#)
 Alexander Williams, Anthony Degleris, Yixin Wang, Scott Linderman, Stanford University 231

III-68. [BehaveNet: behavioral video embedding and neural analysis toolbox](#)
 Matthew R Whiteway, Eleanor Batty, Shreya Saxena, Dan Biderman, Taiga Abe, Simon Musall, Winthrop Gillis, Jeffrey E Markowitz, Anne Churchland, John Cunningham, Sandeep Robert Datta, Scott Linderman, Liam Paninski, Columbia University 231

III-69. [Drosophila small object detectors trigger stopping with a novel, displacement-sensitive algorithm](#)
 Ryosuke Tanaka, Damon Clark, Yale University 232

III-70. [Inter-area interactions encode real-time adaptive biasing of performance during motor skill learning](#)
 Lucas Tian, W Hamish Mehaffey, Michael S Brainard, Massachusetts Institute of Technology 233

III-71. [Neural state space geometry in human motor cortex underlying speaking different phonemes](#)
 Guy Wilson, Sergey Stavisky, Donald Avansino, Leigh Hochberg, Jaimie Henderson, Krishna Shenoy, Shaul Druckmann, Stanford University 233

III-72. [Real-time discovery of effective dynamics from streaming noisy neural observations](#)
 Yuan Zhao, Josue Nassar, Il Memming Park, Stony Brook University 234

III-73. [Counterintuitive effects of excitatory and inhibitory perturbations in spiking networks](#)
 William Podlaski, Nuno Calaim, Christian Machens, Champalimaud Centre for the Unknown 234

III-74. [Deep Nose: Using artificial neural networks to represent the space of odors](#)
 Ngoc Tran, Sergey Shuvaev, Daniel Kepple, Alexei Koulakov, Cold Spring Harbor Laboratory 235

III-75. [Neural encoding characteristics match behaviorally measured prior expectations in visual speed perception](#)
 Lingqi Zhang, Alan Stocker, University of Pennsylvania 235

III-76. [Network principles predict motor cortex population activity structure across movement speeds](#)
 Shreya Saxena, Abigail Russo, John Cunningham, Mark Churchland, Columbia University 236

III-77. Rapid confidence-based hierarchical and counterfactual reasoning Mahdi Ramadan, Mehrdad Jazayeri, Massachusetts Institute of Technology	236
III-78. A canonical correlation analysis model of the cortical microcircuit David Lipshutz, Siavash Golkar, Yanis Bahroun, Anirvan Sengupta, Dmitri Chklovskii, flatiron Institute . . .	237
III-79. Investigating the ability of astrocytes to drive neural network synchrony Gregory Handy, Alla Borisyuk, University of Pittsburgh	237
III-80. Effects of structured neural correlations in population coding: beneficial or detrimental? Seyedamin Moosavi, Magalie Tatischeff, Bingyue Zhu, Hideaki Shimazaki, Kyoto University	238
III-81. A RNN that runs forward and backward: A new approach to biologically plausible sequence learning Huzi Cheng, Joshua Brown, Indiana University Bloomington	239
III-82. Generalized theory of dynamic balance with arbitrary low-rank structure Itamar D Landau, Haim Sompolinsky, Hebrew University of Jerusalem	239
III-83. The complete synaptic-resolution connectome of a <i>Drosophila</i> brain center for navigation Vivek Jayaraman, Shin-ya Takemura, Romain Franconville, Hannah Haberkern, Chuntao Dan, Brad Hulse, Marcella Noorman, Ruchi Parekh, Lou Scheffer, Stephen Plaza, Ann Hermundstad, Sandro Romani, Gerry Rubin, Dan Turner-Evans, HHMI Janelia Research Campus	240
III-84. Timescales of ongoing activity reflect local connectivity and are modulated during attention Roxana Zeraati, Yanliang Shi, Marc Gieselmann, Nicholas Steinmetz, Tirin Moore, Alexander Thiele, Tatiana Engel, Anna Levina, University of Tübingen	241
III-85. Probabilistic feedback updates priors in sequential decisions with accumulated evidence Khanh Nguyen, Zachary Kilpatrick, Kresimir Josic, University of Houston	241
III-86. Correlation-based spatial layout of deep neural network features generates ventral stream topography Eshed Margalit, Hyodong Lee, Tiago Marques, James J DiCarlo, Dan Yamins, Stanford University	242
III-87. Spectral cues are essential for the auditory azimuthal map in the mouse superior colliculus Shinya Ito, Yufei Si, David Feldheim, Alan Litke, University of California, Santa Cruz	243
III-88. Efficient coding in large networks of spiking neurons with tight-balance Jonathan Timcheck, Jonathan Kadmon, Surya Ganguli, Stanford University	243
III-89. Hippocampal replay rapidly and repeatedly adapts to reconfigurations of barrier wall structure in a changing complex maze John Widloski, David Foster, University of California, Berkeley	244
III-90. Top-down attention as efficient perceptual inference Wiktor Młynarski, Gasper Tkacik, Institute of Science and Technology Austria	244
III-91. Subpopulation coding reveals a mechanism for improved information flow through cortical circuits Matthew Getz, Chengcheng Huang, Brent Doiron, University of Pittsburgh	245
III-92. Population coding and network dynamics during OFF responses in auditory cortex Giulio Bondanelli, Thomas Deneux, Brice Bathellier, Srđjan Ostojic, Ecole Normale Supérieure	245
III-93. A motor cortical model of brain-machine interface learning, fast and slow Jorge Menendez, Jay Hennig, Matt Golub, Emily Oby, Aaron Batista, Steven Chase, Byron Yu, Peter Latham, Gatsby Computational Neuroscience Unit, University College London	246
III-94. A novel method to encode sequences in a computational model of speech perception Meropi Topalidou, Gregory Hickok, University of California, Irvine	246
III-95. Bayesian inference through attractor dynamics in medial entorhinal cortex Malcolm Campbell, Alexander Attinger, Samuel Ocko, Surya Ganguli, Lisa Giocomo, Stanford University . . .	247
III-96. Separable Manifold Geometry in Macaque Ventral Stream and DCNNs SueYeon Chung, Joel Dapello*, Uri Cohen, James J DiCarlo, Haim Sompolinsky, Columbia University . . .	247

Posters III

III-97. Saccades enhance perception via interactions between visual and motor signals in area V1 Seth Akers-Campbell, James Niemeyer, Michael Paradiso, Brown University	248
III-98. Neural correlations in a cortical circuit that generates behavioral variability Galen Lynch, Joseph Scherrer, Harbaljit Sohal, Natasha Denisenko, Ed Boyden, Michale Fee, Massachusetts Institute of Technology	249
III-99. Functional stripes in V2 are a byproduct of interleaved central-peripheral visuotopy Madineh Sedigh-Sarvestani, Kuo-sheng Lee, Juliane Jaepel, Rachel Satterfield, Nicole Shultz, David Fitzpatrick, Max Planck Florida Institute of Neuroscience	249
III-100. Dynamics and computation with anti-leaky integrate-and-fire neurons Paul Manz, Sven Goedeke, Raoul-Martin Memmesheimer, University of Bonn	250
III-101. Hierarchical tensor partitioning and tiling for learning coupled multiscale neuronal dynamics Gal Mishne, Hadas Benisty, Maria Lavzin, Shahar Levy, Ronen Talmon, Ron Meir, Jackie Schiller, Ronald Coifman, University of California, San Diego	251
III-102. Dynamic resource allocation during reinforcement learning accounts for ramping and phasic dopamine activity Minryung Song, Sang Wan Lee, KAIST	251
III-103. Dynamic network representation of invariant external context in medial entorhinal cortex Isabel Low, Alex Williams, Malcolm Campbell, Tucker Fisher, Scott Linderman, Lisa Giocomo, Stanford University	252
III-104. The neural code for face memory in macaque IT and perirhinal face patches Liang She, Marcus Benna, Stefano Fusi, Doris Tsao, California Institute of Technology	253
III-105. Maintenance of visual working memory by a visuomotor cortical loop Ivan Voitov, Tom Mrsic-flogel, University College London	253
III-106. Rapid sensorimotor adaptation through cortical input control Nicolas Meirhaeghe, Hansam Sohn, Mehrdad Jazayeri, Massachusetts Institute of Technology	254
III-107. Functional connectivity graph estimation from nonsimultaneous calcium imaging recordings Giuseppe Vinci, Genevera Allen, Gautam Dasarathy, Rice University	254
III-108. Context-dependent representation of the decision-making process in the monkey prearcuate gyrus Mohamad Saleh Esteki, Marcus Benna, Stefano Fusi, Roozbeh Kiani, New York University	255
III-109. The role of thalamus in maintenance of sequential cortical activity Nader Nikbakht, Michale Fee, Massachusetts Institute of Technology	255
III-110. A common cause for multiple suboptimalities in perceptual decision-making Diksha Gupta, Carlos D Brody, Princeton University	256
III-111. Fast gating and slow amplification of relevant evidence in frontal cortex during flexible decision-making Marino Pagan, Vincent Tang, Mikio Aoi, Jonathan Pillow, Carlos D Brody, Princeton University	257
III-112. Reorganization of neuronal population codes in the ACC underlies long-term memory retrieval Meytar Zemer, Noa Brande-Eilat, Liron Sheintuch, Alon Rubin, Ofer Yizhar, Yaniv Ziv, Weizmann Institute of Science	257
III-113. Brain-inspired replay in artificial neural networks Gido van de Ven, Andreas Tolias, Baylor College of Medicine	258
III-114. Stimulus encoding in the lateral intraparietal parietal cortex during categorization depends on training history Kenneth Latimer, David Freedman, The University of Chicago	258
III-115. A computational model for border ownership Miguel Lazaro-Gredilla, Rajeev Rikhye, Dileep George, Wolfgang Lehrach, Vicarious AI	259
III-116. High-capacity pattern labeling and robust retrieval with internally modular attractor networks Ila Fiete, Rishidev Chaudhuri, Massachusetts Institute of Technology	259

III-117. [AutoLFADS: A large-scale neural network training framework for generalized estimation of single-trial population dynamics](#)
 Mohammad Reza Keshtkaran, Raghav Tandon, Diya Basrai, Sarah L Nguyen, Raeed H Chowdhury, Hansem Sohn, Mehrdad Jazayeri, Lee E Miller, Chethan Pandarinath, Emory University 260

III-118. [Understanding the functional and structural differences across excitatory and inhibitory neurons](#)
 Sun Minni, Li Ji-An, Theodore Moskovitz, Grace Lindsay, Kenneth Miller, Mario Dipoppa, Guangyu Robert Yang, Columbia University 261

III-119. [Antagonism normalizes input to the olfactory system in freely breathing animals](#)
 Joseph Zak, Venkatesh Murthy, Harvard University 261

III-120. [Spatio-temporal dynamics of cortical cholinergic signaling](#)
 Sweyta Lohani, Andrew Moberly, Mike Higley, Jessica Cardin, Yale University 262

III-121. [Ongoing correction of temporally perturbed spike sequences by fluctuating scalar modulation](#)
 Rich Pang, Adrienne Fairhall, University of Washington 262

III-122. [The caudate nucleus controls coordinated, reward-dependent adjustments to perceptual decisions](#)
 Takahiro Doi, Yunshu Fan, Joshua Gold, Long Ding, University of Pennsylvania 263

III-123. [Inhibitory inputs onto premotor neurons regulate vocal timing in an avian model of vocal turn-taking](#)
 Jonathan Benichov*, Philipp Norton*, Susanne Schreiber, Daniela Vallentin, Max Planck Institute for Ornithology 263

III-124. [Neural circuit strategy for categorical selection in the midbrain](#)
 Nagaraj Mahajan, Shreesh Mysore, Johns Hopkins University 264

Abstracts

Abstracts for talks appear first, in order of presentation; those for posters next, in order of poster session and board number. An index of all authors appears at the back.

T-1. Deep reinforcement learning and its neuroscientific implications

Matthew Botvinick^{1,2}

BOTVINICK@GOOGLE.COM

¹DeepMind

²Princeton University

The last few years have seen some dramatic developments in artificial intelligence research. What implications might these have for neuroscience? Investigations of this question have, to date, focused largely on deep neural networks trained using supervised learning, in tasks such as image classification. However, there is another area of recent AI work which has so far received less attention from neuroscientists, but which may have more profound neuroscientific implications: Deep reinforcement learning. Deep RL offers a rich framework for studying the interplay among learning, representation and decision-making, offering to the brain sciences a new set of research tools and a wide range of novel hypotheses. I'll provide a high level introduction to deep RL, discuss some recent neuroscience-oriented investigations from my group at DeepMind, and survey some wider implications for research on brain and behavior.

T-2. Computation-through-dynamics for the temporal control of behavior

Mehrdad Jazayeri

MJAZ@MIT.EDU

Massachusetts Institute of Technology

A central aim of systems neuroscience is to understand mental computations in terms of neurobiological mechanisms. Recent experimental advances in animal models have begun to investigate the neural basis of cognition at progressively higher resolutions, for example by monitoring and perturbing specific cell-types and synapses. However, characterizing the neurobiological details is not enough. The bridge between brain and cognition has to be built upon a rigorous mathematical foundation that can explain how the structure and dynamic of neural activity across populations of neurons support mental computations. Here, I will describe our efforts integrating cognitive neuroscience theories with in-vivo experiments in non-human primates and in-silico experiments in recurrent neural network models to understand how the brain performs Bayesian inference.

T-3. The role of the striatal indirect pathway/D2R+ spiny projection neurons in choice and rejection

Linda Wilbrecht

WILBRECHT@BERKELEY.EDU

University of California, Berkeley

Bilateral optogenetic stimulation experiments implicate the dorsal striatal dopamine D1 receptor expressing (D1R+) direct pathway spiny projection neurons (SPNs) in promoting action (such as running) and indirect pathway D2R+ SPNs in freezing, stopping and aversion. There is ongoing interest in how activity in the striatal direct and indirect pathways can sculpt more cognitive aspects of choice selection and suppression or rejection. To experimentally probe the function of the D1R and D2R+ SPNs in choice we performed a series of experiments using optogenetics and chemogenetics in the dorsal and ventral striatum in a lateralized two alternative choice task and a non lateralized multiple choice foraging task. Our studies implicate a role for activity in the dorsal striatum in choice selection and the ventral striatum in reinforcement of a past choice with pathway specific differences. Specifically, for D1R+ SPNs we found that unilateral optogenetic stimulation of dorsal striatal D1R+ SPNs could drive a contralateral bias in left vs. right action selection within a trial while ventral striatum D1R+ SPNs stimulation had robust effects on the action chosen in the subsequent trial. In these ventral striatum experiments, D1R+ SPN excitation reinforced previous actions while D1+ SPN inhibition reinforced the alternate action. Data from D2R+ SPNs manipulation generated more unexpected results. We found that dorsal striatal D2R+ SPNs could bias choice in an ipsilateral direction (when stimulated unilaterally) and could both speed and slow action depending on congruence with trial history. Chemogenetic excitation (but not inhibition) of the D2R+ indirect pathway in the dorsal striatum disrupted learned choice suppression in a foraging task. Finally, we found in the ventral striatum that optogenetic excitation of D2R+ SPNs could drive reinforcement of past choices in the context of a decision making task but did not support intracranial self stimulation outside of a complex task. These data suggest we need to consider more nuanced models of the striatal D2R+ SPNs within the context of the larger basal ganglia circuit anatomy and examine function outside of the context of stopping or freezing.

T-4. Automating parameter inference in theoretical neuroscience

John Cunningham

JPC2181@COLUMBIA.EDU

Columbia University

A cornerstone of neuroscience is the circuit model: a system of equations that captures a hypothesized neural computation. The operation of these circuits critically depends on model parameters, and the historical gold standard has been to analytically derive the relationship between model parameters and computational properties. With larger, more biologically realistic, and more complex models, this gold standard is infeasible, often resulting in ad hoc approaches to understanding the relationship between model and computation. I will first describe our recent work bringing deep generative models to bear on this problem, inferring distributions of parameters that induce specified computational properties. I will then point to our use of this approach across the stomatogastric ganglion, primary visual cortex, superior colliculus, and task-solving recurrent neural networks. Lastly, I will demonstrate how to deploy this approach (and others) effortlessly and at scale using a new platform-as-a-service for neural data analyses. Taken together, this talk points towards the use of modern machine learning and modern compute infrastructure for sophisticated interrogation of biologically relevant models.

T-5. Circuits for learning, predicting, and deciding in *Drosophila*

Marta Zlatic

ZLATICM@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

Dopaminergic neurons (DANs) drive learning across the animal kingdom, but the upstream circuits that regulate their activity and thereby learning remain poorly understood. We provide a synaptic-resolution connectome of the circuitry upstream of all DANs in a learning center, the mushroom body (MB) of *Drosophila* larva. We discover afferent sensory pathways and a large population of neurons that provide feedback from MB output neurons and link distinct memory systems (aversive and appetitive). We combine this with functional studies of DANs and their presynaptic partners and with comprehensive circuit modelling. We find that DANs compare convergent feedback from aversive and appetitive systems which enables the computation of integrated predictions that may improve future learning. Computational modelling reveals that the discovered feedback motifs increase model flexibility and performance on learning tasks. The connectome, together with the functional and modelling studies provides exciting opportunities for elucidating the biological implementation of reinforcement learning algorithms.

T-6. Evidence for task agnostic feedback in visual cortex

Hendrikje Nienborg
NEI/NIH

HNIENB@GMAIL.COM

Neural responses in visual cortex are influenced not only by the feed-forward retinal input, but also by feedback that depends on an animal's behavioral goal. The computational role of such task-dependent feedback is subject to debate. One general view is that this ability to process stimuli selectively is a mechanism that supports perceptual inference. That is, it is a way to incorporate prior knowledge about the relevance of the respective sensory information at the level of the sensory representation. This view implies that the task- and behavioral relevance predicts the modulation by feedback. I will present data that test this view in population recordings in mid-level visual cortex. In contrast to the above prediction our results provide evidence for a common feedback mechanism across different tasks that appears agnostic to the specific task performed. While this simplifies the mechanistic implementation, it raises further questions about the computational role of such feedback.

T-7. The complex behaviors of the 'simple' cerebellar circuit

Megan Carey
Champalimaud Centre for the Unknown

MEGAN.CAREY@NEURO.FCHAMPALIMAUD.ORG

Nearly all of our actions require us to precisely coordinate the movement of multiple body parts in space and time. The cerebellum is critical for coordinated motor control and learning, but how cerebellar outputs enable complex, coordinated movement is still poorly understood. Our lab has taken a behavior-centric approach to this problem, focusing on cerebellar contributions to locomotor coordination. first, I will briefly describe our efforts to establish a quantitative framework for locomotion in freely walking mice and identify specific, cerebellum-dependent features of locomotor coordination. Second, I will describe a form of cerebellum-dependent locomotor learning, in which mice adapt their locomotor patterns to achieve a more symmetrical gait while walking on a split-belt treadmill. These two approaches have provided strong behavioral evidence for a particularly important role for the cerebellum in the control of interlimb coordination. finally, I will present recent electrophysiological recordings from Purkinje cells in locomoting mice, which are beginning to provide insights into how cerebellar outputs ensure accurate and coordinated complex, whole-body behaviors.

T-8. The architecture of exploration

Sam Gershman
Harvard University

SAM.GERSHMAN@GMAIL.COM

The trade-off between exploration and exploitation is one of the fundamental dilemmas in reinforcement learning. I review tractable algorithmic solutions to this dilemma and how they might be implemented in the brain. Recent experiments demonstrate that humans use sophisticated uncertainty-guided exploration algorithms.

T-9. Human planning in large state spaces

Wei Ji Ma
New York University

WEIJIMA@NYU.EDU

As deep reinforcement learning has revolutionized the artificial intelligence of planning, our lack of understanding of human planning in large state spaces has come into stark contrast. Chess once seemed promising as an experimental paradigm to study such planning, but this promise was never realized. My lab uses tasks of intermediate complexity that allow for rigorous computational modeling. Our most developed paradigm is four-in-a-row ([basvanopheusden.github.io](https://github.com/basvanopheusden)), a variant of tic-tac-toe. I will describe a computational model of human play in this task, inspired by best-first search and fitted to individual subjects' moves. The model predicts moves in unseen positions, decisions under novel instructions, eye and cursor locations, and response times. The model further allows us to analyze the effects of time pressure and expertise, as well as the balance between model-based and model-free systems. I will describe collaborations that use four-in-a-row and involve mobile game data, human neuroimaging, monkey physiology, vmPFC lesion patients, and developmental data. I will describe other tasks in the same space, some of which are designed to study collaborative planning and effective helping. In summary, I hope to convince you that studying human planning in tasks of intermediate complexity in conjunction with AI-inspired algorithms is a promising direction in computational cognitive science.

T-10. Parallel social information processing circuits are differentially impacted in autism

Gul Dolen
Johns Hopkins University

GUL@DOLENLAB.ORG

Parallel processing circuits are thought to dramatically expand the network capabilities of the nervous system. Magnocellular and parvocellular oxytocin neurons have been proposed to subservise two parallel streams of social information processing, which allow a single molecule to encode a diverse array of ethologically distinct behaviors, although to date direct evidence to support this hypothesis is lacking. Here we provide the first comprehensive characterization of magnocellular and parvocellular oxytocin neurons, validated across anatomical, projection target, electrophysiological, and transcriptional criteria. We next used novel multiple feature selection tools in Fmr1 KO mice to provide direct evidence that normal functioning of the parvocellular but not magnocellular oxytocin pathway is required for autism relevant social reward behavior. Finally, we demonstrate that autism risk genes are uniquely enriched in parvocellular oxytocin neurons. Taken together these results provide the first evidence that oxytocin pathway specific pathogenic mechanisms account for social impairments across a broad range of autism etiologies.

T-11. Multiple maps for navigation

Lisa Giocomo
Stanford University

GIOCOMO@GMAIL.COM

Over the last several decades, the tractable response properties of parahippocampal neurons have provided a new access key to understanding the cognitive process of self-localization: the ability to know where you are currently located in space. Defined by functionally discrete response properties, neurons in the medial entorhinal cortex and hippocampus are proposed to provide the basis for an internal neural map of space, which enables animals to perform path-integration based spatial navigation and supports the formation of spatial memories. My lab focuses on understanding the mechanisms that generate this neural map of space and how this map is used to support behavior. In this talk, I'll discuss how learning and experience shapes our internal neural maps of space to guide behavior.

T-12. Connectivity and computation in olfaction

Rainer Friedrich
FMI Basel

RAINER.FRIEDRICH@FMI.CH

The olfactory bulb (OB) transforms correlated patterns of activity across its input channels, the olfactory glomeruli, into decorrelated and normalized patterns of activity across its output neurons, the mitral cells. This transformation can be described as a whitening of activity patterns which is thought to support pattern classification in olfactory cortex. Whitening cannot be explained by global scaling or thresholding operations but requires structured interactions between specific subsets of neurons. We measured odor-evoked activity throughout the OB of a zebrafish larva and subsequently reconstructed the full wiring diagram by volume electron microscopy. This “functional connectomics” approach revealed an overrepresentation of triplet connectivity motifs that privileges multisynaptic reciprocal inhibition among mitral cells with similar tuning. Tuning-dependent multisynaptic connectivity specifically suppressed activity of output neurons that made high contributions to pattern correlations. Connectivity in the OB therefore mediates contrast reduction, rather than contrast enhancement, which resulted in feature suppression. This connectivity was necessary and sufficient to reproduce whitening of natural input patterns in generic network models. Hence, whitening in the OB is achieved by higher-order structure in the wiring diagram that is adapted to natural input patterns. These results provide direct insights into the network mechanism underlying a fundamental neural computation and illustrate the potential of “functional connectomics” approaches to analyze complex structure-function relationships in neuronal circuits.

T-13. Cortical networks for flexible decisions during spatial navigation

Christopher Harvey
Harvard University

CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

We study flexible decision-making in mice during spatial navigation in virtual reality environments. I will present approaches to study this topic using analyses of population dynamics from calcium imaging movies, optogenetic perturbations, and modeling of behavioral variability. We have developed behavioral tasks in which mice integrate internally stored contextual information, in the form of short-term memories, with sensory cues to guide navigation choices. We have used optogenetics-based screening approaches to identify regions of the mouse cortex involved in these tasks, leading to a particular focus on posterior parietal cortex, retrosplenial cortex, and V1. I will present findings on the spatial structure of encoding across posterior parts of cortex during flexible navigation decisions. I will discuss potential functions for posterior cortical areas in flexibly mapping sensory and internal information onto navigation actions.

T-14. flexible neuronal dynamics during categorization task switching in artificial and biological neural networks

Krithika Mohan
Oliver Zhu
David Freedman

The University of Chicago

KRMOHAN@UCHICAGO.EDU
ZHOU@UCHICAGO.EDU
DFREEDMAN@UCHICAGO.EDU

Animals exhibit remarkable flexibility in their behavior by efficiently adapting to changes in their environment. Although much research has quantified how artificial and biological networks adapt to changes in sensory inputs or motor outputs, it is unknown how networks adapt to changes in cognitive demands, such as working memory, within the framework of a single task. To advance our understanding of flexibility in biological networks, we trained monkeys to perform categorization task switching while recording neural responses from the posterior parietal cortex. Given a contextual cue, monkeys categorized visual motion stimuli by either making a rapid eye movement to a colored target associated with the category or by making a delayed hand movement to indicate whether sequentially presented stimuli belonged to the same category. To discover the range of viable solutions for task switching, we studied artificial recurrent neural networks (RNNs) trained to perform the same task as behaving animals. In artificial and biological neural networks, we find similar patterns of category selectivity in both tasks, but surprisingly, categorical encoding is more binary and more persistent in the delayed matching task. We hypothesize that attractor dynamics underlying the generation of persistent activity is a candidate mechanism by which neural networks reduce the dimensionality of stimulus representations. Because only the delayed task requires maintaining category information in working memory, attractor dynamics appears to compress category-related information to a simpler, binary format by collapsing all directions within a category toward a single population state. Indeed, analysis of the fixed-point structure of trained RNNs shows that following stimulus onset, sample categories produce stable attractors in the delayed matching task but not in the rapid categorization task. Together, our findings suggest that neural networks mediate flexible task switching through generalized category representations that are reorganized through attractor dynamics that underlie working memory-based computations.

T-15. Dissecting modularity and redundancy in multi-regional circuits by population recording and modeling

Byungwoo Kang*¹
Guang Chen*²
Nuo Li**²
Shaul Druckmann**¹

¹Stanford University
²Baylor College of Medicine

BKANG@STANFORD.EDU
GUANG.CHEN@BCM.EDU
NUO.LI@BCM.EDU
SHAULD@STANFORD.EDU

The brain plans volitional movements before they are executed. Preparatory activity is distributed across multiple brain regions, but it is unclear what is the function of these parallel representations. Nor is it known how interactions between brain regions mediate preparatory activity. Activity in the mouse anterior lateral motor cortex (ALM) instructs specific upcoming movements. ALM preparatory activity is robust to large-scale perturbations and this robustness is dependent on ALM interhemispheric interactions. We examined ALM interhemispheric interactions using bilateral silicon probe recordings and optogenetic perturbations. Preparatory activity was correlated across the hemispheres, i.e., in single trials, activity in both hemispheres predicted the same future movement. Therefore, the motor plan was redundantly represented in both hemispheres. Surprisingly, perturbation revealed different interhemispheric interactions across mice. In some mice, the two hemispheres were strongly coupled and silencing one hemisphere substantially reduced preparatory activity in the other hemisphere. In other mice, the two hemispheres appeared modular and weakly coupled: each hemisphere could independently maintain preparatory activity without the other. The distinct interhemispheric interactions were results of learning. Modular organization increased the robustness of preparatory activity. When one hemisphere suffered a perturbation, signals

from the other intact hemisphere could help restore correct activity. Building predictive models of multi-regional neuronal dynamics, we were able to infer these distinct interhemispheric interactions across mice from control trial preparatory activity. A state-dependent gating of information flow between hemispheres was inferred by our models and it predicted the robustness of behavioral performance. finally, using artificial RNNs, we identified conditions in which distinct interhemispheric interactions and robustness against perturbations can arise. These results reveal degeneracies in multi-regional circuits and present experimental and computational approaches to dissect redundant and modular representations across brain regions. Such organization may enable dynamical error- correction that allows multi-regional circuits to produce robust behaviors under challenging circumstances.

T-16. flexible neural control of motor units revealed via latent factor models

Najja Marshall
Joshua Glaser
Sean Perkins
Larry Abbott
John Cunningham
Mark Churchland
Columbia University

NJM2149@COLUMBIA.EDU
JIG2124@COLUMBIA.EDU
SP3222@COLUMBIA.EDU
LFA2103@COLUMBIA.EDU
JPC2181@COLUMBIA.EDU
MC3502@COLUMBIA.EDU

Voluntary movement requires communication from cortex to the spinal cord, where motor units (MUs) relay neural drive to muscles. The canonical description of MU control rests upon two foundational tenets. first, cortex cannot control MUs independently; rather, it supplies a shared one-dimensional 'common drive' to each muscle's MU pool. Second, MU firing rates are sigmoidal functions of that common drive. The point where each function rises is determined by MU 'size' (number of muscle fibers activated when the motor neuron fires). Thus, MUs are recruited in a consistent order as force rises. These tenets comprise the decades-old 'size principle': MU activities are a prespecified function of a one-dimensional descending command.

We re-examined both central tenets using a novel task, multi-channel cortical microstimulation, and intramuscular recordings from multiple single MUs. We trained two rhesus macaques to produce force profiles at multiple frequencies (0-3 Hz). In separate experiments, cortical stimulation during task performance induced small transient increases in force. The size principle predicts that MU activities should be a reliable function of a single underlying latent variable. Yet this was not the case during either natural force production or stimulation-induced forces. For example, a given MU could be preferentially active during static forces, and another during 3 Hz forces. Furthermore, stimulation through different electrode channels could recruit MUs independently of one another. We developed a probabilistic latent factor model to assess whether empirical MU activity could be accounted for by a single latent factor, with each MU having a unique, monotonically increasing link function. We found that single-factor models were insufficient during both natural performance and artificial stimulation. Instead, multiple factors were needed to account for the diversity of MU responses. Thus, contrary to classical predictions, MUs are flexibly controlled to meet task demands.

T-17. Evidence of a memory trace in motor cortex after short-term learning

Darby Losey¹
 Jay Hennig¹
 Emily Oby²
 Matt Golub³
 Patrick Sadtler²
 Kristin Quick²
 Stephen Ryu³
 Elizabeth Tyler-Kabara²
 Aaron Batista²
 Byron Yu¹
 Steven Chase¹

DLOSEY@ANDREW.CMU.EDU
 JHENNIG@ANDREW.CMU.EDU
 EMILYOBY@GMAIL.COM
 MGOLUB@STANFORD.EDU
 PATRICK.T.SADTLER@GMAIL.COM
 KRISTINMQUICK@GMAIL.COM
 SEOULMANMD@GMAIL.COM
 ELIZABETH.TYLER-KABARA@CHP.EDU
 AARON.BATISTA@PITT.EDU
 BYRONYU@CMU.EDU
 SCHASE@ANDREW.CMU.EDU

¹Carnegie Mellon University

²University of Pittsburgh

³Stanford University

Does learning a new task change the neural activity patterns that the brain uses to perform a previously learned task? We hypothesized that neural activity used to re-perform an original task would remain appropriate for the new task. A difficulty in addressing this hypothesis is that, in most circumstances, the causal relationship between neural activity and behavior is unknown. To overcome this, we leveraged an intracortical brain-computer interface (BCI), where the mapping between neural activity and behavior is specified by the experimenters. This experimental paradigm allows us to quantify how appropriate neural activity is for a given mapping, even when that mapping is not currently being used. Experiments utilized an “ABA” block design, where monkeys proficiently used an original mapping A in a baseline period, and then the mapping was switched to a new mapping B that had to be learned through trial and error. After several hundred trials, mapping A was reinstated in a washout period. To evaluate whether learning a new mapping leaves a memory trace in motor cortex, we quantified the extent to which neural activity patterns produced when the monkey was controlling the cursor with mapping A were appropriate for mapping B. We found evidence of a memory trace, in that neural activity during the late portion of the washout period was more appropriate for mapping B than the neural activity during the baseline period. This result held when we controlled for factors such as behavioral differences between the baseline and washout periods. Overall, our findings suggest that learning a new task brings neural activity to a novel solution for the original task that also provides a benefit for the new task. This may be a mechanism by which the brain could more rapidly learn when re-exposed to the same perturbation, a phenomenon known as “savings”.

T-18. Separation of preparatory neural states when learning multiple arm-movement dynamics

Xulu Sun
 Daniel O’Shea
 Matt Golub
 Eric Trautmann
 Stephen Ryu
 Krishna Shenoy
 Stanford University

XULU2@STANFORD.EDU
 DJOSHEA@STANFORD.EDU
 MGOLUB@STANFORD.EDU
 ETRAUTMANN@GMAIL.COM
 SEOULMANMD@GMAIL.COM
 SHENOY@STANFORD.EDU

The ability to learn new motor skills without interfering with old ones is essential for us to acquire and maintain a broad motor repertoire. Recent psychophysics studies found that when people plan for or imagine different movements associated with different curl fields, they can learn multiple skills without interference that would otherwise hurt learning (Sheahan et al. 2016 and 2019). Here we ask when learning multiple motor skills, are preparatory neural states of these skills separated to form different motor plans and subsequently reduce interference between

neural dynamics that generate distinct movements? We designed a curl force field task to investigate the neural correlate of learning multiple motor skills (here, arm-movement dynamics). We trained two rhesus macaques to learn different curl fields sequentially within the same session or over multiple sessions (block design of the task is similar to Sun et al. 2019 COSYNE abstract). We recorded neural activity in dorsal premotor and primary motor cortex using Utah-arrays, V-probes, and Neuropixels probes which gave us access to hundreds of neurons per curl field condition. We applied PCA to preparatory activity and examined the subspace spanned by the top 10 PCs that captured over 90% of the total variance. We sought response patterns most strongly present in the data and asked whether the separation of preparatory states appeared in the most prominent neural activity features. We found that after learning one curl field applied to only one reach target, preparatory neural states of reaching to all targets uniformly shifted away from the before-learning states. Preparatory states of learning different curl fields were separated by multiple uniform shifts that showed distinct geometrical relationships depending on curl field types. The uniform shifts raise the possibility that interference between neural dynamics of learning multiple motor skills is reduced by having separated preparatory neural states.

T-19. Long-range inhibition mediates decision making in the superior colliculus

Jaclyn Essig
Josh Hunt
Gidon Felsen

University of Colorado

JACLYN.ESSIG@CUANSCHUTZ.EDU
JOSHUA.HUNT@CUANSCHUTZ.EDU
GIDON.FELSEN@CUANSCHUTZ.EDU

Decision making is a fundamental process of the nervous system for generating goal-directed behaviors. The mid-brain superior colliculus (SC) contributes to sensorimotor decision making by integrating cortical and subcortical inputs to guide orienting movements of the eyes, head, and body towards spatial goals. The SC is topographically organized and encodes for specific regions in retinotopic space, however the underlying circuitry for how the SC selects where to orient is unknown. Multiple models of excitatory/inhibitory interactions have been proposed to describe SC function, but these are based on cellular anatomy, ex-vivo slice physiology, and in-vivo recordings in the absence of behavior, or on recordings during behavior from unknown cell types. Here, we record and manipulate the activity of GABAergic neurons in mice performing a spatial choice task to determine the functional role of inhibition during spatial choice. We trained mice to select a left or right reward port based on a binary odor mixture. Importantly, after odor delivery, mice wait for a “go tone” before orienting to the reward port, giving us access to neural activity during the decision (i.e., the “choice epoch”). We hypothesized that GABAergic neurons would shape spatial choice locally by inhibiting SC motor output neurons promoting contralateral choice, and therefore predicted that these cells would be most active before an ipsilateral choice. However, optogenetic identification (i.e., “optotagging”) and activation of channelrhodopsin-expressing GABAergic neurons revealed that GABAergic neurons are active before contralateral choices and driving their activity during the “choice epoch” biases mice to select the contralateral port. A biologically restricted attractor model recapitulated our behavioral results and supports a role for long-range inhibitory interactions between the left and right SC. Our work has revealed a neural mechanism for how the SC implements spatial choice.

T-20. Spontaneous activity and multi-sensory integration in the developing higher-order cortex

Marina Wosniack¹
 Jan Hendrik Kirchner¹
 Paloma Maldonado²
 Nawal Zabouri²
 Christian Lohmann²
 Julijana Gjorgjieva¹

MARINA-ELAINE.WOSNIACK@BRAIN.MPG.DE
 JAN.KIRCHNER@BRAIN.MPG.DE
 P.MALDONADO@NIN.KNAW.NL
 N.ZABOURI@NIN.KNAW.NL
 C.LOHMANN@NIN.KNAW.NL
 GJORGJIEVA@BRAIN.MPG.DE

¹Max Planck Institute for Brain Research

²Netherlands Institute for Neuroscience

In the developing rodent brain, the primary visual (V1) and somatosensory (S1) cortices exhibit patterns of spontaneous activity even before the onset of sensory experience. Multiple studies provide evidence towards a general developmental profile of spontaneous activity: early global and dense spontaneous events transition to more localized and sparse patterns typical of adult networks. Interestingly, before eye-opening spontaneous events in S1 are sparser than their counterparts in V1, implying an earlier maturation of S1 relative to V1. The spontaneous activity in V1 and S1 propagates to higher cortical areas, and one important common target is the rostralateral area (RL). Area RL integrates multi-sensory inputs (visual and tactile), and more than 60% of its neurons are bimodal. Furthermore, the visual and tactile maps are topographically aligned in RL. Spontaneous activity is known to instruct various aspects of organization in the developing brain. Here we explore the hypothesis that the combined activity of V1 and S1 contains instructive cues for the correct refinement of connections in RL. We combined a detailed analysis of in vivo wide-field recordings of spontaneous activity in the V1- S1-RL region and a computational model to investigate (1) the emergence of bimodal RL neurons and (2) the role of more mature S1 events in the topographic alignment of the visual-tactile map in RL. We found that bimodal RL neurons develop when a critical amount of correlation exists between events in V1 and S1. By simulating a scenario where S1-RL connectivity is initialized with a topographical bias while V1-RL is randomly connected, we found that events in S1 can "teach" the topography to V1-RL connections more effectively when S1 events are sparser. Examining in vivo spontaneous activity, we found that the correlation between events from V1 and S1 is sufficiently high to instruct the alignment of their maps.

T-21. Learning rate adjustments in a volatile environment by mouse medial prefrontal cortex

Huriye Atilgan
 Cayla Murphy
 Hongli Wang
 Alex Kwan
 Yale University

HURIYE.ATILGAN@GMAIL.COM
 CAYLA.MURPHY@YALE.EDU
 HONGLI.WANG@YALE.EDU
 ALEX.KWAN@YALE.EDU

Animals rely on the experience of past choices and outcomes to guide current decisions. But how should they weigh recent evidence relative to distant events in learning? Human and non-human primate studies suggest that volatility of the environment – which could be inferred by the animal from the statistics of the reward history – is a crucial factor in optimizing the learning rate. Namely, in a volatile environment, we want to update often and use information from the recent past. By contrast, in a stable environment, we may integrate over many trials for a more reliable estimate of the world. Still unknown, however, is the extent to which mice are capable of such meta-reinforcement learning. Moreover, causal evidence for the potential neural substrates is lacking.

In this study, head-fixed mice performed a two-armed bandit task. By design, the reward probabilities switched after a varying number of trials, leading to differences in volatility. Analyses of the performance, including fits to reinforcement learning algorithms, indicated that mice dynamically adjusted their learning rate based on estimates

of the reward statistics. To determine the neural substrate, we tested mice before and after excitotoxic lesion of the cingulate and premotor regions (Cg1/M2) of the medial prefrontal cortex. Unilateral lesion diminished learning rate adjustments. Intriguingly, deficits were specific to switches towards the contralateral side, suggesting aberrant volatility estimation from contralateral rewards.

Competing theories posited that volatility estimates could be used for computations during action selection, value updating, or both processes. To disambiguate these possibilities, we optogenetically silenced Cg1/M2 activity either before or after the decision in each trial. We found that post-choice, but not pre-choice, inactivation reproduced the learning deficits. Altogether, the results demonstrate that mice are sensitive to the statistics of the reward history, and the post-decisional adjustment to learning depends crucially on Cg1/M2.

T-22. Sequential Component Analysis (SCA)

Virginia Rutten¹
Alberto Bernacchia²
Guillaume Hennequin²

GIN.RUTTEN@GMAIL.COM
AB2349@CAM.AC.UK
GJE.HENNEQUIN@GMAIL.COM

¹Gatsby Computational Neuroscience Unit, University College London

²University of Cambridge

From molecules to whole organisms, the dynamics of natural living systems depart from “thermodynamic equilibrium” [Gnesotto, 2018]. Statistically, the state trajectories produced by such systems are non-reversible, i.e. they do not have equal likelihood of occurring in the reverse temporal direction even in stationary regimes. For example, neural systems exhibit anisotropic waves of activity (e.g. during development), produce precisely ordered spike sequences (e.g. episodic memory), or generate rotational patterns of activity (e.g. motor control). Despite the prevalence and importance of sequential neural activity, there is a relative paucity of methods for exploring the spatio-temporal structure of irreversibility in multivariate time series. Here, we introduce Sequential Components Analysis (SCA), a simple yet effective and scalable method for doing this.

SCA performs a systematic analysis of spatio-temporal covariances (all time pairs and unit pairs), and extracts the spatio-temporal modes of activity that contribute most to sequentiality. We highlight the main distinguishing features of the method using a toy example, and apply it to monkey M1 motor activity as well as rat hippocampal data where we show that sequential features separate navigational memories better than mere principal components.

T-23. Nonlinear computations in semi-balanced networks

Cody Baker
Vicky Zhu
Robert Rosenbaum

CBAKER9@ND.EDU
RZHU@ND.EDU
ROBERT.ROSENBAUM@ND.EDU

University of Notre Dame

An approximate balance between excitation and inhibition is widely observed in cortical recordings. How does this balance shape neural computations and stimulus representations? This question is often approached using computational models of neuronal networks that realize a dynamically stable balanced state. These “balanced network models” predict a linear relationship between stimuli and population responses in contrast to the nonlinearity of cortical computations. In addition, we show that every balanced network architecture admits some stimuli that break the classical balanced state. We find that these breaks in balance push the network into a “semi-balanced state” characterized by excess inhibition to some neurons, but an absence of excess excitation. We develop a theory of semi-balanced networks, showing that the corresponding semi-balanced state is unavoidable in networks driven by multiple stimuli and is more consistent with recorded data. Stimulus representations in semi-balanced networks are nonlinear and have a direct, mathematical relationship to artificial neural networks with rectified linear activations. This relationship allows us to construct and train biologically realistic, semi-balanced spiking

neural networks that implement nonlinear functions, such as the XOR function, and classify hand-written digits. In summary, the strong and dense coupling of cortical circuits combined with the presence of time-varying stimuli imply that cortical circuits are in a semi-balanced state. Our analysis of this state shows a direct correspondence to artificial neural networks and therefore has extensive implications for the computational properties of cortical circuits.

T-24. Gradient-based learning with Hebbian plasticity in structured and deep neural networks

Dina Obeid
Cengiz Pehlevan
Harvard University

DINA OBEID@SEAS.HARVARD.EDU
CPEHLEVAN@SEAS.HARVARD.EDU

Synaptic plasticity is widely accepted to be the mechanism behind learning in the brain's neural networks. A central question is how synapses, with access to only local information about the network, can still organize collectively and perform circuit-wide learning in an efficient manner. In single-layered and all-to-all connected neural networks, local plasticity has been shown to implement gradient-based learning on a class of cost functions that contain a term that aligns the similarity of outputs to the similarity of inputs. Whether such cost functions exist for networks with other architectures is not known. In this work, we introduce structured and deep similarity matching cost functions, and show how they can be optimized in a gradient-based manner by neural networks with local learning rules. These networks extend Foldiak's Hebbian/Anti-Hebbian network to deep architectures and structured feedforward, lateral and feedback connections. Credit assignment problem is solved elegantly by a factorization of the dual learning objective to synapse specific local objectives. Simulations show that structured and deep similarity-based unsupervised networks learn meaningful features.

T-25. Neural mechanism for illusory motion perception from stationary patterns

Margarida Agrochao
Ryosuke Tanaka
Damon Clark
Emilio Salazar-Gatzimas
Yale University

MARGARIDA.AGROCHAO@YALE.EDU
RYOSUKE.TANAKA@YALE.EDU
DAMON.CLARK@YALE.EDU
EMILIO.SALAZARCARDOSO@YALE.EDU

Visual motion detection is one of the most important computations performed by visual circuits in the brain. Yet, we perceive vivid illusory motion in periodic, stationary luminance gradients, such as the Fraser-Wilcox illusion. This illusion is shared by diverse vertebrate species, likely reflecting common motion detection algorithms. The theories proposed to explain this illusion have remained difficult to test. Here, we demonstrate that the fruit fly *Drosophila*, like humans and other vertebrates, perceives illusory motion in periodic, stationary luminance gradients. This illusory motion percept was abolished by genetic silencing of the earliest motion detectors in the fly brain, T4 and T5 neurons. In vivo calcium imaging revealed that T4 and T5 neurons encode the sharp edges of the stationary luminance gradients, as well as those of other stationary patterns, in a complementary, contrast-specific manner. A simple model for motion perception allowed us to predictably manipulate the illusory percept by selectively silencing either T4 or T5 neurons, thus creating flies with either a stronger or reversed illusory motion percept. These results show that, in flies, the illusion arises from responses of direction-selective neurons to sharp, stationary edges, which are imperfectly cancelled by downstream subtractive mechanisms that processes light and dark signals asymmetrically. We further demonstrate that, in humans, selective adaptation of the visual system to moving light or dark edges differently affects the illusory motion percept, paralleling the

results of silencing of T4 or T5 in flies. Overall, our results suggest that the illusory motion is a manifestation of a set of convergent strategies of motion estimation across phyla.

T-26. Sensory control of neuromodulation enables flexible selection of behavioral states

Ni Ji
Gurrein Madan
Guadalupe Fabre
Alyssa Dayan
Casey Baker
Ijeoma Nwabudike
Steven flavell

SEARCHINJINNI@GMAIL.COM
GURREIN@MIT.EDU
LIZEBEL1294@GMAIL.COM
DAYANALYSSA@GMAIL.COM
CMBAKER@MIT.EDU
INWABUDIKE@GMAIL.COM
FLAVELL@MIT.EDU

Massachusetts Institute of Technology

To benefit survival, animal behavior must tune in closely to dynamical changes in the sensory world. However, behavioral states, such as foraging or mating, need to persist over long time scales to ensure proper execution. How does the nervous system balance behavioral stability with sensorimotor flexibility? To address this fundamental question, we investigated a foraging behavior in *C. elegans* where two competing behavioral states (exploration or exploitation) are jointly controlled by neuromodulation and the sensory input. Using state-of-the-art imaging technology, we monitored circuit-wide neural activity in free moving animals and uncovered robust neural correlates for the two foraging states. We then combined genetic analysis with circuit imaging to dissect the functional architecture of the circuitry. We first identified a mutual inhibitory loop through which the two foraging states are kept stable and mutually exclusive. Next, we developed a supervised machine learning approach to uncover stereotyped circuit activity patterns that predict state transitions. This approach led us to uncover key neurons that control the transition between competing circuit states. Using an ethologically relevant foraging assay, we revealed that these neurons are the critical nexus that channel dynamic sensory input to the downstream neuromodulatory circuit. Together, these results reveal a functional circuit architecture whereby two mutually inhibiting nodes receive excitatory drive from a common input pathway (see #4 in Significance Statement). We termed this circuit architecture the Incoherent Mutual Inhibition (IMI) motif. Lastly, we built a minimal model and used dynamical systems theory to reveal that, for a wide range of parameters, the IMI motif allows the circuit to rapidly switch to the appropriate state in response to an instructive input signal. Incorporated into an agent-based foraging model, the IMI circuit promotes well-timed state transitions that enhances the efficiency of the foraging behavior.

T-27. Dissecting feedforward and feedback interactions between populations of neurons

Evren Gokcen¹
Joao Semedo¹
Amin Zandvakili²
Christian Machens³
Adam Kohn²
Byron Yu¹

EGOKCEN@ANDREW.CMU.EDU
JSEMEDO@ANDREW.CMU.EDU
ZANDVAKILI@GMAIL.COM
CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
ADAM.KOHN@EINSTEIN.YU.EDU
BYRONYU@CMU.EDU

¹Carnegie Mellon University

²Albert Einstein College of Medicine

³Champalimaud Centre for the Unknown

Brain areas interact through a rich combination of feedforward and feedback signals. Yet the computational role of this bidirectional communication in perception and cognition remains largely unknown. Modern recording

techniques offer data that may elucidate that role: We can now simultaneously record many neurons across multiple brain regions. Still, disentangling inter-area communication remains a challenging problem: How do we distinguish population signals that are relayed across brain areas from those that are not? And among the relayed signals, how do we distinguish feedforward from feedback? Toward that end, we propose a novel dimensionality reduction framework, Time-Delay Gaussian Process Canonical Correlation Analysis (TD-GPCCA), which dissects the time-varying activity in each brain area into two sets of low-dimensional trajectories: (1) A within-area set, and (2) An across-area set, which can be dissected further into feedforward and feedback components. This dissection requires three tasks: dimensionality reduction, temporal smoothing, and the discovery of time-lags, or delays, between activity in each area. TD-GPCCA executes these tasks jointly. We first demonstrate that our method accurately recovers ground truth parameters and latent trajectories in realistic-scale synthetic data. Then, we use TD-GPCCA to study the interactions between simultaneously recorded populations in areas V1 and V2 of the macaque monkey. Across multiple stimuli and recording sessions, we find that (1) Nominally feedforward (V1 to V2) trajectories carry stronger stimulus-related “signal” than nominally feedback ones (V2 to V1); and (2) Virtually all signal in V2 is attributable to communication with V1, whereas V1 appears to keep a portion of its signal private. Overall, this work establishes a framework for dissecting feedforward and feedback interactions between populations of neurons in distinct brain areas.

T-28. Tracking information flow in subnetworks of mouse visual cortex

Xiaoxuan Jia
Joshua Siegle
Shawn Olsen

XIAOXUANJ@ALLENINSTITUTE.ORG
JOSH@ALLENINSTITUTE.ORG
SHAWN@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Characterizing highly recurrent information propagation between cortical ensembles requires measuring large populations of interacting neurons, combined with computational tools that can reveal the logic of these interactions. Here we simultaneously recorded the spiking activity of hundreds of neurons across six hierarchically organized regions of the mouse visual cortex (areas: V1, LM, RL, AL, PM, AM). By analyzing functional interactions between pairs of neurons, we uncovered distinct ensembles that differentially participate in feedforward and feedback processes. One ensemble (the ‘driver’ ensemble) leads network activity, whereas the other (the ‘driven’ ensemble) follows it. These ensembles are distributed across all cortical areas we studied, but ‘driver’ neurons are enriched in lower-order areas, such as V1 and LM, and ‘driven’ cells are enriched in higher-order areas, such as PM and AM.

To investigate the role of these two ensembles, we quantified their response latency, information content, and temporal precision during visual stimulation. Response latencies in the ‘driver’ ensemble systematically increased along the visual hierarchy. Strikingly, average response latency of the ‘driven’ ensemble in lower areas of the hierarchy (V1, LM) was slower than the ‘driver’ ensemble at the top of the hierarchy (AM, PM). This suggests that visual signals first propagate along the hierarchy through feedforward connections within the ‘driver’ ensemble (~3ms) followed by activation of the ‘driven’ ensemble, likely through feedback and lateral connections. We found that stimulus-related information systematically decreased along the hierarchy within the ‘driver’ ensemble, as reflected by changes in stimulus-dependent information, temporal precision, and explainable variance of neural response to sensory inputs. In general, the ‘driver’ ensemble represented more information about external stimuli than the ‘driven’ ensemble in each area. These findings show that separate functional ensembles within the same network have distinct properties and are positioned to mediate different stages of information flow during sensory processing.

T-29. Assembly structure expands the dimension of shared variability in cortical networks

Danielle Rager¹
 Sanjeev Khanna²
 Matthew Smith¹
 Brent Doiron²

DMRAGER@GMAIL.COM
 SANJEEVBKHANNA@GMAIL.COM
 MATTSMITH@CMU.EDU
 BDOIRON@PITT.EDU

¹Carnegie Mellon University

²University of Pittsburgh

Cortical circuits often receive multiple inputs from upstream populations with non-overlapping tuning preferences. Both the feedforward and recurrent architectures of the receiving cortical layer will reflect this input tuning. We study how neuronal variability propagates through a hierarchical cortical network receiving multiple, independent, tuned inputs. Micro-electrode array recordings were simultaneously collected from the same hemisphere of primary visual cortex (V4) and prefrontal cortex (PFC) of a non-human primate engaging in a two-dimensional spatial working memory task. Neurons recorded from a single PFC hemisphere preferred spatial locations distributed across both visual hemifields. A single PFC hemisphere therefore receives lateralized, disjoint inputs from both visual hemispheres. Consistent with prior studies, the shared variance of V4 activity from one hemisphere was dominated by a single latent dimension. However, significantly more latent dimensions ($3 \leq d \leq 8$) were required to explain 95 percent of the shared variance of simultaneously recorded PFC activity. To understand this dimensionality expansion between V4 and PFC, we construct a two-layer spiking network with assembly structures reflecting the laterally-tuned V4 inputs. Our network consists of two, one-dimensional V4 hemispheres projecting disjointly to neurons in a recurrent PFC layer. Excitatory (E) and inhibitory (I) PFC model neurons are assigned a visual hemifield preference and recurrently connect to neurons preferring the same hemifield with greater probability. We show that the recurrent population's ability to maintain global E/I balance in the presence of disjoint inputs depends on the recurrent clustering strength. Moreover, we demonstrate that in a moderately-clustered regime, recurrent neural activity has high-dimensional shared variability owing to a multistable solution in network activity. Restricting the dimensionality analysis to activity from one stable state recovers the low-dimensional structure inherited from V4 variability. Our model thus introduces a framework for thinking of high-dimensional cortical variability as "time-sharing" between low-dimensional, tuning-specific circuit dynamics.

T-30. Isolating a locus for reach adaptation in the cerebellar cortex

Dylan Calame
 Matt Becker
 Abigail Person
 University of Colorado

DYLAN.CALAME@CUANSCHUTZ.EDU
 MATT.BECKER@CUANSCHUTZ.EDU
 ABIGAIL.PERSON@CUANSCHUTZ.EDU

A leading hypothesis of cerebellar function is that the cerebellum generates an internal model to predict upcoming body kinematics, facilitating feedforward motor control and smooth, accurate movements. Previous work has shown that cerebellar output is graded in response to the speed of paw during reaching movements in mice, acting to decelerate the limb accurately to endpoint [1]. How upstream elements of the cerebellar circuit learn to predict where this endpoint is within the reach remains unclear. We hypothesized that predictive cues carried by cerebellar mossy fibers (MF) are used by Purkinje cells (PC) to sculpt reach kinematics. To test this hypothesis we characterized PC encoding of reach parameters, then analyzed reach adaptation to optogenetic perturbation up and downstream from PCs. PC recordings in the cerebellar primary fissure in headfixed mice performing a skilled reaching task with paw tracking showed that simple spikes (SS) are broadly modulated during reach. Using multilinear LASSO regression of limb kinematics with PC SS rate, we show that PC firing rates (FRs) can be accurately modeled by the kinematics that follow, indicating that PCs make salient predictions to guide the limb during reaching. Next, we sought to understand how perturbing PC predictions during movement affect reach trajectories and whether perturbations can be learned by the cerebellum over time. When stimulating on

every trial, early reach trajectories were acutely affected by MF stimulation, but over time the reaches adapted to match pre-stimulation trajectories. When stimulation was turned off, reaches showed aftereffects that corrected over time. By contrast, analogous optogenetic perturbation of cerebellar output does not show adaptation to stimulation but does reveal after-effects, indicating that reach adaptation takes place at sites between the MFs and cerebellar nuclei, rather than downstream brain or spinal structures.

T-31. Identical encoding of flexible and automatic motor sequences in the dorsal lateral striatum

Kevin Mizes
Bence Olveczky
Harvard University

KMIZES@G.HARVARD.EDU
OLVECZKY@FAS.HARVARD.EDU

To achieve specific goals, our motor system must quickly learn new movements, flexibly reorganize existing ones, and consolidate repeated motor sequences. Consider the feats of a concert pianist. After extensive training, the pianist can flexibly perform any number of sonatas or etudes when provided with sheet music. But to prepare for a concert, she extensively practices a single piece, making it automatic, effortless, and robust to mistakes. Although both the flexible and automatic performances can be kinematically similar, it is generally believed that they are controlled by different neural circuits (Ashby et al., 2010). In particular, flexible and automatic movements are thought to be differentiated in the dorsal lateral striatum (DLS), an essential node of the motor system and a major input to the basal ganglia. However, the specific mechanisms, as well as how these mechanisms generalize across flexible and automatic movements, remain poorly understood. One popular hypothesis is that the DLS dedicates specific circuitry to represent a well-trained sequence, forming what's called a motor chunk, and reducing cognitive load (Graybiel, 1998). In an alternative hypothesis, the striatum encodes features relevant to the ongoing movement, independent of sequence learning (Desmurget and Turner, 2010). To arbitrate between these hypotheses, we recorded from populations of single neurons in the DLS while rats performed the same movement sequence both in a flexible and automatic context. In contrast to the generally accepted view that DLS differentiates flexible and automatic movements, we found that neural representation of similar movement kinematics was invariant to the context (i.e. automatic or flexible). Further characterization of the encoding scheme in the DLS suggested that units primarily represent the single motor elements, and not any sequence specific or ordinal position information. These results are consistent with a low-level kinematic role of DLS independent of automaticity, cognitive load, or behavioral flexibility.

T-32. Exploring the hidden state of Coma

Sima Mofakham¹
Adam Fry²
Josue Nassar³
Joseph Adachi⁴
Charles Mikell¹

SIMA.MOFAKHAM@STONYBROOKMEDICINE.EDU
ADAM.FRY@MOUNTSINAI.ORG
JOSUE.NASSAR@STONYBROOK.EDU
JOSEPH.ADACHI@STONYBROOKMEDICINE.EDU
CHARLES.MIKELL@STONYBROOKMEDICINE.EDU

¹Stony Brook University Hospital

²Icahn School of Medicine at Mount Sinai

³Stony Brook University

⁴Stony Brook University School of Medicine

Severe traumatic brain injury (TBI) normally results in loss of consciousness or coma. Treatment of coma following TBI is limited primarily by a lack of understanding of the underlying mechanisms supporting consciousness. Several functional imaging studies have identified thalamic coordinated activity across cortical networks as the seat of consciousness, however, direct recording data following TBI are sparse. To address this gap, under our

approved IRB, we performed direct depth electrode recordings of the FPN in comatose patients following TBI. We utilized time-frequency analysis along with dynamical system methodology to characterize the latent state of FPN dynamics, and its evolution as patients regain consciousness. We found that the complexity of these electrocorticography (ECoG) recordings correlate with the level of consciousness after TBI as well as patient recovery. On the other hand, our imaging studies revealed that the complexity and variability of the ECoG has a direct link to the integrity of the thalamocortical loop. We found that the underlying dynamical system governing cortical activity (depth cortical contact in PFC) in the patient with bilateral thalamic injury, who did not recover consciousness, was confined within a periodic structured. We characterized the extent to which single-pulse stimulation changed the trajectory of this attractor. Conversely, the phase-space in the healthy control and the recovering comatose patient were high dimensional, high entropy and unpredictable. Understanding the characteristics of this attractor and its response to stimulation can guide future neuromodulatory therapies to facilitate the return of consciousness.

T-33. Deep neural network modeling of visuomotor transformations during social interaction

Adam Calhoun
Benjamin Cowley
Jonathan Pillow
Mala Murthy
Princeton University

ADAM.CALHOUN@GMAIL.COM
BCOWLEY@PRINCETON.EDU
PILLOW@PRINCETON.EDU
MMURTHY@PRINCETON.EDU

Visual cues in the environment are high-dimensional and complex, and animals must, in general, parse these cues to drive responsive behaviors. Here we investigate coding at a visual bottleneck in the *Drosophila melanogaster* brain during natural behavior. During courtship, male flies chase female flies and use dynamic visual cues to pattern their locomotion and song production behaviors. These cues are processed via neurons that project from the optic lobe to the brain, and form a small number of processing channels. Previous work suggests that each channel carries information about a particular aspect of visual motion (Wu et al., eLife 2016). However, we lack a framework to determine how many visual features an animal is using during a particular behavior. To solve this problem, we trained a deep neural network (DNN) to model the mapping from pixel-level visual inputs to behavior. Prior work had shown that a small number of abstract features about the female (e.g., distance) could explain much of the variation in song behavior (Coen et al., Nature 2014), leading us to expect the model to use only a few channels to embed visual input. Surprisingly, by silencing each embedding channel separately, we found that many individual channels contribute to behavior. To confirm these predictions in the fly brain, we silenced each of the 22 channels (sets of lobula columnar neurons) that form a bottleneck between the eye and the central brain of *Drosophila*. Surprisingly, we found that rather than using only one or two channels as would be expected from other behaviors in *Drosophila*, courtship behavior required almost all visual channels, consistent with the DNN. Overall, by modeling behavior from sensory input to motor output, our framework illustrates how the brain extracts useful sensory features to guide complex behavior.

T-34. Multi-region network models of brain-wide interactions during adaptive and maladaptive state transitions

Matthew Perich¹
Aaron Andalman²
Eugene Carter¹
Kanaka Rajan¹
Karl Deisseroth²

MPERICH@GMAIL.COM
AARON.ANDALMAN@GMAIL.COM
GANANDA@GMAIL.COM
KANAKA.RAJAN@MSSM.EDU
DEISSERO@STANFORD.EDU

¹Icahn School of Medicine at Mt. Sinai

²Stanford University

During prolonged periods of stress, animals can switch from active to passive coping strategies to manage effort-expenditure. Such normally adaptive behavioral state transitions can become maladaptive in psychiatric disorders such as depression. Here, we studied the neural mechanisms underlying behavioral state transitions using whole-brain, cellular-resolution neural recordings from larval zebrafish during maladaptive passivity induced by inescapable stress. We developed multi-region recurrent neural network (RNN) models to infer brain-wide interactions driving such maladaptive behavior. Models were constrained by training to match experimental data across two levels simultaneously: (a) large-scale neural dynamics from more than 10,000 neurons and (b) the behavior of the organism (tail movements). Analysis of the trained RNN models revealed a specific change in inter-area connectivity between the habenula (Hb) and raphe nucleus – two regions previously, and individually, implicated in maladaptive passivity – in response to inescapable stress. Next, we analyzed the dynamics within the neural manifolds capturing Hb and raphe population activity. We defined a subspace within the Hb activity that captured communication with the raphe nucleus (Hb-raphe subspace). At the start of the stressful experience, the Hb-raphe subspace responded immediately; however, RNN models identified no early or fast-timescale change in the interactions between these regions. Later in the stressful experience, as the animal lapsed into passivity, the responses within the Hb-raphe subspace decreased. This later transition was accompanied by a concomitant change in the effective inter-area interactions between the raphe and Hb inferred from the RNN weights. This innovative combination of network modeling and population dynamics analysis points to dual mechanisms with distinct timescales driving the adaptive to maladaptive behavioral state transition: early behavioral response to stress could be mediated by reshaping the neural dynamics within a preserved network architecture, while long-term state changes correspond to altered connectivity between brain regions.

T-35. Accurate angular integration with only a handful of neurons

Marcella Noorman
Vivek Jayaraman
Sandro Romani
Ann Hermundstad

NOORMANM@JANELIA.HHMI.ORG
VIVEK@JANELIA.HHMI.ORG
ROMANIS@JANELIA.HHMI.ORG
HERMUNDSTADA@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

Ring attractor networks have been studied theoretically for over two decades as a putative model of the mammalian head-direction system ([1,2]; fig 1A). Such networks rely on large numbers of neurons to accurately integrate angular velocity in order to maintain a precise representation of heading. However, recent findings in the fruitfly *Drosophila melanogaster* suggest that this computation can be robustly carried out by a small network consisting of relatively few neurons ([3-6]; fig 1B). These findings call into question our current understanding of continuous attractors and their implementation in small circuits.

In this work, we dissect the performance of a ring attractor model of the *Drosophila* heading network [5], with a focus on the model's shortcomings. In particular, we show how the model's failures in integration—including the inability to integrate small velocity input, variable integration of constant velocity input, and drift in the absence of velocity input—can all arise in the limit of small numbers of neurons. Motivated by the observation that these failures are seldom seen in a walking fly's heading network [3,5], we then mathematically derive conditions under which such a constrained system can overcome these shortcomings and achieve the accuracy of the unconstrained system. Together, these results show how angular velocity integration can be accurately performed by small networks. When combined with electron-microscopy-based reconstruction of the cellular- and synaptic-resolution connectome of the fly heading network, together with ongoing physiology experiments, these results enable us to test whether and how such conditions for accurate performance are achieved mechanistically in the real circuit.

T-36. Preexisting hippocampal network dynamics constrain optogenetically induced place fields

Sam McKenzie¹
 Gyorgy Buzsaki²
 Roman Huszar²
 Daniel English³
 Kanghwan Kim⁴
 Euisik Yoon⁴

SAM.A.MCKENZIE@GMAIL.COM
 GYORGY.BUZSAKI@NYUMC.ORG
 RH2618@NYU.EDU
 NEURODAN83@VT.EDU
 KHANKIM@UMICH.EDU
 ESYOON@UMICH.EDU

¹New York University Langone Medical Center

²New York University

³Virginia Tech

⁴University of Michigan

Neuronal circuits face a fundamental tension between maintaining existing structure and changing to accommodate new information. Memory models often emphasize the need to encode novel patterns of neural activity imposed by “bottom-up” sensory drive. In such models, learning is achieved through synaptic alterations, a process which potentially interferes with previously stored knowledge. Alternatively, neuronal circuits generate and maintain a preconfigured stable dynamic, sometimes referred to as an attractor, manifold, or schema, with a large reservoir of patterns available for matching with novel experiences. Here, we show that incorporation of arbitrary signals is constrained by pre-existing circuit dynamics. We optogenetically stimulated small groups of hippocampal neurons as mice traversed a chosen segment of a linear track, mimicking the emergence of place fields, while simultaneously recording the activity of stimulated and non-stimulated neighboring cells. Stimulation of principal neurons in CA1, but less so CA3 or the dentate gyrus, induced persistent place field remapping. Novel place fields emerged in both stimulated and non-stimulated neurons, which could be predicted from sporadic firing in the new place field location and the temporal relationship to peer neurons prior to the optogenetic perturbation. Circuit modification was reflected by altered spike transmission between connected pyramidal cell – inhibitory interneuron pairs, which persisted during post-experience sleep. We hypothesize that optogenetic perturbation unmasked sub-threshold, pre-existing place fields. Plasticity in recurrent/lateral inhibition may drive learning through rapid exploration of existing states.

T-37. A map of object space in primate inferotemporal cortex

Pinglei Bao
 Liang She
 Mason McGill
 Doris Tsao

PLBAO.USC@GMAIL.COM
 LIANGSHE@CALTECH.EDU
 MASON.B.MCGILL@GMAIL.COM
 DORTSAO@CALTECH.EDU

California Institute of Technology

How is the representation of complex visual objects organized in inferotemporal (IT) cortex, the brain region responsible for object recognition? Areas selective for a few categories such as faces, bodies, and scenes have been found, but large parts of IT lack any known specialization, leading to uncertainty over whether any general principle governs IT organization. Here, we used fMRI, microstimulation, electrophysiology, and deep networks to investigate the organization of macaque IT. We built a low dimensional object space to describe general objects using a deep network. Responses of IT cells to a large set of objects revealed that single IT cells are projecting incoming objects onto specific axes of this space. Remarkably, cells were anatomically clustered into four networks according to the first two components of their preferred axes, forming a map of object space. This map was repeated across three hierarchical stages of increasing view invariance, and cells comprising these maps collectively harbored sufficient coding capacity to reconstruct arbitrary objects. These results provide a unified picture of IT organization in which category-selective regions are part of a coarse map of object space.

T-38. A map for odors and place in posterior piriform cortex

Cindy Poo¹
 Gautam Agarwal²
 Niccolo Bonnachi³
 Zachary Mainen³

CINDYPOO@GMAIL.COM
 AGARWALLED@GMAIL.COM
 NICCOLO.BONACCHI@NEURO.FCHAMPALIMAUD.ORG
 ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

¹Champalimaud Centre for the Unknown
²University of California, Berkeley
³Champalimaud Foundation

Animals rely on olfaction for a wide range of natural behaviors. To support behaviors such as foraging and navigation, the brain must integrate odor information with an internal model of the spatial environment to produce flexible behavior. In order to understand neural circuits underlying the integration of sensory and spatial information, we designed an odor-cued spatial choice task for freely-moving rats, where rats use odor information for allocentric navigation. The highly recurrent and plastic nature of primary olfactory (piriform) cortex (PCx), together with abundant connections with higher order structures, make PCx an excellent site to investigate interactions between sensory information and associational or cognitive processes. We recorded from posterior PCx (pPCx) in rats with stable task performances (n=995 neurons, 3 rats, 44 sessions). We found that individual pPCx neurons were not only odor-selective, but also fired differentially to the same odor sampled at different locations, forming an "odor-place map". Spatial locations could be successfully decoded from simultaneously recorded pPCx population. Strikingly, we found that location representation in pPCx did not depend on direct sensory drive and was maintained across behavioral contexts. Using simultaneous recordings in hippocampus (HPC) and coherence analysis of pPCx spike times and HPC local field potentials (LFPs), we found that HPC projections preferentially drive those pPCx neurons with higher spatial selectivity. Together, our results reveal an associative odor-place map in pPCx that is well-suited to guide olfactory navigation, and challenge the conventional view of how spatial information is represented in the brain. A cognitive map of space in primary sensory cortex allows incoming sensory evidence to directly update an internal model, and could be critical for a wide range of spatial behaviors where seamless and dynamic integration of sensory information and cognitive maps are critical.

T-39. Rapid representational drift in primary olfactory cortex

Carl Schoonover
 Andrew fink
 Richard Axel
 Columbia University

CES2001@COLUMBIA.EDU
 AF2243@COLUMBIA.EDU
 RA27@COLUMBIA.EDU

Primary olfactory (piriform) cortex has been traditionally hypothesized to facilitate the identification of odorants. A region whose principal role is to support stimulus identification should produce a readout that varies little over the animal's lifetime. For example, absent learning-induced plasticity, odorant representations in the olfactory bulb (the principal input to piriform) are stable over several weeks. Similarly, sensory maps in primary neocortical regions, like retinotopy in V1, are grossly stable. We asked whether piriform, which lacks an analogous sensory map, exhibits comparable stability. For this we developed a strategy to achieve long-term recordings from fixed populations of single-units in anterior piriform and measured odor tuning across weeks. We have made three observations. (1) Odorant representations are profoundly reorganized on a timescale of weeks: a linear classifier trained using responses recorded on one day performs nearly at chance when tested on responses recorded one month later. (2) Odorant representations are stabilized if the animal has frequent, regular experience with the stimuli. But (3) this stabilization lasts only so long as regular experience is maintained; otherwise, previously stabilized representations become labile once again. Piriform therefore seems poorly suited for facilitating odor identification over an organism's lifetime. Instead it exhibits key features of a fast learning system: regular, non-reinforced, experience is sufficient to produce a memory trace (a stable representation), however, this trace is not burnt in and is rapidly overwritten absent frequent experience. We propose that piriform is a highly plastic

system that learns statistical regularities in the olfactory environment but lacks the capacity for long-term storage of that information. This view produces two predictions that we are currently testing: (1) odor identification over the animal's lifetime is implemented in a parallel system without requiring piriform, and (2) a slow learning system exists downstream to store long-term memories initially encoded in piriform.

T-40. Manipulating synthetic optogenetic odors reveals the coding logic of olfactory perception

Edmund Chong¹
 Monica Moroni^{2,3}
 Christopher Wilson¹
 Shy Shoham¹
 Stefano Panzeri²
 Dmitry Rinberg¹

EDMUND.W.CHONG@GMAIL.COM
 MONICA.MORONI@IIT.IT
 CHRISTHEWILSON@GMAIL.COM
 SHOHAM@NYU.EDU
 STEFANO.PANZERI@IIT.IT
 RINBERG@NYU.EDU

¹New York University

²Istituto Italiano di Tecnologia

³University of Trento

How does neural activity generate perception? The spatial identities and temporal latencies of activated units correlate with external sensory features, but finding the subspace of activity that is consequential for perception, remains challenging. We trained mice to recognize synthetic odors: optogenetically-driven spatiotemporal patterns of glomerular activity in the olfactory bulb. We then performed precise spatial or temporal perturbations on trained patterns and measured how recognition changes. Changes in recognition reflect the perceptual relevance of the modified feature (or groupings of features). We modeled recognition as the matching of glomerular activity to learned templates, and uncovered what forms a perceptually-meaningful pattern template: activation sequences ordered by latencies relative to each other, with surprisingly minimal effect of sniff. Within templates, spatially-identified glomeruli contribute additively, with larger contributions from earlier-activated glomeruli. Template matching with these perceptually-meaningful features can account for animals' responses, with the degree of mismatch predicting changes in recognition. The model accurately generalizes to novel spatio-temporal manipulations of patterns, and produces non-linear responses that resemble the non-linear responses in the data. This is the first report to our knowledge, that not only establishes a causal role for neural activity sequences in perception, but also uncovers the perceptually-relevant coding schemes governing these sequences. Hence, by performing precise and parametric manipulation of glomerular activity and quantifying effects under a common metric, we derived a unifying model that explains how odor perception arises from structured glomerular activation. Our synthetic approach reveals the fundamental logic of the olfactory code, and provides a general framework for testing links between sensory activity and perception.

T-41. A neuro-economic model for rats' willingness to work for water

Pamela Reinagel
 University of California, San Diego

PREINAGEL@UCSD.EDU

Given alternatives, animals choose options with larger rewards, and work harder when those options are offered. But in nature a resource such as water may be scarce for long times, so animals must be willing to work harder precisely when the reward for work is smaller. Mechanisms for ensuring sufficient but not excessive effort in this context are poorly understood. We investigated this in rats performing work to obtain water. Rats had access to the task 24 hours/day but no other hydration source. Task difficulty was fixed and the reward size (ml/trial) was varied. Only one reward size was offered at any given time, and this value was stable for weeks. When reward size was stably low, rats did more total work per day – even though this required investing more effort for smaller

rewards. The simplest explanation would be that rats consume the same quantity of water per day, regardless of its price in effort, but this was not the case. When the cost of water was high, rats consumed enough for health, but when the cost was low, rats consumed twice that amount or more. These results mirror “back-bending labor supply curves” and “price elasticity of demand” in classical economics, and can be explained by Utility Maximization. We present a model in which the “Net Marginal Utility” in this theoretical framework is interpreted as a time-varying signal in the brain that causally drives a switch from work to rest state; this stochastic model replicates the temporal dynamics of trials. Our behavioral data allow us to fit the utility functions in detail. The curve we obtain for the Marginal Utility of Water does not track physiological hydration, but rather bears a striking resemblance to the activity of neurons in SFO and MnPO, which we suggest encode this variable.

T-42. The role of frontal cortex in multisensory decision-making

Philip Coen
Miles Wells
Peter Zatka-Haas
Matteo Carandini
Kenneth Harris
University College London

PIPCOEN@GMAIL.COM
MILES.WELLS@UCL.AC.UK
P.ZATKA@UCL.AC.UK
M.CARANDINI@UCL.AC.UK
KENNETH.HARRIS@UCL.AC.UK

The ability to combine visual and auditory cues to localize objects in space is critical to many organisms, whether prey, predator, or pedestrian crossing the street. Inactivation experiments in cat have suggested that cortex is critical for audiovisual spatial localization, but the temporal precision and reversibility of these inactivations were limited by the techniques available. By using mice, we can overcome these limitations by pairing rigorous psychophysics with comprehensive and reversible inactivation. But recent work suggests that mice do not perform optimal audiovisual integration, displaying unimodal dominance when two cues are presented in conflict. Do mice combine audiovisual cues independently in a spatial context, and do they use cortex to do so?

To answer this question, we developed a two-alternative forced choice task where head-fixed mice turn a wheel to indicate whether a stimulus appeared on the left or right. The stimuli can be auditory, visual, or a combination of the two, presented in coherent or conflicting locations. In this task, mice showed no unisensory dominance, and integrated auditory and visual cues independently in all conditions. In fact, a simple additive model, with independent weights for audio and visual stimuli, completely predicted the mean response of mice to all cue combinations. We proceeded to optogenetically inactivate different spots across dorsal cortex on individual trials while mice performed this task. Perturbations of the additive model were analyzed with a hierarchical Bayesian approach and identified distinct roles for different cortical regions. Finally, we recorded over 2000 neurons in these regions to determine how audiovisual information is encoded at the population level during behavior.

Our results demonstrate that mice independently combine audiovisual cues during spatial localization, indicating that the underlying neural framework is homologous to other mammals. Further, we establish that different regions of cortex have distinct roles in this multimodal decision process.

T-43. Task representation in evoked and spontaneous activity in the visual cortex

Marton Hajnal¹
 Gergo Orban²
 Peyman Golshani³
 Duy T Tran³
 Michael C Einstein³
 Pierre-Olivier Polack⁴

HAJNAL.MARTON@WIGNER.MTA.HU
 ORBAN.GERGO@WIGNER.MTA.HU
 PGOLSHANI@MEDNET.UCLA.EDU
 DUYTRAN1125@UCLA.EDU
 MICKY.EINSTEIN@GMAIL.COM
 POLACK.PO@RUTGERS.EDU

¹HAS Wigner Research Centre for Physics

²Wigner RCP

³University of California, Los Angeles

⁴Rutgers University

Sensory areas of the cortex display rich activity even without stimulus presentation, which has been linked to the representation of expectations. Recent work, however, demonstrated that a large portion of the activity recorded in a passive animal when stimulus is absent is related to behavioural activity. Here we argue that in task engaged animals a task-relevant variable is reliably represented in the intertrial interval, when stimulus is not presented. We trained mice to perform a pair of tasks where the same set of multi-modal stimuli were used in different contexts: in one task context the animals were expected to make decisions based on visual stimulus content and ignore the auditory stimulus and in the other roles were reversed, thus decoupling task context, stimulus identity, and behavioural outcome. Unit recordings were obtained from V1. We found that task context was represented not only during stimulus presentation but a task representation consistent with on-stimulus activity was present during pre-stimulus activity as well. Representation of task context was achieved through recruiting a population overlapping with the one representing the visual stimulus but establishing an orthogonal representational space. This task-related activity was a major component of the activity recorded without stimulus and spanned only a low dimensional subspace. Context was invariantly represented in the same subspace throughout the session while transitioning to the other context. The task design enabled us to assess choice related activity independently. We found a representation of choice that emerged early in the trial, prior to actual behaviour and this choice representation resided again in a linearly independent subspace of the population activity space. Taken together, our results demonstrate that a structured activity pattern in V1 that is not related to the stimulus but is consistent whether stimulus is or is not presented reflects cognitive variables relevant to task execution.

I-1. Chimera states as epileptic seizure predictors

Nuttida Rungratsameetaweemana¹
 Claudia Lainscsek²
 Sydney S Cash^{3,4}
 Terrence Sejnowski²

NUTTIDA.RUNGRAT@GMAIL.COM
 CLAUDIA@SALK.EDU
 SCASH@MGH.HARVARD.EDU
 TERRY@SALK.EDU

¹University of California, San Diego

²Salk Institute for Biological Studies

³Massachusetts General Hospital

⁴Harvard University

Worldwide, more than 50 million people live with epilepsy, a common neurological disorder characterized by hyper-synchronous abnormalities of neurons. Precise seizure characterization is key to developing a successful intervention as seizure predictions are challenging. One way to further our understanding of epilepsy is through examining the transition from desynchronized to highly-synchronized cortical states. It has been proposed that in the brain, a hybrid dynamic or chimera state where both coherent and incoherent domains coexist may also be present. Through computational simulation, past studies have demonstrated that these states can be observed in a variety of dynamical systems across a wide range of coupling topologies and coupling strengths [1]. Based on

these findings, the studies proposed a possible link between chimera states and epileptic seizures. The present study investigated this question and specifically hypothesized that the incoherence-coherence transition that leads to epileptic seizures can be explained by the induction and dynamics of chimera-like states. To test this, we first examined the nature of chimera states by systematically analyzing their spatio-temporal dynamics in two networks: 1) chaotic Rossler oscillators, and 2) Fitzhugh-Nagumo (FN) neurons. While the Rossler system was chosen due to its relative simplicity in between-unit interactions, the FN model sufficiently describes the process of activation and deactivation dynamics of a spiking neuron despite being computationally tractable [2]. The second part of the present study investigated chimera states in electrocorticography (ECoG) data from patients with intractable focal epilepsy. In one patient, we found that each seizure was reliably preceded by chimera states. Therefore, such states were 100% seizure predictive. Additionally, our analyses revealed that the brain stayed in these chimera states for up to two hours prior to a seizure. This is, to the best of our knowledge, the first direct evidence of chimera states in the human epileptic brain.

I-2. Measuring and modeling a fine spatial scale clustering of orientation tuning in L2/3 of mouse V1

Peijia Yu¹
 Yuhan Yang²
 Daisong Pan²
 Jiang Lan Fan²
 Chengcheng Huang³
 Mario Dipoppa⁴
 Kenneth Miller⁴
 Na Ji²
 Brent Doiron³

PEIJIA.YU@ANDREW.CMU.EDU
 YUHANY@BERKELEY.EDU
 DAISONG_PAN@BERKELEY.EDU
 KEVINFAN@BERKELEY.EDU
 HUANGC@PITT.EDU
 MD3681@COLUMBIA.EDU
 KEN@NEUROTHEORY.COLUMBIA.EDU
 JINA@BERKELEY.EDU
 BDOIRON@PITT.EDU

¹Carnegie Mellon University

²University of California, Berkeley

³University of Pittsburgh

⁴Columbia University

The absence of spatial organization in orientation tuning had been thought as a major feature of the rodent primary visual cortex (V1). However, recent experimental discoveries have been revisiting and challenging this view. Population imaging studies have suggested that nearby neurons in the layer 2/3 (L2/3) of mouse V1 tend to have stronger tuning similarity than that of distant neuron pairs, indicating a localized spatial clustering of stimulus feature preference (Ringach et al. 2016, Jimenez et al. 2018, Kondo et al. 2016). However, the spatial scale of clustering is still in debate: either spread over hundreds of microns (Ringach et al. 2016), or limited to the scale of tens of microns (Kondo et al. 2016). Those differences could reflect distinct scales of local feedforward/recurrent cortical connectivity, so an accurate measurement of the spatial profile of local clustering will shed light on the underlying neuronal circuits, yielding way to circuit-based mechanisms of visual processing in rodent V1. Here using two-photon calcium imaging, we measured the orientation tuning properties of L2/3 neurons in mouse V1. We found a significant spatial clustering of tuning, but horizontally localized in only approximately $20\mu m$, which is typically the average distance between horizontally neighboring neurons. To understand this narrow clustering, we explored a spiking neuron network model of L2/3 and L4 of mouse V1. Building on past models with broad recurrent wiring over $\sim 200\mu m$ (Rosenbaum et al., 2017; Huang et al., 2019) we additionally considered an excess connecting probability over a narrow $20\mu m$ range. A spatially narrow local tuning similarity matching our data emerges for even weak narrow connectivity, effectively adding only a few extra local connections per neuron. Our combined experimental and modeling work argue for a fine spatial scale of wiring between adjacent neurons in mouse V1.

I-3. Computational cognitive requirements of random decision problems

Pablo Franco
 Karlo Doroc
 Nitin Yadav
 Peter Bossaerts
 Carsten Murawski
 University of Melbourne

JPFRANCO123@GMAIL.COM
 KARLO.DOROC@UNIMELB.EDU.AU
 NITIN.YADAV@UNIMELB.EDU.AU
 PETER.BOSSAERTS@UNIMELB.EDU.AU
 CARSTENM@UNIMELB.EDU.AU

The decisions people face in their everyday lives involve solving problems of varying computational cognitive requirements. Computational resource requirements have been studied widely using computational complexity theory. However, it is an open question whether this framework applies to human decision-makers. Here we propose and extend a computational complexity measure of the expected computational requirements of solving a random instance (i.e. random case) of a decision problem. This measure, which we refer to as instance complexity (IC), quantifies cognitive resources required to make a decision based on a small number of features of the instance. It has been shown that IC is related to computational resource requirements, in particular, compute time, of electronic computers. In this study, we test whether IC also applied to humans. More specifically, we investigate whether IC is a generalisable measure, for humans, of the expected computational requirements of solving a problem. For this purpose, we conducted a set of experiments in which participants solved a set of instances of one of three widely studied computational problems, Traveling Salesperson, Knapsack and Boolean Satisfiability. Instances varied in their IC. We show that participants expended more effort on instances with higher IC but that decision quality was lower in those instances. Together, our results suggest that IC can be used to measure the expected computational requirements of solving random instances of a problem, based on an instance's mathematical properties. Our results provide further support of the conjecture that computational complexity is inherent to computational problems. Moreover, our results generate future avenues for research, based on IC, that could play a crucial role in understanding the cognitive resource allocation process in the brain.

I-4. Mechanisms of top-down attentional control in thalamic reticular circuits

Qinglong Gu¹
 Norman Lam²
 John D Murray²
¹Yale School of Medicine
²Yale University

QINGLONG.GU@YALE.EDU
 NORMAN.LAM@YALE.EDU
 JOHN.MURRAY@YALE.EDU

The thalamus serves as a relay of peripheral sensory signals to cortex, but a growing literature has also characterized its flexible engagement in higher-order cognitive functions according to task demands. For instance, attentional goals can exert top-down control of thalamic gain to amplify task-relevant signals and filter task-irrelevant signals depending on cognitive needs. However, the dynamical circuit mechanisms supporting this cognitive flexibility are poorly understood. Here, we developed a biophysically-constrained computational model of a thalamic-reticular circuit, composed of thalamocortical and TRN spiking neurons, and applied it to model top-down attentional modulation in cognitive task paradigms. Circuit dynamics arise through an interplay between the active neuronal conductances and synaptic excitatory-inhibitory interactions. Synaptic parameters are constrained by experiments, with the circuit operating in an awake in-vivo dynamical regime exhibiting asynchronous spontaneous activity and spindle activity evoked by transient perturbations. We found that the model can capture a number of electrophysiological features of thalamic-reticular circuits in awake behaving animals performing cognitive tasks involving flexible top-down control of attention on sensory processing. For example, a task can provide trial-by-trial attentional cues for the subject to base their response on stimuli from one modality (vision vs. audition) while ignoring the distractor modality. Specifically, we found increasing top-down inhibition to TRN can disinhibit response gain of thalamus, thereby improving the ability of downstream regions to detect and discriminate sensory stimuli. Conversely, decreasing thalamic gain via TRN inhibition can improve task performance

by filtering out the task-irrelevant distractor modality. Furthermore, we characterized how inhibitory dysfunction in the model impacts neural dynamics and cognitive behavior, which is relevant to multiple neuropsychiatric disorders and pharmacology. In summary, our biophysically-based thalamic-reticular circuit model reproduces multiple experimental phenomena including engagement in flexible attentional modulation of sensory processing. The model makes dissociable predictions for electrophysiological and perturbative experiments studying of cognitive function.

I-5. Time-dependent simultaneous configural and elemental odorant mixture processing

Tingkai Liu
Chung-Heng Yeh
Aurel A Lazar
Columbia University

TL2747@COLUMBIA.EDU
C.YEH@COLUMBIA.EDU
AUREL@EE.COLUMBIA.EDU

Natural odorant scenes are complex landscapes comprising mixtures of volatile compounds. Successful navigation in such environments requires two seemingly incompatible mixture processing schemes: 1) configural processing in which mixtures are perceived as a holistic odor object, and 2) elemental processing in which components of a mixture are individually identifiable. By modeling the *Drosophila Melanogaster* olfactory system with biophysical neuron and synapse models, we advanced a feedback gain control circuit architecture that reconciles the two schemes via spatio-temporal dynamics. Particularly, we investigated the PN-KC-APL circuit formed by Projection Neurons (PNs), Kenyon Cells (KCs) and the Anterior Paired Lateral (APL) neuron. Our model transforms the PN combinatorial code into a high-dimensional sparse KC representation via 1) an expansion by random projections, 2) the KC spiking non-linearity and 3) the APL feedback gain control. We modeled pure odorant stimuli as concentration-modulated affinity vectors, and mixtures as weighted summations between pure odorant components. We show that, due to the feedback gain control exerted by the APL neuron and the KC spiking mechanism, the steady-state KC output is an input-independent sparse combinatorial code with consistently 5-10% active channels, leading to a configural code which exhibits the same computational power in similarity search tasks as models from previous works. Importantly, preceding the steady-state phase, the circuit exploits the transient gradient encoding in previous layers of the olfactory pathway, as well as the APL temporal dynamics to encode mixture components elementally. This elemental encoding phase further facilitates attention-driven active sampling, where identities of odorant components recovered during the elemental phase drive selective feedback inhibition on KC channels, thereby accentuating representation of a particular component in steady-state. Our model demonstrates that the time-dependent simultaneous elemental and configural mixture processing can be traced back to the spatio-temporal dynamics of the early olfactory circuit.

I-6. Inferring low-dimensional latent descriptions of animal vocalizations

Jack Goffinet
John Pearson
Richard Mooney
Duke University

JACK.GOFFINET@DUKE.EDU
JOHN.PEARSON@DUKE.EDU
MOONEY@NEURO.DUKE.EDU

Vocalization is an essential medium for social and sexual signaling in most birds and mammals. Consequently, the analysis of vocal behavior is of great interest to fields such as neuroscience and linguistics. A standard approach to analyzing vocalization involves segmenting the sound stream into discrete vocal elements, calculating a number of handpicked acoustic features, and then using the feature values for subsequent quantitative analysis. While this approach has proven powerful, it suffers from several crucial limitations: first, handpicked acoustic features may miss important dimensions of variability that are important for communicative function. Second,

many analyses assume vocalizations fall into discrete vocal categories, often without rigorous justification. Third, a syllable-level analysis requires a consistent definition of syllable boundaries, which is often difficult to maintain in practice and limits the sorts of structure one can find in the data. To address these shortcomings, we apply a data-driven approach based on the variational autoencoder (VAE), an unsupervised learning method, to the task of characterizing vocalizations in two model species: the laboratory mouse (*Mus musculus*) and the zebra finch (*Taeniopygia guttata*). We find that the VAE converges on a parsimonious representation of vocal behavior that outperforms handpicked acoustic features on a variety of common analysis tasks, including representing acoustic similarity and recovering a known effect of social context on birdsong. Additionally, we use our learned acoustic features to argue against the widespread view that mouse ultrasonic vocalizations form discrete syllable categories. Lastly, we present a novel “shotgun VAE” that can quantify moment-by-moment variability in vocalizations. In all, we show that data-derived acoustic features confirm and extend existing approaches while offering distinct advantages in several critical applications.

I-7. Model metamers reveal that deep neural network invariances differ from human perceptual invariances

Jenelle Feather
Alex Durango
Ray Gonzalez
Josh McDermott

JFEATHER@MIT.EDU
DURANGO@MIT.EDU
RAYGON@MIT.EDU
JHM@MIT.EDU

Massachusetts Institute of Technology

Deep neural networks currently provide the best predictive models of the visual and auditory systems. The success of these models is often attributed to learned representational transformations that instantiate invariances similar to those of human perceptual systems, plausibly reflecting shared constraints from natural recognition tasks. Here, we formalize a behavioral measure to directly test whether the invariances learned by artificial neural networks match those of human perception. We synthesized “model metamers” – physically distinct stimuli that produce the same activations at a particular stage of a task-optimized neural network. To generate model metamers we performed gradient descent on a noise signal to minimize the difference between the activations it produced (at a particular model stage) and those produced by a natural stimulus. By definition, model metamers for one stage of a network produce the same activations at all subsequent model stages, and the same decision, and are thus equally recognizable to the model. We then measured whether humans could recognize the synthetic metamers as well as they recognize the corresponding natural signal, as one would expect if the model invariances mirror those of biological sensory systems. Although model metamers for early layers were recognizable to humans, metamers for late model stages were consistently unrecognizable. This general result held across several different network architectures and across image and audio recognition tasks. However, metamers could be made more recognizable with architectural modifications that reduced aliasing from pooling operations, providing hope that additional biological or engineering constraints might force model invariances to better conform to those of biological sensory systems. The results show that despite replicating aspects of behavior and neural responses, present-day deep neural networks learn invariances that deviate markedly from those of biological sensory systems. Metamers may help guide future model refinements to reduce or eliminate these discrepancies.

I-8. Prospective representation of navigational goals in orbitofrontal cortex

Raunak Basu
Hiroshi Ito

RAUNAKBASU.IITKGP@GMAIL.COM
HIROSHI.ITO@BRAIN.MPG.DE

Max Planck Institute for Brain Research

Two critical computations that animals require for goal-directed navigation are evaluation of their current location

and decision regarding the next destination. Previous research has identified cells in hippocampus and surrounding cortices that fire when rats visit specific locations in an environment, such as place, grid, or border cells, providing a rich representation of the rat's current location. Several studies have also demonstrated that these neurons in hippocampus and entorhinal cortex sometimes exhibit brief spike sequences that correlate with the animal's planned trajectories from current location to future locations. However, generation of such trajectory sequences may require distinct representation of goal location in the brain, and where and how that goal is represented has been a long-standing question in neuroscience. We reasoned that goal representation is generated from decision-making processes in the brain as the animal tries to choose the most valuable location out of available options. To test this hypothesis, we designed a linear track with multiple reward sites, in which rats are required to alternate between two given sites to obtain rewards. After three consecutive successful alternation trials between the two given sites, the reward sites are changed to new locations, thereby ensuring continuous update of goal representation. We recorded from neurons in the orbitofrontal cortex (OFC), a frontal cortical region critical for value-based decision making, when rats performed this task. Our results pointed to two key roles for OFC in navigation. First, OFC neurons differentially encode individual reward locations, providing information about the animal's current reward site. Second, OFC neural population representation switch from the animal's current location to the future destination before the animal starts its journey towards the goal. As further support of this idea, we also found that temporary silencing of OFC results in significant impairment of the animal's performance in this navigation task.

I-9. Different types of plasticity and interneurons shape circuit dynamics after sensory deprivation

Leonidas Richter
Julijana Gjorgjieva

LEONIDAS.RICHTER@BRAIN.MPG.DE
GJORGJIEVA@BRAIN.MPG.DE

Max Planck Institute for Brain Research

Diverse plasticity mechanisms acting in a parallel and orchestrated manner shape cortical circuit dynamics. Especially during the critical period, manipulating sensory experience can induce a high degree of plasticity. Monocular deprivation, the sustained closure of one eye, induces profound synaptic changes in feedforward and recurrent circuits of the primary visual cortex already after two consecutive days. We investigated deprivation-induced plasticity in model networks of spiking neurons with one or two subtypes of inhibitory interneurons. We parametrically investigated the effects of multiple experimentally observed synaptic changes on activity regulation in the model networks and compared the emerging activity to measured activity in the visual cortex *in vivo*. Our modeling shows that the operating regime of the circuit - either stabilized by inhibition or stable without it - and the interaction of multiple interneuron subtypes determine how synaptic changes shape activity. In a rate model, we analytically derive the dependence of the network's response on the regime and interneuronal interactions. Our results also generalize to plasticity induced by brief whisker deprivation in the primary somatosensory cortex. Although this deprivation paradigm acts on single neuron excitability as opposed to synaptic plasticity, we found that the effect on network activity is very similar. Together, our results show that to properly interpret the dynamical role of developmental plasticity in neural circuits, we need to take into account the operating regime of those circuits. Even seemingly disparate circuit changes can, for a matched operating regime and global circuit properties, give rise to similar control of circuit dynamics, pointing to more general principles of activity regulation through the coordination of multiple plasticity mechanisms

I-10. STP controls non-linear responses and contrast-invariant tuning in spatially structured networks

Laura Bernaez Timon
Tatjana Tchumatchenko

BERNAEZ.LAURA@GMAIL.COM
TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

Max Planck Institute for Brain Research

Neurons in the primary visual cortex (V1) encode the orientation and contrast of visual stimuli through changes in firing rate. Their activity typically peaks at a preferred orientation and decays to zero at the orientations that are orthogonal. This activity pattern is scaled by contrast but its shape is preserved. Mathematically, this contrast-invariance translates into the factorization of orientation selectivity and contrast sensitivity functions.

Interestingly, contrast-invariant tuning is observed at the population level in V1, which poses additional questions. In particular, it is not clear how the activity of different neurons with diverse orientation selectivity and non-linear contrast sensitivity combine to give rise to contrast-invariant population tuning.

We propose that a major modulation of population response to orientation and contrast occurs at the synapses and is mediated by short-term plasticity (STP). To test this hypothesis, we examine the effect of STP in balanced networks with distance-dependent connectivity both in a rate model and in spiking simulations. In both models, STP controls neurotransmitter release as a function of the presynaptic firing rates, which translates into changes in synaptic strength and in population activity. We find that STP adjusts the network gain and enables the circuit to process the same input sub-, supra-, or linearly depending on the properties of the synapses. We also demonstrate that only synapses that release neurotransmitter in a power-law like support the emergence of contrast-invariance at the network level, while other types of synapses permit the non-linearity but break contrast-invariance.

In summary, we show that synaptic plasticity controls non-linear network responses and mediates the emergence of particular population tuning patterns in networks with distance-dependent connectivity. Our results, therefore, connect the physiology of individual synapses to the network level and may help understand the establishment of population selectivity.

I-11. Cognitive context alters cortical involvement in identical perceptual decisions

Charlotte Arlt¹
Roberto Barroso-Luque²
Ningjing Xia¹
Christopher Harvey¹

CHARLOTTEARLT@GMAIL.COM
ROBALU94@GMAIL.COM
NXIA@G.HARVARD.EDU
CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

¹Harvard University
²University of Chicago

A goal of decision-making studies is to understand the causal relationships between cortical areas and specific behavioral tasks. In some cases, such relationships have been difficult to establish because systematic tests of the factors determining whether a cortical area is necessary for a task have not been performed. Here we show that identical perceptual decisions require different cortical areas depending on the cognitive context, including past experiences. We trained mice to associate visual cues with navigation actions as they locomoted through a virtual reality environment. We required mice to perform identical perceptual decisions, but varied the cognitive context from simple (two fixed associations) to complex (decisions in the presence of association switches, many associations, or delay epochs). Optogenetic silencing of cortical association areas – posterior parietal cortex (PPC), retrosplenial cortex (RSC), or anterior cingulate cortex (ACC) – had only modest effects on decision performance in the simple cognitive context. In contrast, silencing these regions led to large performance decrements when the identical decisions were embedded in complex cognitive contexts. Decision-making relied on PPC and RSC in all complex contexts, whereas ACC was necessary specifically when working memory was required. Sur-

prisingly, a mouse's prior experience with complex cognitive contexts had profound and long-lasting effects on the requirement of cortical regions. Association regions were required for decisions in a simple context even weeks after exposure to a complex cognitive context, but were not required for identical decisions if mice lacked such prior experience. Our results reveal that cognitive context, including from long-ago experience, can profoundly impact the neural implementation of simple perceptual decisions, even to the degree of which brain areas are needed. Therefore, studies of the involvement of brain regions in decision-making must take into account the cognitive context and experience of the subjects, which has often been ignored or not reported.

I-12. Hierarchical neural network models that more closely match primary visual cortex tend to better explain higher level visual cortical responses

Tiago Marques
Martin Schrimpf
James J DiCarlo

TMARQUES@MIT.EDU
MSCH@MIT.EDU
DICARLO@MIT.EDU

Massachusetts Institute of Technology

Object recognition relies on the hierarchical processing of visual information along the primate ventral stream. Artificial neural networks (ANNs) recently achieved unprecedented accuracy in predicting neuronal responses in different cortical areas and primate behavior. Here, we extended this approach by testing hundreds of different models to quantitatively ask how well they explain primate primary visual cortex (V1) across a wide range of experimentally characterized functional properties. We found that, for some ANNs, individual artificial neurons in early and intermediate layers have functional properties that are remarkably similar to their biological counterparts, and that the distributions of these properties over all neurons approximately match the corresponding distributions in primate V1. Still, none of the candidate models was able to account for all the functional properties, suggesting that current network architectures might not be capable of fully explaining primate V1 at the single neuron level. Since some ANNs have "V1 areas" that more precisely approximate primate V1 than others, we investigated whether a more brain-like V1 model also leads to better models of higher-level areas, such the inferior temporal (IT) cortex, which supports object recognition behavior. Indeed, we found that, over a set of 30 ANN models optimized for object recognition, V1 similarity was positively correlated with neural predictivity along the ventral stream, particularly IT. This result supports the widespread view that the complex visual representations required for object recognition are derived from low-level functional properties, but it also demonstrates – for the first time – that working to build better models of low-level vision has tangible payoffs in explaining high-level neural properties and the behavior they support. Moreover, the set of functional V1 benchmarks we constructed can be used as a gradient to search for better models of V1, which will likely result in better models of the primate ventral stream.

I-13. Algorithmically engineered synaptic plasticity rules

Basile Confavreux¹
Friedemann Zenke²
Everton J Agnes¹
Tim Vogels¹

BASILE.CONFAVREUX@GMAIL.COM
FRIEDEMANN.ZENKE@FMI.CH
EVERTON.AGNES@CNCB.OX.AC.UK
TIM.VOGELS@CNCB.OX.AC.UK

¹University of Oxford

²Friedrich Miescher Institute for Biomedical Research

Synaptic plasticity is fundamental for the brain's life-long learning abilities. In computational models, plasticity is often implemented as rules that take into account pre- and postsynaptic firing patterns at a given synapse and change its efficacy accordingly. The most common approach to finding such rules has been manual construction by trial-and-error, identifying simple relationships that capture data from experimental recordings. However, these rules are typically under-constrained by the available data. They only provide a partial picture of the various

mechanisms at work. Indeed, testing any such-derived rule in new settings (e.g., large network models) rarely produces stable or functional results. We thus sought to assist the search for more complete biologically plausible plasticity rules by way of ‘learning to learn’.

We developed an algorithmic meta-learning framework to automatically shape learning rules in function space. Specifically, our algorithm iteratively refines learning rules and evaluates them on how effectively they transform an initially random network into one with a desired structure or function.

We showcase the validity of our approach in single neuron models with plastic afferents, for which we re-discover previously described plasticity rules. First, we re-discover ‘Oja’s rule’, a simple Hebbian plasticity rule that tunes the synapses of a two-layer rate model according to the direction of largest variability of its input. By comparing the architecture constructed by our candidate rules with the results of Oja, we define a loss function to minimize, and successfully obtain the original rule from a randomly initialized learning rule.

We explore similar approaches in single integrate-and-fire neurons with tuned excitatory, and random inhibitory inputs. Here, the loss function of our algorithm is the residual excitation-inhibition imbalance. We expect the algorithm to reproduce analogs of previously published inhibitory synaptic plasticity rules. Visit our poster to see if we succeed.

I-14. Sensorimotor processing for tracking action outcomes in the zebrafish serotonergic system

Takashi Kawashima¹

Ziqiang Wei²

Sujatha Narayan²

Liam Paninski³

Misha Ahrens²

TAKASHI.KAWASHIMA@WEIZMANN.AC.IL

WEIZ@JANELIA.HHMI.ORG

NARAYANS@JANELIA.HHMI.ORG

LIAM@STAT.COLUMBIA.EDU

AHRENSM@JANELIA.HHMI.ORG

¹Weizmann Institute of Science

²HHMI Janelia Research Campus

³Columbia University

To enable animals to learn from the consequences of their actions, how does the brain distinguish stimuli caused by an animal’s actions from stimuli that are passively observed? The dorsal raphe nucleus (DRN) of the serotonergic (5-HT) system in zebrafish, which underlies short-term motor learning, responds specifically to the visual flow generated by the animal’s own movement: neurons encode distance traveled during swimming, but do not respond to visual flow when the animal is still (Kawashima et al., Cell, 2016). To uncover neuronal mechanisms of this computation at the network and cellular level, we monitored multiple variables in 5-HT DRN neurons in zebrafish during virtual reality motor learning tasks: subthreshold membrane potential and spiking using voltage imaging (Abdelfattah et al., Science, 2019), and glutamatergic and GABAergic input into the dendrites using neurotransmitter indicators (Marvin et al., Nature Methods, 2019). Through a combination of these experimental approaches and computational analysis and modeling, we uncovered an unexpected implementation of neural gating. Swim actions triggered GABA release that transiently hyperpolarized 5-HT neurons. This was followed by a voltage rebound from inhibition, during which visually-tuned glutamatergic input triggered spiking that encoded the distance travelled during swimming. The visual responses depended on rebound from GABAergic inhibition, as it was absent following ablation of GABAergic DRN neurons. In particular, computational models indicated that the network requires cellular rebound to reproduce this gating effect. In summary, sensorimotor computation in the DRN is implemented by a temporally-precise, behavior-locked sequence of inhibition and excitation, that together with cellular mechanisms allows serotonergic neurons to encode the consequences of actions. This neuronal mechanism may underlie more general forms of gated communication between brain areas.

I-15. Landmark stability needed for formation but not expression of spatial maps in retrosplenial cortex

Dhruba Banerjee
Zaneta Navratilova
Bruce McNaughton
Sunil Gandhi

University of California, Irvine

DHRUBAB@UCI.EDU
ZNAVRATI@UCI.EDU
BRUCEMCN@UCI.EDU
SPGANDHI@UCI.EDU

Visual cues are useful for orienting in space and guiding navigation. Because not all cues remain fixed at a single location in the environment, learning to distinguish dynamic cues from stable landmarks can be important when creating a cognitive map. Brain imaging studies suggest the human retrosplenial cortex (RSC) selectively responds to stable landmarks. This role is plausible given RSC's connectivity to the hippocampal formation and visual cortex in both primates and rodents, but the cellular circuit basis remains unclear. Recent publications have revealed hippocampal-dependent spatial-tuning in a subset of mouse RSC cells. We hypothesize these spatial cells would show activity patterns modulated by the stability of visual cues. To test this, we trained tetO-GCaMP6s (Camk2a-tTA driver) mice to run in a head-fixed virtual reality setup. Using two photon imaging we monitored the activity of excitatory RSC cells. Mice ran in environments with only stable visual landmarks (n=4), or a combination of stable and dynamic visual cues (n=3). Confirming prior studies, we found a prominent population of spatially tuned cells spanning the entire familiar environment with only stable landmarks. On the other hand, mice running in familiar environments with moving visual cues developed no spatially tuned cells in the region without stable landmarks. When the mice transitioned from a familiar stable environment to an entirely new stable environment, the RSC population remapped completely. Moving previously stable landmarks within an already familiar environment, however, did not disrupt the spatial tuning of RSC cells. From these results we conclude that RSC requires stable visual cues to form a spatial representation, but once formed, the spatial tuning is robust to landmark perturbations. In ongoing work, we seek to understand the circuit mechanisms by which RSC integrates experiences to identify stable landmarks and how this information is used by other cortical circuits to inform navigation.

I-16. Balancing excitation and inhibition in spiking networks with input dependent inhibitory plasticity (idip)

Klara Kaleb
Victor Pedrosa
Claudia Clopath
Imperial College London

KK1918@IC.AC.UK
V.PEDROSA15@IMPERIAL.AC.UK
C.CLOPATH@IMPERIAL.AC.UK

Inhibitory plasticity is thought to be instrumental in establishing and maintaining the excitatory (E)/ inhibitory (I) balance in neural networks. In experiments in the rodent sensory cortices, neuronal activity was shown to be unperturbed for a brief period after sensory deprivation due to rapid disinhibition caused by loss of excitatory input to the PV interneurons. This suggests that inhibitory neurons might be able to sense local network activity and make compensatory changes, and hence exhibit input dependant inhibitory plasticity (IDIP). Here we use a computational framework to study the effects of such a non-Hebbian inhibitory learning rule in a spiking neural network. We show that IDIP is able to stabilise recurrent network dynamics, as well as preserve network firing rate heterogeneity and stimulus representation. Interestingly, in an associative memory task, IDIP facilitated persistent activity for a period of time after memory encoding, in line with some experimental data. Furthermore, IDIP in combination with place tuned input can explain formation of active and silent place cells in the hippocampus, as well as place cells remapping following optogenetic silencing of active place cells. Hence establishment of global balance with IDIP has diverse functional implications and may be able to explain experimental phenomena across different brain areas.

I-17. Multiscale associative memory recall by modulation of inhibitory circuits

Tatsuya Haga¹
Tomoki Fukai²

TATSUYA.HAGA@RIKEN.JP
TOMOKI.FUKAI@OIST.JP

¹RIKEN Center for Brain Science

²Okinawa Institute of Science and Technology

The brain recognizes information in different scales and resolutions. For example, language consists of hierarchically structured information representations (paragraph>sentence>word, or semantic structures like plant >fruit >apple) and we can flexibly change the scale of our recognition. How neural networks construct and flexibly process such multiscale information representations has been poorly understood. Here, we demonstrate that a single neural network can offer multiscale information representations based on the graph structure of associative links between memory items. We propose a generalized Hopfield-type associative memory model which can be implemented by excitatory circuits with Hebbian learning between items, and the combination of local and global inhibitory circuits. This network model “analyses” the cluster structure of associative links, and the network flexibly recalls information representation corresponding to each item (microscopic scale), a cluster of items (intermediate scale), and a cluster of interconnected clusters (macroscopic scale) by modulating the ratio between strengths of local and global inhibition. Through mathematical analyses and numerical simulations, we show that recalled activity patterns in this network model correspond to eigenvectors of graph Laplacian, and that modulations of inhibitory circuits control the activation of eigenvectors to change the scale of representations. Eigenvectors of graph Laplacian represent global and local structures in the graph (Belkin & Niyogi, 2003), and were previously used for optimal graph partitioning (Shi & Marik, 2000) and segmentation of images and speech signals into multiple chunks (Tung et al., 2010; Bach & Jordan, 2013). Our model extends these properties to provide a novel framework for multiscale information processing in a single network (c.f., hierarchical networks with different time constants such as Yamashita & Tani, 2008). Furthermore, our model offers an insight into the computational role of inhomogeneous neuromodulations of interneuron circuits, which is a focus of recent experimental studies (e.g. Pi et al., 2013; Fu et al., 2014).

I-18. How do time and uncertainty motivate information seeking?

Ethan Bromberg-Martin¹
Takaya Ogasawara²
Yang-Yang Feng²
J Kael White²
Kaining Zhang²
Ilya Monosov²

NEUROETHAN@GMAIL.COM
TAKAYA.OGASAWARA@WUSTL.EDU
YANGYANG.WUSM@GMAIL.COM
JKAELOWHITE@GMAIL.COM
KAINING@WUSTL.EDU
ILYA.MONOSOV@GMAIL.COM

¹Washington University in St Louis

²Washington University School of Medicine

Humans and animals can be strongly motivated to seek information about uncertain future rewards, even when they cannot use this information to influence the outcome. There has been considerable debate over how to modify current theories of reinforcement and value-based decisions to explain this “non-instrumental” information seeking. Underlying this debate is one fundamental, outstanding question: what algorithm does the brain use to compute the subjective value of information (SVOI)?

We hypothesized that the brain computes SVOI with a simple mechanism: aversion to the state of being uncertain about future rewards, such that information is valued because it resolves uncertainty. If so, the brain must compute SVOI using two key variables: (1) how early information will arrive in advance of the reward outcome, and (2) how much uncertainty it will resolve. Furthermore, by measuring how SVOI varies across reward distributions, we could infer the mathematical function the brain uses to compute reward uncertainty – another outstanding question in the literature.

To test this, we trained monkeys to perform an economic decision-making task for juice rewards. Each option had

visual features indicating its probability distribution of juice rewards, whether it would show an informative cue to reveal the reward size in advance, and the information and reward delivery times. Monkeys willingly traded juice for information, letting us measure its subjective value.

The best-fitting model of behavior implied that monkeys indeed compute SVOI using both information timing and uncertainty resolution, and estimate uncertainty using a specific functional form not proposed in existing literature. In neuronal recordings, key signatures of these uncertainty and value computations were found in a subcortical motivational pathway from the ventral pallidum (VP) to the lateral habenula (LHb), arising as animals made their decisions. Our results place strong constraints on the computational and neuronal mechanisms that value information in economic decisions.

I-19. Insights from a deep convolutional neural network into mid-level representation in visual cortex

Wyeth Bair
Dean Pospisil
Anthony Bigelow
Dina Popovkina
Taekjun Kim
University of Washington

WYETH0@UW.EDU
DEANP3@UW.EDU
TONY.BIGELOW1@GMAIL.COM
DINA4@UW.EDU
TAEKJUN@UW.EDU

Artificial neural networks (ANNs) instantiate complex visual systems that differ from the brain but nevertheless contain visual representation learned under pressure to categorize natural images in the face of a limited number of unique kernels in early stages. We examined the visual representation in AlexNet, a popular convolutional neural network (CNN) that, after training, contains V1-like front-end filters and mid-level units with V4-like shape tuning and translation invariance. We uncovered an intriguing pattern of results that bears on differential encoding of boundary and surface properties and that links with observations from visual cortex. Recent studies show that fill-outline invariance (FOI) is weak in V4 (i.e., shape selectivity changes when filled shapes are replaced by their outlines) and that many V4 neurons preferentially encode texture. Utilizing stimuli from those studies, we found that most CNN units have high FOI compared to V4 and that FOI relates strongly to texture selectivity, color selectivity and dynamic range within the CNN, which contains partially segregated achromatic shape vs. color/texture pathways that merge downstream. Our findings suggest that CNN units multiplexing color/texture information with shape are more brain-like, or at least V4-like, raising the important question of why ANNs develop such a strong, invariant, achromatic shape representation that is apparently absent in V4. In analogy to connectomics studies, we have the full ANN wiring diagram, but also full access to network activity. Leveraging this, we find that the composition of rectification and weight-covariance in convolutional kernels across layers can explain diversity in invariance and a distributed representation along the achromatic-form vs. color-texture axis in the network. We are testing these predictions in vivo and across more varied ANN architectures. The use of ANNs as predictive model organisms may be increasingly important as they become more brain-like, incorporating recurrence and operating on more complex, realistic tasks.

I-20. Probabilistic successor representations allow for flexible behaviour

Jesse Geerts^{1,2}
Kimberly Stachenfeld³
Neil Burgess¹

JESSEGEERTS@GMAIL.COM
STACHENFELD@GOOGLE.COM
N.BURGESS@UCL.AC.UK

¹University College London
²Sainsbury Wellcome Centre
³DeepMind

The effectiveness of Reinforcement Learning (RL) depends on an animal's ability to assign credit for rewards to the appropriate preceding stimuli or events. One aspect of understanding the neural underpinnings of this process involves understanding what sorts of stimulus representations support generalisation – understanding which aspects of stimuli are rewarding, so that learning about one stimulus tells you about others that share some of its properties. The Successor Representation (SR) enforces generalisation over stimuli that predict similar outcomes, and has become an increasingly popular model because it provides a normative explanation of both behaviour and neural phenomena in rodent hippocampus. Next to generalisation, another dimension of credit assignment involves understanding how animals handle uncertainty about learned associations, using probabilistic methods such as Kalman Temporal Differences (KTD). Combining these approaches, we propose using KTD to estimate a distribution over the SR. KTD-SR captures uncertainty about the estimated SR as well as covariances between different long-term predictions. We applied KTD-SR to a decision making task that probes an agent's ability to respond flexibly to changes in the reward function and in the transition structure, respectively. We show that, because KTD-SR keeps track of the SR covariance information, it exhibits partial transition reevaluation as humans do. Unlike previous work, our model explains this result using a single learning system. We conclude by discussing future applications of the KTD-SR as a model of the interaction between predictive and probabilistic animal reasoning.

I-21. Multivariate phase coupling networks across brain regions using torus graphs

Josue Orellana¹
 Natalie Klein¹
 Scott Brincat²
 Earl Miller²
 Robert Kass¹

JOSUE@CMU.EDU
 NEKLEIN@ANDREW.CMU.EDU
 SBRINCAT@MIT.EDU
 EKMILLER@MIT.EDU
 KASS@STAT.CMU.EDU

¹Carnegie Mellon University

²Massachusetts Institute of Technology

Functional connectivity among multiple brain regions is of great interest, and a leading theory is that neural oscillations may facilitate communication across areas. Traditionally, phase locking value (PLV) is used to quantify the association between oscillatory phases extracted from recordings such as LFPs, but PLV is a bivariate measure and cannot successfully identify multivariate network structure. We have developed a powerful multivariate approach based on what we call “torus graphs”. In a torus graph each node represents the phase of a particular LFP, and an edge is placed between two nodes when they show direct functional connectivity (as opposed to indirect connectivity through other nodes). In 24 dimensional LFP data recorded from PFC and three hippocampal subregions (CA3, DG, Sub), torus graphs provided sensible results, while PLV did not. Specifically, we identified a network of beta-band phase coupling where communication between PFC and Hippocampus was mediated by CA3 and Sub. This network suggests that during the stimulus presentation of an associative memory task hippocampus is coupled with PFC through output subregions. A torus graph is an exponential family model for a multidimensional vector of circular random variables (representing phases), and as such it is a maximum entropy model with properties analogous to a Gaussian graphical model (a workhorse in machine learning). However, in the Gaussian case a single scalar can describe positive and negative association between two real-valued variables, while in a torus graph we need two complex numbers to describe association between two phase-angle variables. We have shown that previous models in the literature are special cases of torus graphs, and developed fast estimation and inference methods that have excellent performance in simulations. Torus graphs can be extended using hierarchical latent variables models to make inferences about functional coupling across brain areas from large numbers of oscillatory signals.

I-22. Inferring random network parameters from continuous-time trajectories

Alexander van Meegen
Tobias Kuhn
Moritz Helias
Research Centre Jülich

A.VAN.MEEGEN@FZ-JUELICH.DE
T.KUEHN@FZ-JUELICH.DE
M.HELIAS@FZ-JUELICH.DE

Intriguing features of cortical dynamics are reproduced by random neural networks—from fluctuation driven activity to chaos transitions. In a seminal study, Sompolinsky, Crisanti, and Sommers showed that such random networks furthermore admit an analytical treatment which predicts chaotic activity in a certain parameter regime (Sompolinsky et al., 1988). Recently, these results received considerable attention and were substantially extended (Kadmon and Sompolinsky, 2015; Mastrogiuseppe and Ostojic, 2017; Schuecker et al., 2018; Muscinelli et al., 2019). Here, we consider the inverse problem: Given the activity produced by a random network, what are the underlying parameters?

This is a challenging problem because the data consists of trajectories that are continuous in time. Our approach is based on a novel bridge between the path integral approach used by Sompolinsky, Crisanti, and Sommers (Crisanti and Sompolinsky, 2018), and the mathematically rigorous approach by Ben Arous and Guionnet, based on large deviation theory (Arous and Guionnet, 1995). The key result is that the log-likelihood of the empirical measure takes the form of a Kullback-Leibler divergence. Minimizing the Kullback-Leibler divergence yields a simple condition that allows a straightforward inference of the network parameters.

The analytical form of the likelihood opens the door for many applications. Here, we use it to infer the network parameters. Other possible applications are variational inference (What is the best random network explaining the data?) or hypothesis testing (Is the data generated by a random network?). Our result already generalizes to arbitrary single-unit dynamics and coupling functions; we are confident that it is possible to generalize the theory to networks with multiple populations and spiking neurons.

I-23. Local synaptic balancing for network robustness

Christopher Stock
Sarah Harvey
Samuel A Ocko
Surya Ganguli
Stanford University

CHSTOCK@STANFORD.EDU
HARVEYS@STANFORD.EDU
SOCKO@STANFORD.EDU
SGANGULI@STANFORD.EDU

A neural network that performs a task is often thought to implement a particular, deterministic input-output map. Generally, there exist many configurations of synaptic weights which compute the same deterministic input-output map and which, therefore, would enjoy equivalent task performance in the absence of noise. Yet, the assumption of deterministic computation hardly holds in biology, where single neurons compute in a highly noisy information environment. Two networks which in the absence of noise would compute the same input-output map may perform quite differently when noise is present. It is of considerable interest, then, to not only 1) better understand the set of weight configurations which yield equivalent task performance in the deterministic setting, but also 2) locate configurations which yield improved task performance in the presence of noise. Unfortunately, given the complexity and nonlinearity of tasks and of networks, analytical answers to these questions have been hard to find.

Here we present an analytical method, which we call local synaptic balancing, to explore entire sets of deterministically equivalent networks (addressing question 1) while finding weight configurations which are more robust to noise (addressing question 2). Importantly, our method does not rely on task-based network training—or for that matter a learning signal of any kind—but, rather, exploits natural symmetries in commonly-studied neural network formulations. Key to our approach is constructing an integrable dynamical system on the space of synaptic weight matrices, leading to many conserved quantities, which in turn imply conservation of global dynamics while

increasing network robustness in a mean-field sense.

Remarkably, our exploration of weight space corresponds directly to weight updates that are locally implementable by biological neurons, raising the possibility that biological neural networks implement local synaptic balancing as a means to improve the robustness of computation.

I-24. Spectral regularization in biological and artificial neural networks

Josue Nassar¹
Piotr A Sokol¹
SueYeon Chung²
Kenneth Harris³
Il Memming Park¹

JOSUE.NASSAR@STONYBROOK.EDU
PIOTR.SOKOL@STONYBROOK.EDU
SUEYEON.CHUNG@COLUMBIA.EDU
KENNETH.HARRIS@UCL.AC.UK
MEMMING.PARK@STONYBROOK.EDU

¹Stony Brook University

²Columbia University

³University College London

Recently, Stringer et al.2019 used large-scale recordings to show that the mouse visual cortex exhibits high-dimensional responses to visual stimuli with a power-law-like eigenspectrum $\lambda_n \propto n^{-\alpha}$. Surprisingly, the exponent α was close to the theoretical boundary 1, but not smaller where the representational smoothness shatters. The authors conjectured that this representation is ideal for large neural networks, as it allows for an optimal trade off between robustness and generalization. They also conjectured that the presence of adversarial examples in artificial neural networks (ANN) is due to the insufficiently fast spectral decay (Szegedy 2013). As ANNs are becoming a cornerstone in theorizing about neural computations, the presence of adversarial attacks is worrying from a computational neuroscience perspective because humans are robust against them.

Building on Stringer et al.'s conjecture we investigate whether power-law-like spectral regularization of representations in ANNs makes them robust to adversarial attacks. We hope that in doing so, we will better understand the role of high-dimensional codes in the mammalian visual cortex. first, we show that adversarially robust ANNs have better power-law exponent as suggested by theory. Second, we developed a novel explicit spectral regularization for wide ANNs inspired by neuroscience and show that it improves robustness against adversarial attacks. finally, we prove that in the limit of infinite neurons, an ANN whose neural code decays slower than n^{-1} is not Lipschitz continuous, allowing small perturbations in the input space to cause arbitrarily large deviations in the neural code. This extends the existing theory and provides an insight into ANNs, especially with multiple wide-layers. If one intermediate layer violates Lipschitz continuity, then so will the following, even if the power-law exponent at the final layer is in the redeemable range. In conclusion, our theory and analysis brings new insights into both biological and artificial NNs that support Stringer and coworker's conjecture.

I-25. IronClust: Scalable and drift-resistant spike sorting for long-duration, high-channel count recordings

James Jun¹
Jeremy Magland¹
Catalin Mitelut²
Alex Barnett¹

JAMESJUN@GMAIL.COM
JMAGLAND@FLATIRONINSTITUTE.ORG
MITELUTCO@GMAIL.COM
ABARNETT@FLATIRONINSTITUTE.ORG

¹flatiron Institute

²Columbia University

Spike sorting enables a direct measurement of neural population activity at a single-cell, single-spike resolution based on their extracellular voltage waveforms. Silicon probes are widely used to record action-potential events on

multiple adjacent electrodes, but probe movements in the brain often result in spike sorting errors requiring manual curation efforts. Latest-generation neural probes (e.g. Neuropixels) offer high channel count, density, and spatial coverage that could potentially improve the spike sorting accuracy by tracking the probe movement in the brain tissue. We developed a reliable, scalable, and usable spike sorting package (IronClust) that is resistant to probe drift by applying a KNN-graph density clustering, anatomical similarity matching, and parallel computing. Our automated workflow is comprised of detection, clustering, and merging steps where 1) detection step performs filtering, event detection, and feature extraction, 2) clustering step includes anatomical similarity matching, KNN-graph discovery, and density clustering, 3) automated merging based on KNN-graph and waveform similarity. Sorting accuracy was independently verified by the SpikeForest framework based on an extensive collection of real and simulated recordings with ground truth. IronClust achieves scalable operation by performing parallel computation in three dimensions (time, channels, unit index) successively at each sorting step based on a wide range of hardware platforms (multi-core CPU, GPU, and compute clusters). We developed an interactive user interface for efficient manual inspection and correction. IronClust supports interoperability with other frameworks such as SpikeInterface and SpikeForest by providing python wrappers, and exporters for manual sorting using other sorters (Phy, Klusters, and JRCLUST). In conclusion, IronClust offers the best combination of speed and accuracy for various types of electrodes among ten sorters compared. A GPU-workstation can achieve real-time speed up to ~2000 channels, or even faster with access to an HPC cluster. We recommend IronClust for analyzing long-duration, high-channel-count probe recordings containing drift.

I-26. A common model explaining flexible decision making, grid fields, and cognitive control

Payam Piray
Nathaniel D Daw
Princeton University

P.PIRAY@GMAIL.COM
NDAW@PRINCETON.EDU

A central difficulty for reinforcement learning (RL) in sequential tasks is that the value of an action now depends on which actions you take afterwards. Thus, optimal actions are coupled across states. We argue that this fact underlies a pattern of challenges for RL models in neuroscience to explain both the brain's flexibility (to replan, transfer, and define useful neural representations of long-run contingencies) and its inflexibility (such as cognitive control and Pavlovian biases). Building on recent advances in control engineering, we propose a new model for decision making in the brain, termed linear RL, which connects a wide range of seemingly disparate empirical and theoretical phenomena. The model replaces the classic iterative Bellman optimization with an approximate variant that softly maximizes around a default, stochastic policy distribution. This allows it accurately and inexpensively to approximate the optimal policy using linear operations and a cached map of long-run state expectancies, which unlike earlier proposals is independent of the optimized policy and current goals. This independence allows the model to explain animals' ability to solve replanning problems that otherwise would require exhaustive, biologically unrealistic model-based tree search. It also explains how stable representations of low-frequency relationships over space (as in the entorhinal grid code) can be useful for choice even under changing goals. Finally, because the optimization works by assigning a cost to deviation from the assumed default policy, it offers an answer to the question why organisms treat certain ("control demanding") actions as costly, and a pattern of explanations of cognitive control and Pavlovian response biases.

I-27. Recurrent neural network models of multi-area computation underlying decision-making

Michael Kleinman¹
 Chandramouli Chandrasekaran²
 Jonathan Kao¹

MICHAEL.KLEINMAN@UCLA.EDU
 CCHANDR1@BU.EDU
 KAO@SEAS.UCLA.EDU

¹University of California, Los Angeles

²Boston University

Cognition emerges from the coordination of computations in multiple brain areas. However, elucidating these coordinated computations within and across brain regions is challenging because intra- and inter-areal connectivity are typically unknown. To study coordinated computation, we trained multi-area recurrent neural networks (RNNs) to discriminate the dominant color of a checkerboard and output decision variables reflecting a direction decision, a task previously used to investigate decision-related dynamics in dorsal premotor cortex (PMd) of monkeys. We found that multi-area RNNs, as opposed to single RNNs, reproduced the decision-related dynamics observed in PMd during this task. The RNN solved this task by a novel mechanism where RNN dynamics integrated color information on an axis subsequently readout by an orthogonal direction axis. Direction information was selectively propagated through preferential alignment with inter-areal connections, while the color information was filtered. These results suggest that cortex uses modular computation to generate minimal but sufficient representations of task information, selectively filtering unnecessary information between areas. Finally, we leverage multi-area RNNs to produce empirically testable hypotheses for computations that occur within and across brain areas, enabling new insights into distributed computation in neural systems.

I-28. Neuronal representations in mice barrel cortex and ANNs performing a whisker discrimination task

Ramon Nogueira
 Chris Rodgers
 Randy Bruno
 Stefano Fusi
 Columbia University

RN2446@COLUMBIA.EDU
 CCR2137@COLUMBIA.EDU
 RB2604@COLUMBIA.EDU
 SF2237@COLUMBIA.EDU

Rodents, and in particular mice, scan objects by whisking, which allows them to form an internal representation of their environment. Although simple whisker-based detection tasks can typically be solved by a linear combination of sensory evidence across whiskers, challenging object recognition tasks might require a more complex transformation of these cues across input channels and time. In order to elucidate the behavioral and neural mechanisms underlying complex object identification, mice were trained on a novel task where they had to discriminate concave from convex shapes. Mice learned the task well, but because the task is challenging, their performance plateaued at 75.7% correct on average, still significantly above chance. From the pattern of cumulative whisker contacts we were able to decode both with linear and non-linear classifiers shape identity (~70%) and mice choice (60%). Additionally, we trained a recurrent artificial neural network (ANN) to perform an analogous “whisker-based” discrimination task. In particular, we trained two networks to perform a simple and a complex discrimination task, which required linear and nonlinear cue combination across input channels, respectively. In order to characterize how neurons in biological and artificial networks represent the input variables that are necessary to perform the task, we characterized the encoding properties of each network by fitting encoding models with different levels of complexity to each dataset. We found that non-linear encoding models were better models of neuronal activity for the barrel cortex and for the ANN trained on the complex task. On the other hand, the activity of units in the ANN trained for the simple task was equally explained by linear and non-linear encoding models. Overall, our results suggest that neural networks can adopt different representational strategies depending on the task they have been optimized for, producing nonlinear (and therefore high-dimensional) representations for more complex tasks.

I-29. Simultaneous dimensionality reduction and deconvolution for calcium imaging recordings

Tze Hui Koh¹
 William Bishop²
 Takashi Kawashima²
 Misha Ahrens²
 Steven Chase¹
 Byron Yu¹

TZEHUIK@ANDREW.CMU.EDU
 BISHOPW@JANELIA.HHMI.ORG
 KAWASHIMAT@JANELIA.HHMI.ORG
 AHRENSM@JANELIA.HHMI.ORG
 SCHASE@ANDREW.CMU.EDU
 BYRONYU@CMU.EDU

¹Carnegie Mellon University

²HHMI Janelia Research Campus

Dimensionality reduction is commonly applied to study the spiking activity of populations of neurons underlying motor control, decision-making, and more. In recent years, calcium imaging has emerged as a powerful tool for recording the simultaneous activity of populations of identified cell types. Because the fluorescence from calcium imaging represents a transformed, smoothed version of the underlying electrical activity of the neurons, an important question is whether the same dimensionality reduction techniques applied to electrophysiological recordings would be appropriate for calcium imaging recordings. We developed a novel method, Calcium Imaging Linear Dynamical System (CILDS), that unifies dimensionality reduction and deconvolution in a single probabilistic framework. We compared CILDS with 1) dimensionality reduction performed directly on recorded fluorescence, and 2) a two-stage method in which deconvolution is applied to each fluorescence trace to estimate spiking activity, then dimensionality reduction performed on this activity. In simulation, we identified regimes in which each method most accurately recovered the ground truth latent variables by systematically varying parameters such as noise level, timescales of latent variables, and calcium dynamics. CILDS is better able to separate latent variables from calcium dynamics than the other two methods, particularly as the latent timescales decrease and the number of neurons increase. We then applied these methods to calcium imaging recordings from the dorsal raphe nucleus in larval zebrafish performing a motor learning task. We found that CILDS extracted latent variables that more accurately predicted the activity of held-out neurons than the other two methods. More broadly, this work can guide the use of dimensionality reduction in calcium imaging experiments.

I-30. Proprioceptive tuning emerges in a task-driven hierarchical deep neural network model

Kai Sandbrink¹
 Pranav Mamidanna²
 Claudio Michaelis²
 Mackenzie Mathis¹
 Matthias Bethge²
 Alexander Mathis¹

KAIJSANDBRINK@GMAIL.COM
 PRANAV.MAMIDANNA@GMAIL.COM
 CLAUDIO.MICHAELIS@UNI-TUEBINGEN.DE
 MATHIS@ROWLAND.HARVARD.EDU
 MATTHIAS.BETHGE@UNI-TUEBINGEN.DE
 AMATHIS@FAS.HARVARD.EDU

¹Harvard University

²Eberhard Karls Universitat Tubingen

Proprioception is critical for sensing and controlling the body, yet no canonical hierarchical model of the system exists. Task-driven modeling has provided important insights into other sensory systems. However, unlike for vision and audition, databases of relevant proprioceptive stimuli are not readily available. Here we propose a proprioceptive MNIST-like task. We generated a large scale dataset of human arm trajectories as the hand is tracing the alphabet in 3D space. Using a human musculoskeletal arm model together with muscle spindle models we simulated the input to the spinal cord for this task. Due to the multi-segmental nature of the arm, and the distributed nature of muscles, the spinal cord only has implicit information about which character is traced. SVMs indeed perform poorly at solving this task (23% multiclass accuracy). Artificial neural networks with sufficiently many layers (>3) could robustly solve this task (>99% multiclass accuracy). We analyzed the networks' units

and discovered directional movement tuned neurons akin to the primate somatosensory cortex (S1). Intriguingly, comparing the trained networks with the same network before training (i.e. with randomly initialized weights but the same architecture) yielded surprising results: Neurons in networks before training also show high levels of tuning for kinematic features. Strikingly, differences remain: The trained networks' kinematic representations demonstrate higher tuning invariance across workspaces, and the distribution of preferred directions evens out along the processing pathway in the trained networks. By contrast, the populations in the control networks retain the input muscle spindles' biased distribution. Crucially, S1 also has uniformly distributed preferred tuning just like middle layers of the trained model, but not in the untrained model (Rao's spacing test). We can thus recapitulate aspects of the uniform, direction-selective representation of the work space in primates with a simple character recognition task and task-driven modeling.

I-31. Inter-joint coupling dynamics of walking *Drosophila* from complete 3D leg kinematics

Pierre Karashchuk
 Evyn Dickinson
 Katie Rupp
 Bingni Brunton
 John Tuthill
 University of Washington

PIERREK@UW.EDU
 EDICKINS@UW.EDU
 RUPPP@UW.EDU
 BBRUNTON@UW.EDU
 TUTHILL@UW.EDU

Locomotion requires precisely coordinated sequences of movement. A walking fruit fly, *Drosophila melanogaster*, must coordinate thirty joints across six legs at a step frequency up to 20 Hz. fly walking is thought to be driven by neural circuits that produce rhythmic patterns of activity, termed central pattern generators (CPGs). Previously, CPGs have been used to explain the coordination patterns across insect legs. However, walking requires co-ordination not only across legs, but also between joints within each leg. It has not previously been possible to investigate inter-joint coupling in *Drosophila* due to the challenges of tracking fly leg kinematics with high speed and precision. In this project, we developed new automated 3D tracking methods and comprehensively analyzed full leg joint kinematics to uncover the joint-joint coupling architecture of fly walking.

Previous work has modeled insect leg joint dynamics using coupled oscillators. Such models suggest the two joint angles are produced either by a single underlying oscillator or through two separate but coupled oscillators, predicting strong coupling or joint angles syncing with each other at a stable fixed point, respectively. These predictions do not hold consistently in our data. The phases of the joint angles are not strongly coupled: they deviate by as much as a quarter cycle. Furthermore, although the middle and back legs appear to have stable fixed points, the front legs do not form a stable fixed point. Therefore, some mechanism beyond a centrally driven pattern generator also contributes to within-limb joint coupling in the fly. We propose that joint angles within each limb are coordinated independently through proprioceptive feedback as the fly walks.

I-32. Unsupervised learning of odor sequences in primary olfactory cortex

Andrew Fink
 Carl Schoonover
 Richard Axel
 Columbia University

ANDREW.JP.FINK@GMAIL.COM
 CES2001@COLUMBIA.EDU
 RA27@COLUMBIA.EDU

Classical results in experimental psychology (Blodgett/Tolman, Sokolov, Brogden, Groves/Thompson, Berlyne) have documented that animals continuously learn complex features in their sensory environment. This faculty does not depend on external reinforcement to drive learning, only the existence of structure in the world. How does an organism form a model of the world's statistics when the vast majority of the time it is neither rewarded

nor punished for doing so? We have developed a behavioral paradigm for the head-fixed mouse that permits observation of exploration and non-reinforced learning simultaneous with longitudinal recording of $\sim 10^2$ single units. Mice presented a sequence of two neutral odorants A->B initially investigate the stimuli, and then habituate following repeated experience over three days. Recordings in primary olfactory (piriform) cortex show that over this same time period the magnitude of the population response to the A->B sequence decreases by nearly tenfold, with responses during the B epoch nearly abolished. However, B presented alone elicits strong activity; and if B is omitted from the A->B sequence, the piriform responds vigorously to its absence, producing a representation that matches the expected timing of the omitted stimulus even though there is no odorant molecule in the naris. Taken together these concordant behavioral and electrophysiological results show that response magnitude in both behavior and in piriform activity scales inversely with the degree to which an event is predictable. This suggests that piriform representations do not strictly reflect physical stimuli, but rather a model of the statistics of the olfactory environment, formed over the course of unsupervised learning. Familiar stimuli that contain little new information are filtered out, whereas the content of unanticipated events is transmitted to regions downstream. We are currently elucidating the neuronal mechanism by which this model is formed with a view towards developing new approaches to unsupervised learning in silico.

I-33. The hourglass organization of the *C. elegans* connectome

Constantine Dovrolis¹
 Kaeser Sabrin¹
 Yongbin Wei²
 Martijn van den Heuvel²

CONSTANTINE.DOVR@GMAIL.COM
 KMSABRIN@GMAIL.COM
 YONGBIN.Y.WEI@GMAIL.COM
 MARTIJN.VANDEN.HEUVEL@VU.NL

¹Georgia Institute of Technology
²Vrije Universiteit Amsterdam

We approach the *C. elegans* connectome as an information processing network that receives input from about 90 sensory neurons, processes that information through a highly recurrent network of about 80 interneurons, and it produces a coordinated output from about 120 motor neurons that control the nematode's muscles. We focus on the feedforward flow of information from sensory neurons to motor neurons, and apply a recently developed network analysis framework referred to as the "hourglass effect". The analysis reveals that this feedforward flow traverses a small core ("hourglass waist") that consists of 10-15 interneurons. These are mostly the same interneurons that were previously shown (using a different analytical approach) to constitute the "rich-club" of the *C. elegans* connectome. This result is robust to the methodology that separates the feedforward from the feedback flow of information. The set of core interneurons remains mostly the same when we consider only chemical synapses or the combination of chemical synapses and gap junctions. The hourglass organization of the connectome suggests that *C. elegans* has some similarities with encoder-decoder artificial neural networks in which the input is first compressed and integrated in a low-dimensional latent space that encodes the given data in a more efficient manner, followed by a decoding network through which intermediate-level sub-functions are combined in different ways to compute the correlated outputs of the network. The core neurons at the hourglass waist represent the information bottleneck of the system, balancing the representation accuracy and compactness (complexity) of the given sensory information.

I-34. Experimental demonstration of single neuron specificity during under-actuated neurocontrol

Samuel Brown¹
 Jason Ritt²
¹Boston University
²Brown University

SAMUELGB@BU.EDU
 JRITT@BU.EDU

Population-level neurocontrol has been advanced predominately through the miniaturization of hardware, such as MEMS based electrodes. However, miniaturization alone may not be viable as a method for single-neuron-specific control within large ensembles, as it is typically infeasible to create electrode densities approaching 1:1 ratios with the neurons whose control is desired. That is, even advanced neural interfaces will likely remain underactuated, where there are fewer inputs (electrodes) within a given tissue than there are outputs (neurons). A complementary “software” approach may help to bridge the gap by allowing individual electrodes to independently control multiple neurons simultaneously. Here we extend an underactuated control schema previously validated in theoretical analysis and simulation (Ching & Ritt, 2013), that uses stimulus strength-duration tradeoffs to activate a target neuron while leaving non-targets inactive. We test this schema in mice in vivo, by independently controlling pairs of cortical neurons receiving common optogenetic input. Initial results show neurons can be specifically and independently controlled following a short (~3 min) system identification procedure. However, when using static stimulation parameters, nonstationarities degrade performance on a timescale of 2-5 minutes. We developed an adaptive control procedure that repeatedly fits recorded neurons to stochastic Integrate and fire (IAF) models based on the deviation of expected and observed spiking in blocks of trials, and selects optimal stimulation parameters in the next block from the updated models. This adaptive recalibration repeats periodically, allowing the controller to adapt to both short- and long-term changes in neural parameters. We find the adaptive approach can maintain control over much longer time periods (>20 minutes).

I-35. Predicting functional similarity of neural networks from architectural features

Adam Haber
Elad Schneidman
Weizmann Institute of Science

ADAMHABER@GMAIL.COM
ELAD.SCHNEIDMAN@WEIZMANN.AC.IL

The combination of novel molecular and imaging tools is transforming our understanding of brain connectivity, and would enable the quantitative study of the relations between structural connectivity and network function in the brain. Current approaches to analyzing ‘connectomes’ rely mainly on graph-theoretical tools and structural features of networks. However, focusing on network topologies downplays the complex nonlinear dynamics of single neurons and networks, and how the function of networks depend on their stimuli. We present a novel approach for comparing networks in terms of their function, by quantifying the similarity of their population responses to the same stimuli. Specifically, we simulated the responses of many thousands of networks of 4, 15, and 50 neurons to different classes of stimuli. We then mapped the space of network architectures and their function, and studied the relation between functional similarity and structural features. We show that in all cases, a wide set of common graph theory measures of similarity of networks convey very little information about the similarity of networks’ responses. Instead, we learn a functional metric between networks based on all synaptic values between cells, and show that it accurately predicts the similarity of novel networks. Surprisingly, we show that this learned metric is dominated by a sparse set of structural features – the sum of synaptic input each neuron receives and the sum of each cell’s synaptic output. We show that based on these features we can predict the functional similarity of networks of up to 50 cells to a high degree of accuracy. Moreover, we show that the functional organization of networks generalizes to other stimuli. We complement our simulation-based approach with an analytical investigation, and demonstrate that the same structural features emerge in a simple probabilistic model of a neural network, regardless of networks size.

I-36. To what extent can the olfactory cortex be modeled by random connectivity?

William Dorrell
Julien Grimaud
Venkatesh Murthy
Cengiz Pehlevan
Harvard University

DORRELLWEC@GMAIL.COM
JGRIMAUD@G.HARVARD.EDU
VNMURTY@FAS.HARVARD.EDU
CPEHLEVAN@SEAS.HARVARD.EDU

Topographic mappings from sensory organs to cortex are frequently employed by living organisms, for example the retinotopic map from retina to V1 preserves spatial topography. Yet similar maps are conspicuously absent from the olfactory system: a wide range of searches have found no recognizable spatial topography in the projection from olfactory bulb to piriform cortex. As a result, it is often posited that this mapping is entirely random whereby each piriform neuron samples from a random subset of bulb neurons. If, in a simple feedforward model, one follows this claim to its logical completion then, by appealing to the Central Limit Theorem, each neuron's activity should be well fit by a thresholded Gaussian. We tested this claim on neuronal recordings from a mouse piriform cortex and found a surprising degree of correspondence.

We next wondered whether random connectivity could explain bilateral odor integration. In many species a septum separates both nares meaning inhaled odors are physically compartmentalized within each nostril. Hence, cross-over between the left and right side of the olfactory system is exclusively neural. Additionally, the bulb does not send contralateral projections so inter-hemispheric interactions are only cortical. We analytically showed that, if the inter-cortical mapping is random, you would expect the two representations of the same odor from different nostrils to be uncorrelated. However, from neuronal recordings we found the representations were highly aligned, to the extent that a classifier trained to identify ipsilateral odor identities could decode contralateral identity with the same accuracy. We built a network model of the two olfactory cortices and showed that the data was best fit by a significantly Hebbian connectivity matrix. This work reveals that, while randomness could explain the bulb to cortical mapping, it fails to make the interhemispheric leap.

I-37. Rule adherence warps choice representations and increases decision-making efficiency

Becket Ebitz¹
Cindy Tu²
Benjamin Hayden¹

REBITZ@GMAIL.COM
CINDYHFLS@GMAIL.COM
BENHAYDEN@GMAIL.COM

¹University of Minnesota

²Washington University in St. Louis

We respond flexibly to the world around us, changing the policies we use to interact with the world as our experiences, expectations, and goals evolve. However, the computations by which these policies guide behavior remain mysterious. One influential idea is that some simple policies, like stimulus-response rules, could be implemented, in part, through an attentional gate that enhances rule-relevant information at the expense of irrelevant information. However, tests of this hypothesis failed to find any evidence for attentional gating in either sensory cortex or in classic prefrontal attentional control regions (e.g. Mante et al., 2013). Here, we consider the possibility that gating emerges at the level of decision-making—meaning in sensory representations within the limbic circuits associated with value-based decision-making. To test this idea, we recorded from single neurons in the orbitofrontal cortex (OFC), ventral striatum (VS) and dorsal striatum (DS) while macaques performed a rule-based decision-making task where the correct rule changed frequently. We found that implementing rules warped population representations of chosen options: simultaneously expanding rule-relevant neural dimensions and compressing rule-irrelevant ones—exactly as we would expect from an attentional gate. These representational changes could, in theory, make decision-making more efficient because they could facilitate rule-relevant computations at the ex-

pense of irrelevant ones. To test this idea, we used a hidden Markov Model to identify physically identical decisions that maximized both reward and information, but could not be explained by a simple stimulus-response rule. Contrasting rule-based decisions with these rule-free decisions revealed that stimulus-response rules increased the efficiency of decision-making in OFC, VS, and DS: reducing spike counts at the same time as it increased choice information. Together, these results suggest that we implement rules, in part, through an attentional gate that increases the efficiency of decision-making by filtering information and warping the representational structure in decision-making circuits.

I-38. A comprehensive analysis of local and global inputs to motion detection pathways

Luis Giordano Ramos Traslosheros Lopez¹
Marion Silies²

LRAMOS@GWDG.DE
MSILIES@UNI-MAINZ.DE

¹International Max Planck Research School Neurosciences

²Johannes Gutenberg University Mainz

Individuals of several species exploit motion cues in the environment to navigate and detect preys or predators. The algorithmic implementation of the canonical computation of motion detection is within reach in the *Drosophila* visual system. Here, local motion detection arises in the dendrites of T4 and T5 direction-selective neurons that encode the movement of bright and dark edges, respectively. T4 or T5 neurons serve as local motion detectors in each of the 800 columnar units that compose the fly eye and contain the cell types that feed into the direction-selective neurons. To become direction-selective, T5 and T4 neurons are thought to combine signals from three neighboring columns. Strikingly, the main synaptic input to T5 neurons, Tm9, displayed both narrow and wide-field receptive field properties in different studies. Here we show that to understand local motion detection in the fly, we need to incorporate cell types with wide dendritic arbors spanning tens of eye columns. Using *in vivo* calcium imaging and an array of visual stimuli, we demonstrate that Tm9 displays both narrow and wide receptive fields depending on the stimulus, thus solving a previous discrepancy in the field. Using GFP reconstitution across synaptic partners, we mapped inputs onto Tm9 from (wide-field) neurons not mapped in connectomics. We then studied the functional connectivity between Tm9 and novel wide-field and connectomic inputs by optogenetic stimulation of individual cell types, showing that wide-field neurons were inhibitory to Tm9. Furthermore, genetically silencing the wide-field inputs enhanced Tm9 responses to bright and dark light flashes and widened the receptive field to bright stimuli. Downstream of Tm9, activity from Tm9 and its wide-field inputs are required for normal direction-selective responses in both T4 and T5. Thus, wide-field neurons contribute to direction selectivity by sharpening receptive fields of the main inputs to T5 neurons.

I-39. L2/3 and L5 pyramidal neuron somata and apical dendrites exhibit distinct responses to unexpected violations of visual flow

Colleen Gillon¹
 Joel Zylberberg²
 Blake Richards³
 Jerome Lecoq⁴
 Yoshua Bengio⁵
 Timothy P Lillicrap⁶
 Jed Perkins⁴
 Jason Pina²
 Sam Seid⁴
 Carol Thompson⁴
 Ryan Valenza⁴

COLLEEN.GILLON@MAIL.UTORONTO.CA
 JOELZY@YORKU.CA
 BLAKE.RICHARDS@MCGILL.CA
 JEROMEL@ALLENINSTITUTE.ORG
 YOSHUA.BENGIO@MILA.QUEBEC
 TIMOTHYLILICRAP@GOOGLE.COM
 JEDP@ALLENINSTITUTE.ORG
 JAYPINA@YORKU.CA
 SAMS@ALLENINSTITUTE.ORG
 CAROLTH@ALLENINSTITUTE.ORG
 RYANV@ALLENINSTITUTE.ORG

¹University of Toronto

²York University

³McGill University

⁴Allen Institute for Brain Science

⁵Montreal Institute for Learning Algorithms

⁶Google DeepMind

Computational neuroscientists have postulated that the neocortex implements hierarchical predictive coding, whereby each region sends predictions of the inputs it expects to lower-order regions, allowing the latter to efficiently propagate prediction errors only. This hypothesis is supported by the presence of prediction-error-like responses to unexpected stimuli in the somatic compartments of layer 2/3 (L2/3) pyramidal neurons in early cortical areas, such as primary visual cortex (VisP). However, to date, no one has systematically investigated whether other potential sites for prediction-error signals exhibit similar responses, such as L5 pyramidal neuron somata, or pyramidal neuron apical dendrites in L1. Here, we addressed this using 2-photon calcium imaging in the VisP of awake mice, in a passive viewing task. We investigated the responses of L2/3 and L5 pyramidal neuron somata and apical dendrites to a visual flow stimulus punctuated by unexpected disruptions, wherein a subset of the stimulus temporary reversed direction. We found that both L2/3 somata and apical dendrites exhibited prediction-error-like signals, independent of visual flow direction, whereas L5 somata and apical dendrites did not. Our results show that both the somata and the apical dendrites of L2/3 neurons possess information about visual flow violations. Future work will dissect the responses to different types of surprising stimuli in these different compartments, as well as the evolution of these responses over longer periods of time.

I-40. Efficient high-dimensional receptive field inference using a flexible spline basis

Ziwei Huang¹
 Philipp Berens²

ZIWEI.HUANG@UNI-TUEBINGEN.DE
 PHILIPP.BERENS@UNI-TUEBINGEN.DE

¹University of Tübingen

²Institute for Ophthalmic Research, University of Tübingen

Spatio-temporal receptive field (STRF) models are frequently used to approximate the computation implemented by a sensory neuron. Typically, such STRFs are assumed to be smooth and sparse. Current empirical Bayes estimation approaches such as automatic relevance determination (ARD), automatic smoothness determination (ASD) and others encode such prior knowledge into a prior covariance matrix, whose hyperparameters are learned from the data, and thus provide STRF estimates with the desired properties even with little or noisy data. However, empirical Bayes methods are not computationally efficient in high-dimensional settings, as often encountered in sensory neuroscience.

Here we pursue an alternative approach and encode prior knowledge for estimation of STRFs by choosing a set of basis function with the desired properties. We use a natural cubic regression spline basis, which is known as the smoothest possible interpolant. We find that this method provides a good basis for high-dimensional STRFs. On spike recordings from retinal ganglion cells, we show that in the linear STRF model with Gaussian noise, spline-based STRFs are as smooth as STRFs estimated by ASD, provide similar prediction performance but are computationally much more efficient. Adding an L1 penalty allows to achieve smoothness and sparseness at the same time, resulting in very good STRF visualizations. In addition, this method allows efficient recovery of 3D STRFs from noisy two-photon calcium recordings from ganglion cell dendrites. Finally, our spline-based method can be naturally extended to hierarchical subunit models (such as the LNLN-Poisson model). The resulting sub-units are already well-separated and smooth using only 5 minutes of data, while traditional estimation methods require much more data to achieve similar results.

I-41. Inhibitory plasticity rules create a contextual switch via complementary receptive fields

Everton J Agnes
Andrea I Luppi
Tim Vogels

EVERTON.AGNES@CNCB.OX.AC.UK
LUPPI272@GMAIL.COM
TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Synapses from different interneuron types likely follow distinct plasticity rules, with unclear consequences for neuronal function. In sensory areas of the cortex, for example, neurons have been shown to change their activation (i.e., receptive field profile) according to context, or behavioural task. Here, we hypothesise that activity changes are due to disinhibitory effects of different inhibitory populations that have distinct connectivity profiles learned via distinct plasticity rules. To test this hypothesis we simulated a postsynaptic neuron receiving correlated excitatory and inhibitory inputs. Excitatory synapses were kept fixed, following a standard receptive field profile, with preferred and non-preferred inputs signals. Inhibitory inputs were divided into two groups corresponding to two distinct interneuron types, each following a specific biologically inspired plasticity rule. We tested three plasticity rules: (i), Hebbian, which potentiates weights for correlated pre- and postsynaptic spikes and depresses synapses for sole presynaptic spikes; (ii), homeostatic scaling, which multiplicatively depresses and additively potentiates inhibitory synapses according to the postsynaptic activity; and (iii), anti-Hebbian with opposite changes as the Hebbian one. We combined the Hebbian with either homeostatic or anti-Hebbian rules. Synapses following the Hebbian rule developed strong weights for preferred inputs and zero weights for non-preferred inputs. Synapses following the homeostatic rule collapsed to a single value, while anti-Hebbian synapses developed vanishing synapses for preferred and strong ones for non-preferred inputs. When both inhibitory populations were active, inhibition balance excitation regardless of the inhibitory tuning profiles, resulting in uncorrelated postsynaptic responses. Modulating the activity of any given inhibitory population produced strong correlations to either preferred or non-preferred inputs, in line with recent experimental findings that show dramatic context-dependent changes of neurons' receptive fields. Our work sheds light on how presynaptic populations can coordinate on long and short timescales to create diverse and flexible receptive fields.

I-42. Hippocampal CA3 neurons are organized into correlated cell assemblies supporting stable place codes

Liron Sheintuch
Nitzan Geva
Alon Rubin
Yaniv Ziv

LIRONSHEINTUCH@GMAIL.COM
NITZGEVA@GMAIL.COM
ALON.RUBIN@GMAIL.COM
YANZIV@GMAIL.COM

Weizmann Institute of Science

Hippocampal CA1 mostly consists of feed-forward connections, whereas hippocampal CA3 exhibits abundant recurrent connections which were hypothesized to provide auto-associative capabilities, facilitating long-term storage and retrieval of information. While previous electrophysiological studies corroborate such capabilities, it remains unclear how large populations of CA3 neurons cooperate to form memories that are stable over long time periods. Recent advances in optical imaging techniques allow chronic readout of activity from hundreds of simultaneously recorded neurons in freely behaving mice, enabling direct measurement of the collective computational capabilities of hippocampal CA3, and of their development over time. Here, we used Ca²⁺ imaging to longitudinally record from large populations of CA3 and CA1 neurons in mice that repeatedly explored the same environments, and compared the collective coding properties across these two hippocampal subfields. We found that CA3 neurons were highly tuned to position from the first exposure to the environment, and that their spatial tuning was stable over weeks. By contrast, the place codes in CA1 were less informative and less stable when the environment was novel, but became markedly refined over the course of familiarization. Additionally, CA3 place codes were organized into distinct assemblies of highly correlated cells that were orthogonal to other assemblies. Surprisingly, as the environment became familiar, CA1 developed an assembly organization, similar to that of CA3. In both hippocampal subfields, cells that were members of such assemblies exhibited higher long-term spatial tuning stability. Overall, our results suggest that the organization of hippocampal CA3 into correlated cell assemblies yield refined and stable place codes that are gradually inherited by its downstream target CA1 throughout the learning process.

I-43. Dynamical learning of dynamics

Christian Klos
Yaroslav Felipe Kalle Kossio
Sven Goedeke
Aditya Gilra
Raoul-Martin Memmesheimer
University of Bonn

CKLOS@UNI-BONN.DE
YKALLEKO@UNI-BONN.DE
SGOEDEKE@UNI-BONN.DE
ADITYA.GILRA@GMAIL.COM
RM.MEMMESHEIMER@UNI-BONN.DE

The predominant paradigm of learning via rather slow synaptic plasticity is difficult to reconcile with the ability of humans and animals to quickly adapt to novel tasks. A potential solution to this problem is dynamical learning, i.e., learning of new tasks purely by dynamics, typically after appropriate weight pretraining ("learning to learn"). Existing work on dynamical learning does not, however, consider the learning of autonomous trajectories or dynamical systems and mostly relies on continuous teacher feedback.

Here we show how neural networks with fixed connection weights can learn to generate required trajectories or dynamical systems without teacher feedback. We consider recurrent neural networks of rate neurons with a signal output and a context output. Both are fed back into the network, initially together with an error input, given by the difference between signal and target output. During weight pretraining, the networks learn to minimize the error input and to generate a unique context signal for several dynamics from a task family. For the following dynamical learning, we fix the connection weights. Nevertheless, the error input alone suffices to teach unseen dynamics from the task family. By fixing the context signal, we can even remove the error input while the networks continue to generate a close-to-desired signal output. We illustrate the abilities of the proposed learning scheme by a variety of applications, ranging from autonomously generating sinusoidal oscillations to imitating chaotic Lorenz

attractor dynamics. For complex tasks, our dynamical learning outperforms even biologically implausible weight learning rules in terms of learning speed. We use dynamical systems theory to analyze the network mechanisms underlying the learning.

Our theory suggests that biological neural networks may well use dynamical learning for generating even complicated dynamics by imitation – a task relevant for, e.g., the quick learning of movements or short-term memory of sequences.

I-44. Using unsupervised machine learning to understand the role of mPFC in social dominance in mice

Kanha Batra¹
Nancy Padilla-Coreano²
Sebastien Hausmann²
Laurel Keyes²
Ila fiete³
Kay Tye²

KBATRA@ENG.UCSD.EDU
NPADILLA@SALK.EDU
SEBASTIEN.HAUSMANN@EPFL.CH
LKEYES@SALK.EDU
FIETE@MIT.EDU
TYE@SALK.EDU

¹University of California, San Diego

²Salk Institute for Biological Studies

³Massachusetts Institute of Technology

Although social hierarchies are central to successful group dynamics, the neural basis of social dominance remains unknown. Cross-species evidence suggests that the medial prefrontal cortex (mPFC) is involved in social dominance, but how the mPFC encodes this is unknown. Since dominant animals have priority access to resources, to facilitate statistical comparisons we designed a trial-based competition assay wherein mice compete for a reward that is signaled by a cue. Mice were trained to associate a cue with an Ensure reward and paired with a cage-mate of known social rank to compete against. Across the session, dominant mice (as defined via the tube test) won more rewards and occupied the reward port longer. Using wireless-electrophysiology we investigated how mPFC single-unit activity relates to dominance in this assay. Hierarchical clustering of mPFC activity during the cue-onset and reward port-entries (both self and competitor entries) revealed that clusters of neurons encode the outcome of the competition while others encode reward-seeking behavior. Using dimensionality reduction we analyzed neural trajectories of mPFC population activity during the cue and port-entries. At the population level, mPFC encodes reward-seeking behavior more during cues and the largest differences occurred for cue-onsets of winning vs losing trials; moreover, differences were present before the cue onset suggesting a state difference. To reveal if mPFC activity drives dominance using multiple states we used a Hidden Markov Model with logistic state transitions and emissions (“outputs”) as published by Calhoun et al in 2019. We used mPFC activity as input and dominance behavior as output. Using a 5-fold validation, we identified five possible hidden states that explained the behavior better than a static GLM. This study demonstrates that mPFC encodes social dominance behaviors and that there is evidence of multiple hidden states guiding behavior during social competition.

I-45. A theory of visual search for foveated visual systems

R Calen Walshe¹
Jared Abrams²
Wilson S Geisler²

CALEN.WALSHE@GMAIL.COM
JARED.ABRAMS@UTEXAS.EDU
W.GEISLER@UTEXAS.EDU

¹The University of Texas

²The Center for Perceptual Systems and The University of Texas at Austin

Searching the environment for specific targets is a critical capability for humans and other animals. Here we present a principled theory of visual search in arbitrary backgrounds that is based on the statistical properties

of natural images, the retinal falloff in resolution and sampling with eccentricity, the increase in intrinsic location uncertainty with retinal eccentricity, and the prior probabilities of target presence and location within the image. The important theoretical result is that near-optimal fixation selection for arbitrary targets added to natural backgrounds can be obtained with relatively simple, biologically plausible mechanisms. We test the theory here by predicting human covert search behavior in white noise backgrounds. Predictions are generated as follows. First, the effective prior probability distribution on target location is computed from the prior and the intrinsic location uncertainty. Second, the effective amplitude of the target (also dependent on retinal eccentricity) is computed and the target (if present) is added to the background. Third, template responses are computed at each image location by taking the dot product of a template (having the shape of the target) with the image and then adding a random sample of internal noise. Fourth, the responses are correctly normalized by the sum of the internal noise variance and the estimated variance due to external factors (the background statistics). Fifth, the normalized responses are combined with the effective prior on target location to obtain values proportional to the posterior probability. If the maximum of these values exceeds a criterion, the response is that the target is present at the location of the maximum. Human error rates averaged over spatial location correspond to the theoretical predictions. However, humans make a higher proportion of errors near the fovea, possibly due to underestimation of priors near the fovea or to allocating more decoding resources to the periphery.

I-46. Modeling the separate functions of feedforward and feedback pathways in the visual system

Grace Lindsay¹
Tom Mrsic-flogel²
Maneesh Sahani¹

GRACEWLINDSAY@GMAIL.COM
T.MRSIC-FLOGEL@UCL.AC.UK
MANEESH@GATSBY.UCL.AC.UK

¹Gatsby Computational Neuroscience Unit, University College London

²University College London

The ability of the feedforward path of the ventral visual processing stream to perform core object recognition has long been studied. The extensive feedback connections that go from higher visual areas to lower ones (such as those from V2 to V1) have been hypothesized to perform a separate function: iteratively refining feedforward representations to enhance perception under challenging visual conditions. Here, we use two different training paradigms for a recurrent convolutional neural network (CNN) to test whether a model with such a separation of functions best matches data. Models are either trained first to classify clean images using feedforward connections and then feedback is added for noisy images, or the network is trained with feedback connections on clean and noisy images from the start. Models trained both ways are able to use a single set of feedback connections to classify digits under four different types of challenging image degradation. However, only the model with a clear separation of functions captured all of the following experimental findings: behavioral consequences of interrupting feedback processing with a visual mask, the dynamics of responses to occluded stimuli, and physiological effects of cooling higher visual areas. We next asked how the representation of the input image in the model changes as a result of the influence from feedback. We find that feedback tends to boost both foreground and background pixels in the image, but increases the former more than the latter. This work helps to verify the theory that feedback refines feedforward processing, provides important groundwork on how to model feedback connections in a network that can perform visual tasks, and generates hypotheses about how feedback aids classification.

I-47. Naive artificial intelligence

Tomer Barak
Yonatan Loewenstein
Hebrew University of Jerusalem

TOMER.BARAK@MAIL.HUJI.AC.IL
YONATAN@HUJI.AC.IL

Inductive reasoning, the ability to find regularity in examples and use it to infer an underlying rule, is an important cognitive faculty. In humans, inductive reasoning ability is commonly tested in intelligence tests which are widely used as a candidate evaluation tool. Effective inductive reasoning is challenging because it requires the agent to simultaneously address two problems: identify the relevant features in the examples and use these features to identify the underlying rule. Previous computational models that solve intelligence tests have utilized predefined relevant features or predefined / pretrained rules. Here we present an inductive reasoning deep artificial neural network that simultaneously finds both the relevant features and the rules that govern a sequence of examples. We show that the network can be trained to infer the rules behind sequences using a relatively small number of examples. Moreover, we show that even a naive network, randomly initiated, can solve intelligence tests in the absence of any specific training, suggesting it can be used as a model of real-time inductive reasoning machine. These results demonstrate the potential "fluid intelligence" of artificial networks and help us identify the computational challenges of inductive reasoning in humans and animals.

I-48. Ignoring correlations causes a failure of retinal population codes under rod-mediated light levels

Kiersten Ruda¹
Joel Zylberberg²
Greg field¹

KIERSTEN.RUDA@DUKE.EDU
JOELZY@YORKU.CA
FIELD@NEURO.DUKE.EDU

¹Duke University
²York University

Regions downstream of the retina must extract visual information encoded in the activity of retinal ganglion cells (RGCs) from starlight to sunlight conditions. A key aspect of integrating responses from RGC populations is the role of noise correlations, which can enhance or degrade information depending on the structure of the signal and noise. Here we determine how light level-dependent changes in correlated activity impact decoding of retinal output. We recorded from populations of rat RGCs with a large-scale multielectrode array over light levels spanning rod to cone vision and decoded retinal activity using a generalized linear model (GLM). We find that responses are more correlated under rod-mediated conditions corresponding to moonlight, and decoders accounting for these correlations extract 100% more visual information than decoders that assume RGCs are independent. Strikingly, at these low light levels, assuming independence among a local population of RGCs produced worse decoding performance than decoding with a single RGC. In this way, we demonstrate a failure in decoding neural populations when noise correlations are substantial and ignored. To contrast, in cone-mediated light levels, accounting for correlated activity improves decoding by just 20% compared to decoders that ignore correlations, consistent with previous findings [1]. To generalize these results, we created a model of tuned, correlated neurons that highlights conditions under which correlation structure causes population codes to decode worse than single neurons. These findings elucidate that under strong and long-range signal and noise correlations, accounting for correlated activity not only improves visual processing but is necessary to take advantage of population codes. More generally, by showing that correlations affected by light adaptation significantly impact visual signaling, this work demonstrates the importance of context-dependent correlations in sensory processing.

I-49. Synaptic dynamics as a substrate for temporal learning

Alessandro Barri
Martin Wiechert
David DiGregorio
Institut Pasteur

ALESSANDROBARRI@GMAIL.COM
MARTIN.WIECHERT@GMX.DE
DAVID.DIGREGORIO@PASTEUR.FR

Sensations, thoughts and actions are temporally extended events. As such, telling the passage of time is a

fundamental property of brain functions [1]. How time is internally represented in the brain remains, however, elusive. It has been proposed that synaptic short-term plasticity (STP) could subservise temporal representations in neural circuits [2], but has not been linked to behaviourally-relevant circuit computations. The cerebellum is one of the best known circuits to mediate sub-second temporal learning and thought to be acting as an adaptive filter, which requires that sensory information be represented across the cerebellar granule cell (GC) population with differing temporal dynamics (temporal basis), although the mechanism by which this is achieved is unknown. It has recently been shown that sensory information driving GCs is mediated by synapses with diverse synaptic strength and STP, which segregate according to modalities and input pathways [3]. We hypothesize that this diversity of synaptic dynamics is necessary to generate a GC temporal basis. We created a rate-based perceptron model of the CC in which we introduced STP at mossy-fibre (MF)-GC synapses. Numerical simulations show that transient GC population activity elicited by MF-GC STP is sufficient for encoding the passage of time and can be harnessed to learn eye-lid responses. We used a reduced model to derive a simple relationship linking synaptic properties (release probability, Pr) and MF firing-rates to relevant learning timescales. The distribution of timescales available to the system is widened when Pr is positively correlated to pre-synaptic firing-rates. Indeed, we find that optimal temporal learning over multiple timescales is achieved when low Pr MFs exhibit low firing rates, high Pr MFs exhibit high firing rates. Given the ubiquity of STP throughout the brain our results are suggestive of general principles of temporal representations and learning in neural networks.

I-50. A two time-scale normative model of *C. elegans* sensory adaptation and behavior

Hamilton White¹
 Sruti Mallik²
 Baranidharan Raman²
 ShiNung Ching²
 Dirk Albrecht¹

HAWHITE@WPI.EDU
 SRUTI.MALLIK@WUSTL.EDU
 BARANI@WUSTL.EDU
 SHINUNG@WUSTL.EDU
 DALBRECHT@WPI.EDU

¹Worcester Polytechnic Institute

²Washington University in St. Louis

We examine the mechanisms underlying adaptation or habituation of sensory neural responses to repetitive stimuli. Specifically, we observe that early olfactory neurons adapt their peak response to repetitive stimulation with an exponential decay pattern. The rate of decay is proportional to the duty cycle of the stimulus. Since neural function and behavior are canonically linked, we hypothesized that behavioral responses to repeated chemosensory stimulation should adapt as well. Instead, behavior remained constant, suggesting that adaptation in primary sensory responses is matched with sensitization elsewhere in sensory circuits. To explore these dynamics, we posited two timescales of information embedded in these responses: immediate, actionable information associated with the current stimulus, and longer-term residual memory associated with prior stimulus history. We adopted a normative approach to investigate the neural dynamics that might give rise to such responses. By formulating the two timescales into drift-diffusion type decoders, we formally optimized sensory neural dynamics so as to efficiently produce short and long-term latent representations of the stimulus. The model reproduced experimental measurements of adaptation over minutes and neural inactivation over seconds. Succinctly, the goal of the sensory neurons is to drive both immediate actionable representations as well as memory traces that confer information about stimulus history. It turns out that for later trials in the experiment residual memory possesses more concrete information about stimulus identity. Mechanistically, this indicates that a qualitatively inferior immediate stimulus representation produced by reduced activity in the sensory neurons can be compensated by secondary dynamical processes downstream of the sensory neurons. Such dynamics may take the form of interneurons and/or modulatory processes. Thus, the model suggests that composition of information extracted from the stimulus over multiple timescales allow the organisms to maintain reliable behavioral responses despite sensory adaptation.

I-51. Removing uncorrelated noise in system neuroscience using deep interpolation

Jerome Lecoq
Michael Oliver
Natalia Orlova
Joshua Siegle

Allen Institute for Brain Science

JLECOQ@GMAIL.COM
MICHAELO@ALLENINSTITUTE.ORG
NATALIAO@ALLENINSTITUTE.ORG
JOSHS@ALLENINSTITUTE.ORG

Noisy data is a major impediment to any scientific effort. For example, a major limitation for in vivo two photon calcium imaging is the intrinsic presence of Poisson noise. As few fluorescent photons are detected per pixel, the inherent statistics of rare events cause large signal fluctuations. This greatly impairs our ability to detect the activity of individual neurons. Poisson noise is typically averaged out using a region of interest (ROI) paired with temporal binning. This approach, of course, limits the available temporal and spatial resolution. We reasoned that denoising filters could be entirely learned from the data itself. However, in the case of in vivo calcium imaging, it is not a simple matter of learning a mapping from a noisy recorded image to a clean version using deep neuronal networks, as noise-free training data is not readily available. To solve this problem, we combined a recent approach called Noise2Noise (Lehtinen et al., 2018) with the concept of data interpolation and designed a new general-purpose denoising algorithm called Deep interpolation that works without ground truth data. To demonstrate the effectiveness of this approach, we constructed an augmented UNET architecture designed to learn an interpolation function rather than a 1-to-1 mapping. The input to the network is a series of 60 frames from a calcium imaging experiment, 30 consecutive frames before and 30 consecutive frames after a single frame of the movie. The network was tasked with predicting the missing frame in the middle of those 60 frames. Because Poisson noise is uncorrelated, the network cannot possibly predict the noise and its only option is to predict the expected fluorescence of each pixel given the spatio-temporal information from previous and successive frames. We trained this network on 325,000 samples from a database of two-photon movies. Evaluation on test data shows excellent image restoration for every movie frame, while entirely preserving the temporal dynamics of the calcium signal.

I-52. Hearing-impaired deep neural networks replicate behavioral deficits of hearing-impaired humans

Mark Saddler
Jenelle Feather
Andrew Francl
Josh McDermott

Massachusetts Institute of Technology

MSADDLER@MIT.EDU
JFEATHER@MIT.EDU
FRANCL@MIT.EDU
JHM@MIT.EDU

Humans with hearing loss particularly struggle to understand speech in noisy environments, where current hearing aids offer little benefit. The development of more effective assistive devices is hindered by the difficulty of linking physiological and perceptual effects of hearing loss. Damage to different structures in the cochlea is known to alter peripheral signal processing, but connecting these changes to real-world listening difficulties has been limited by the lack of computational models capable of performing real-world tasks. Artificial neural networks optimized to perform auditory recognition tasks from simulated cochlear input have recently been shown to replicate aspects of human auditory behavior. Here, we extend this approach to investigate how outer hair cell (OHC) and auditory nerve fiber (ANF) loss (two common physiological effects of sensorineural hearing loss) can account for difficulties experienced by hearing-impaired listeners. We trained deep neural networks to recognize words from simulated healthy cochlear representations of speech in noise. We then substituted an alternative cochlear model with simulated sensorineural hearing loss and measured how network performance changed. Networks with either OHC or ANF loss replicated behavioral deficits of hearing-impaired listeners: speech recognition performance was degraded, especially at low SNRs. To investigate how plasticity in the central auditory system might allow hearing-

impaired listeners to adapt to their damaged cochleae, we also trained networks with impaired cochlear input. Networks optimized with OHC loss exhibited remarkably unimpaired performance, suggesting that a perfectly plastic auditory system could fully compensate for OHC loss (provided sounds are presented at an audible level via a hearing aid). Networks optimized with ANF loss also exhibited improved performance, but were still significantly impaired relative to normal-hearing networks, indicating that ANF loss produces a peripheral representation that is inherently less informative. The results illustrate how deep learning can provide insight into both normal and abnormal sensory function.

I-53. Steering neural population activity with adaptive closed-loop stimulation system

Sina Tafazoli¹
 Timothy Buschman¹
 Camden MacDowell^{1,2}
 Daniel Che¹
 Kate Letai¹

TAFAZOLISINA@GMAIL.COM
 TBUSCHMA@PRINCETON.EDU
 CAMDENM@PRINCETON.EDU
 ZCHE@PRINCETON.EDU
 KCLETAI@GMAIL.COM

¹Princeton University

²Rutgers RWJ Medical School

Stimulation of neural activity is important for causally testing the relationship between neural activity and behavior (Salzman et al 1990, Grosenick et al 2015). However, stimulation is often limited to a few stimulation sites, using a small set of predefined stimulation patterns. In contrast, recent large-scale recordings have shown information in the brain is represented in the distributed pattern of activity across the neural population. To causally manipulate such distributed representations across the population, new stimulation techniques are needed. This would provide new opportunities for interrogating neural circuits and will form the foundation for the next generation of brain stimulation devices. To begin to develop these techniques, we created a multi-dimensional closed-loop stimulation system that can causally control the activity of populations of neurons. Our approach uses machine learning algorithms to learn the stimulation patterns for navigating within the manifold of neural population activity. Here we show the system can learn to produce specific responses of populations of 5-15 neurons in the visual cortex of awake mice. Importantly, the system learned electrical stimulation patterns that reproduced natural, visually evoked responses. Using an adaptation paradigm, we show electrical stimulation adapted to the presentation of the associated visual stimulus, suggesting electrical stimulation patterns can create behaviorally meaningful representations. Altogether, our results provide the first step in developing a system to causally manipulate populations of neurons to induce specific behaviors.

I-54. Quantifying the role of divisive normalization in contextual modulation of neuronal variability

Dylan Festa¹
 Oren Weiss¹
 Amir Aschner¹
 Adam Kohn¹
 Ruben Coen-Cagli²

DYLAN.FESTA@EINSTEIN.YU.EDU
 OWEISS@MAIL.EINSTEIN.YU.EDU
 AASCHNER@MAIL.EINSTEIN.YU.EDU
 ADAM.KOHN@EINSTEIN.YU.EDU
 RUBEN.COEN-CAGLI@EINSTEIN.YU.EDU

¹Albert Einstein College of Medicine

²Albert Einstein College Of Medicine

Response variability is ubiquitous in cortex and can be modulated by external stimuli. Notably, variability is quenched by stimulus drive, and most reduced for visual neurons for high contrast, large stimuli. Recent work adopted the Gaussian Scale Mixtures (GSM) framework to explain these phenomena as resulting from a sampling-

based neural code in primary visual cortex (V1), suggesting higher response variability reflects larger uncertainty about the stimulus. The GSM framework also captures modulations of V1 firing rate by stimulus contrast, size, and adaptation, through a form of divisive normalization. However, the contribution of normalization to modulation of variability is not well established experimentally, primarily because the normalization signal cannot be measured directly. Here, using manipulations of stimulus size, we clarify how normalization affects variability in the GSM, and quantify this link in awake macaque V1 data, using the recently developed Ratio of Gaussian (RoG) model to estimate the normalization signal from the measured spike counts. In the GSM, normalization drives both stimulus modulation of mean responses and of their variability. We show analytically and with simulations that this reflects inference of a latent variable, called mixer, which represents overall stimulus energy and effectively acts as the divisive normalization signal: the mixer is bigger for larger stimuli, causing both response reduction and stabilization. We then extend the RoG model to measure size-dependent normalization in V1 responses. The RoG provides excellent fits to size-tuning data, and reveals that responses are more reliable for neurons that are more strongly normalized, consistent with GSM predictions. Our results indicate a key role for normalization in response stabilization by contextual stimuli, which could extend to several other experimental manipulations known to engage normalization, underlining the importance of unified, principled models of neural coding that account for both mean response and variability.

I-55. Dopamine firing activity codes reward expectation and drives motivation in a working memory task

Nestor Parga¹
 Stefania Sarno²
 Manuel Beiran³
 Jose Vergara⁴
 Rossi-Pool Roman⁴
 Ranulfo Romo⁴

NESTOR.PARGA@UAM.ES
 UNABUONAORA@GMAIL.COM
 MANUEL.BEIRAN@ENS.FR
 JVERGARA@EMAIL.IFC.UNAM.MX
 ROMANR@IFC.UNAM.MX
 RROMO@IFC.UNAM.MX

¹ Universidad Autonoma Madrid

² Aix-Marseille Universite

³ Ecole Normale Superieure

⁴ Universidad Nacional Autonoma Mexico

It is believed that dopamine (DA) provides a reward prediction error (RPE) signal and motivational drive. However little is known about how DA firing rates behave in cognitively effortful decision-making tasks. We investigated DA firing in monkeys performing a tactile frequency discrimination task that required maintaining a stimulus representation in working memory. Accuracy and reaction times (RTs) were affected by a bias contracting the perceived value of the first frequency to its mean, which also modulated the subjective difficulty of the stimulus classes. Segregating trials in two levels of RTs we found that the phasic DA responses in short- and long-RTs trials were significantly different; the same happened to the two mean delay-period DA activities evaluated at these two levels of RTs. Since RTs reflect motivation and trial difficulty, those differences did not immediately imply a role of DA in motivational drive. Then, to understand DA function we first investigated the activity dependence on the stimulation class (difficulty), finding that: (1) Responses to the second stimulus in short- and long-RT trials depended on the stimulation class in a way consistent with a RPE signal. (2) Responses to the first stimulus coded RPEs only in short-RT trials. (3) Phasic activities did not show any other stimulus dependencies (such as frequency tuning). (4) Delay-period activity was not affected by the first stimulus in any of the two RT levels. Finally, to disambiguate motivation from RPEs, we considered the correlation between phasic DA and RTs on trials of the same class (fixed difficulty) finding phasic DA engagement in motivated behavior on a trial-by-trial basis. Delay-period activity also correlated with the RT. Crucially, it also exhibited a ramp-like increase. Our results suggest roles of phasic DA in learning and motivation and shows that delay-period activity is a motivational signal, plausibly mediating cognitive control.

I-56. Top-down and bottom-up factors regulate cortical dynamics through Vip-Sst mutual inhibition

Daniel Millman
 Gabriel Ocker
 Michael Oliver
 Peter Ledochowitsch
 Jerome Lecoq
 Clay Reid
 Michael Buice
 Saskia de Vries

MILLMAN.DAN@GMAIL.COM
 GABEO@ALLENINSTITUTE.ORG
 MICHAEL@ALLENINSTITUTE.ORG
 PETERL@ALLENINSTITUTE.ORG
 JEROMEL@ALLENINSTITUTE.ORG
 CLAYR@ALLENINSTITUTE.ORG
 MICHAELBU@ALLENINSTITUTE.ORG
 SASKIAD@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Inhibitory interneurons play a major role in establishing the dynamics of cortical microcircuits. In layer 2/3 of cortex, vasoactive intestinal peptide-expressing (Vip) interneurons regulate feedback inhibition of pyramidal neurons through suppression of somatostatin-expressing (Sst) interneurons (Pfeffer 2013). Through this disinhibitory mechanism, Vip interneurons are believed to modulate network dynamics based on the behavioral state of the animal; for instance, Vip cells in mouse visual cortex are reliably active during periods of locomotion (Fu 2014). Vip cells also receive reciprocal inhibition from Sst cells, creating a circuit motif of mutual inhibition between Vip and Sst cells with unknown implications for cortical processing. Behaviorally, activation of Vip cells in visual cortex increases contrast sensitivity whereas activation of Sst cells decreases contrast sensitivity (Cone 2019). We hypothesized that Vip cells amplify responses of pyramidal neurons to weak and behaviorally-relevant stimuli. To this end, we investigated the influence of stimulus contrast and motor behavior (i.e. locomotion) on the visual responses of Vip, Sst, and pyramidal cells in mouse primary visual cortex. We recorded responses to drifting gratings at eight directions and six contrasts during calcium imaging of mouse Cre lines for Vip and Sst as well as pyramidal cells across cortical layers (Cux2: layer 2/3; Rorb: layer 4; Rbp4: layer 5; Ntsr1: layer 6). Sst cells responded exclusively at high contrast with weak direction selectivity, whereas Vip cells responded exclusively at low contrast with strong preference for front-to-back motion that is congruent with self-motion during locomotion. Layer 2/3, but not deeper layer, pyramidal cells responded more strongly at low contrast than high contrast and showed a slight, but significant, bias for front-to-back motion. The responses of all cell types increased during locomotion, with Vip and layer 2/3 pyramidal cells exhibiting greater enhancement at low contrast and Sst cells exhibiting greater enhancement at high contrast.

I-57. The anterior dorsomedial frontal cortex is causally involved in regulating the time constant of evidence accumulation

Thomas Luo
 Carlos D Brody
 Adrian Bondy
 Brian DePasquale
 Princeton University

ZHIHAOL@PRINCETON.EDU
 BRODY@PRINCETON.EDU
 ABONDY@PRINCETON.EDU
 DEPASQUALE@PRINCETON.EDU

Decision-making across multiple domains can be well described by the gradual accumulation of evidence over time, but the neural circuits performing this computation are not yet established. While widespread cortical areas are necessary during decision-making, the computations performed by any given area are not well understood. During a task that requires the gradual accumulation of auditory pulsatile information, a region of rat striatum was previously found to be causally involved in the accumulation process. Retrograde tracing from this striatal region showed dense projections from anterior dorsomedial frontal cortex (admFC, +4.0 mm AP; 0.8 mm ML). Bilateral pharmacological inactivation of admFC and computational modelling of the resulting behavioral deficits indicate a significantly shorter integration time constant after inactivation. Temporally specific optogenetic inactivation impaired performance in a task period well before decision commitment, a result consistent with admFC participating

in evidence accumulation. Neural population recording using Neuropixels probes showed robust and rapid onset of evidence-related dynamics in dorsal medial frontal cortex, in contrast to ventral medial frontal cortex. Finally, optogenetic inactivation of admFC impaired the encoding of preparatory choice signals in striatum. These results indicate admFC is a node in the circuit for evidence accumulation and suggest it may be involved in regulating the time constant of integration.

I-58. Non-linear matrix factorization methods for extracting calcium traces in moving *C. elegans* videos

Conor McGrory
Amin Nejatbakhsh
Erdem Varol
Columbia University

CPM2159@COLUMBIA.EDU
MN2822@CUMC.COLUMBIA.EDU
EV2430@COLUMBIA.EDU

In the past five years, new imaging techniques have made it possible to record 3D volumetric data from whole animals with high speed and single-neuron resolution. When used with calcium indicators, these techniques can be used to observe large-scale patterns of neural activity with high spatial and temporal precision. Extracting traces of neural activity from this data is a difficult problem that has been addressed using a variety of approaches. Recently, non-negative matrix factorization (NMF) methods have been proposed for solving this problem. These methods extract neural activity from imaging data by building a low-rank representation of the data, expressed as a product of two matrices that encode the spatial footprints of neurons and their temporal firing activity respectively. This general approach has proven to be very effective for denoising and demixing calcium signals in a variety of applications. However, one shortcoming of traditional NMF methods is their assumption that the spatial footprint of neurons is invariant across time. In whole-animal recordings of behaving *Caenorhabditis elegans* (*C. elegans*), this assumption is unrealistic due to the non-linear movement of the worm, causing NMF methods to perform poorly at extracting calcium signals. Tracking methods that explicitly model neuron motion also fail at this task, because they require high-fidelity detection of neurons, which is often not realizable in biological signal-to-noise regimes. Instead, to solve this problem, we propose using a method from the machine learning literature called Wasserstein-NMF (W-NMF), which deals with neuron motion by modeling it as a type of noise. We describe how this approach can be adapted to extract calcium traces from moving neurons in imaging data, and demonstrate its effectiveness on synthetic and real whole-brain imaging data of *C. elegans*.

I-59. Emergence of opposite neurons in a circuit for multisensory processing

Ho Yin Chau¹
Xueyan Niu²
Wenhao Zhang³
Tai Sing Lee⁴

HCHAU630@BERKELEY.EDU
XYNIU@CONNECT.HKU.HK
WENHAO.ZHANG@PITT.EDU
TAI@CNBC.CMU.EDU

¹University of California, Berkeley

²The University of Hong Kong

³University of Pittsburgh

⁴Carnegie Mellon University

Congruent (C) and opposite (O) neurons in the dorsal medial superior temporal (MSTd) and the ventral intraparietal (VIP) areas are involved in multisensory processing of visual and vestibular cues to infer heading (self-motion) direction. C neurons prefer congruent visual and vestibular cues of heading direction, integrating them in a near-optimal way to infer heading direction; in contrast, O neurons prefer opposite visual and vestibular cues, and recent theory suggests that the O neurons compute the disparity between these cues to perform causal inference, i.e. determining whether the two cues are produced by the same or different sources. To gain an understanding of how

a neural circuit can develop to perform causal inference in multisensory processing, the present study investigates the mechanisms and conditions under which opposite and congruent tunings can be learnt from multisensory inputs. Hebbian learning can be used to learn congruent tunings, but it does not work for opposite tunings because visual and vestibular cues tend to be consistent in our natural experience. We therefore propose a model in which C and O neurons are in competition, such that O neuron response is suppressed when the visual and vestibular cues are congruent and facilitated when they are opposite. This competition, mediated by inhibitory connections from C to O neurons, can be learned with an anti-Hebbian rule, while bottom-up excitatory connections from visual and vestibular cues can be learned with a Hebbian rule. We found that the learned network can reproduce nearly all the tuning properties and behaviors of the congruent and opposite neurons observed in experiments. Since recent theory suggests that O neurons might play a key role in causal inference or model selection, the circuit proposed might have broader implications beyond understanding multisensory integration.

I-60. Model selection approach for identifying distinct rhythmic entrainment profiles of CA1 interneurons

Pamela Riviere
Lara M Rangel

PDRIVIER@UCSD.EDU
LRANGEL@UCSD.EDU

University of California, San Diego

Neuronal networks receive and respond to many converging inputs. To spatiotemporally organize responses to various sources of information, networks rely on mechanisms capable of filtering, routing, and integrating inputs across cell assemblies. In the hippocampal CA1 subregion, a network critical for associative memory, inhibitory interneurons impose local temporal organization through periodic suppression, ensuring excitability at specific time windows that can be exploited only by exquisitely timed inputs[1]. The timing of CA1 interneuron activity is thus remarkably rhythmic, with spike patterns in distinct frequency bands reflecting their participation in repetitive rhythmic circuit interactions that constrain the subset of effective inputs. The hippocampal network undergoes substantial shifts in rhythmic processing state, yet to date, little is known about the rhythmic circuits that manifest in the CA1 during associative memory tasks, and the interneurons that contribute to shaping the rhythmic architecture of this network to facilitate selective input processing. We investigated CA1 interneuron spike timing relative to the phase of ongoing rhythms as rats performed a context-guided odor association task. Here, we adopt a statistical model selection approach to identify, for each interneuron independently, the simplest combination of local field potential (LFP) rhythms that accurately predicts spiking during correct performance. We find subpopulations of interneurons that exhibited spike timing that was best predicted by either theta[4-12Hz], beta[15-35Hz], low gamma[35-55Hz], or high gamma[65-90Hz] rhythms, effectively classifying interneurons according to engagement in single rhythmic circuits. Unlike the set of descriptive tools currently available to the field, this strategy allows us to weigh the relative contributions of multiple simultaneously occurring circular covariates to the activity of distinct interneuron populations, yielding insights into the plurality of rhythmic circuits that might support associative memory.

I-61. An olfactory pattern generator for functional analysis of neural circuit in *Drosophila* larva

Guangwei Si
Jacob Baron
Yu Feng
Aravinthan Samuel
Harvard University

GSI@FAS.HARVARD.EDU
JBARON@G.HARVARD.EDU
YU.FENG707@DUKE.EDU
SAMUEL@PHYSICS.HARVARD.EDU

For physical stimuli, it is straightforward to engineer patterns to stimulate the sensory circuits and probe the

functional properties. However, this framework is generally not applicable to olfaction, since olfactory stimuli lack a well-defined space. The typical olfactory stimuli, fixed odor panels of pure chemicals, natural scents or mixtures, establish preselected points in olfactory space to be tested in each experiment, which may induce biases when conducting functional analysis of the olfactory circuit. To systematically sample the olfactory sensory representation space, we present a stimuli method that project activity pattern directly onto the olfactory receptor neuron (ORN) ensemble. This method is based on the flexible and precise mixing of an optimized set of primary odors using microfluidics. The primary odors were derived for the *Drosophila* larva to separately target each free variable in the olfactory code, the activity of an individual olfactory receptor neuron type. Our device allows us to program and deliver any mixture of primary odors from all combinatorial possibilities during an experiment, allowing selection of any receptor neuron activity combination on demand. We combine this stimulus method with live imaging the olfactory representations in the *Drosophila* larva from receptor neurons to interneurons of the antennal lobe and Kenyon cells of the mushroom body. With a series of stepwise additive stimuli, we observed compartmentalized calcium activity and heterogeneous superposition property from an individual local interneuron, and various pattern selections from the activities of a population of Kenyon cells. This study will facilitate systematic studies on the population coding and neural processing in the olfaction circuit, and also motivate new frameworks to study the neural processing of unstructured signals.

I-62. Slow manifolds and output attractors as the substrate for working memory in recurrent networks

Elham Ghazizadeh¹
Lawrence Snyder²
ShiNung Ching²

ELHAM@WUSTL.EDU
LSNYDER@WUSTL.EDU
SHINUNG@WUSTL.EDU

¹Washington University in St Louis

²Washington University in St. Louis

Working memory enables the storage and manipulation of information over brief periods of time. This capability is required for many cognitive tasks such as reasoning and problem-solving. But what are the circuit mechanisms that underlie working memory? This issue remains enigmatic and a key broad area of neuroscientific inquiry. A prevailing theory suggests that memories are encoded in attractors of relevant cortical circuits. Each attractor represents a stationary pattern of neural activation that persists in the absence of exogenous input. But the dynamical nature of such attractors is unclear, especially in scenarios involving sequential, unstructured memory demands. In this study, we adopt a line of investigation drawn from machine learning by constructing recurrent, artificial networks to perform a simple yet potentially insightful working memory task: associating sequential stimuli. Concretely, the task for the network is to summate sequentially presented handwritten digits in a trial-based paradigm. We then study the emergent network dynamics and structure to shed light on the mechanistic basis of how this network achieves working memory. Specifically, we are interested in demystifying the encoding of memory representations in invariant features of neural dynamics, such as attractors. Our key finding is that memories are not actually encoded in attractors, but rather in slow, nearly invariant manifolds in the network phase space. Somewhat paradoxically, the outputs of the network (i.e., the summated value emitted at the end of the trial) are encoded as attractors, even though there is ostensibly no memory demand associated with this operation. These results suggest that working memory need not be encoded in dynamical attractors per se, but that there may be a nontrivial phase space geometry underlying such function. Such dynamics are compatible with recent observations that many neurons in working memory circuits 'drop out' of memory representations, with little consequence on subsequent task performance.

I-63. Probabilistic knowledge of event timing to optimize reward seeking behaviors

Eizaburo Doi
 Joshua Dudman
 HHMI Janelia Research Campus

DOIE@JANELIA.HHMI.ORG
 DUDMANJ@JANELIA.HHMI.ORG

When an agent interacts with noisy and dynamic environments, optimal exploitation requires integrating two different kinds of information about the timing of events: a) sensory evidence and b) probabilistic knowledge. Although this is a general Bayesian formulation, its formal application in complex tasks is non-trivial and often remains unsolved. Here we derive how an ideal agent should behave given the true probability model of the environment. Specifically, we examine a doubly stochastic behavioral paradigm in which “whether” and “when” a reward is delivered in each trial is governed by independent probabilistic processes. Since a reward will not necessarily be delivered on a given trial, the agent needs to decide how long to wait before abandoning the trial, which we call the leaving time. If a reward is delivered, the agent would react as soon as possible, and its reaction time should vary with the likelihood of reward timing. We use the sequential analysis to formulate the optimal decision rules for when to react and when to abandon. Importantly, probabilistic knowledge is time-varying and integrated with sensory evidence at every time step. While it is not obvious from the resulting formula, our numerical simulations predict behavioral data in both humans and mice remarkably well. Furthermore, in the zero-noise limit the probabilistic knowledge guiding reaction time can be reduced to a simple and interpretable formula which contains the so-called hazard function as a special case. To conclude, we provide a theoretical analysis which unifies perceptual decision making and probabilistic inference of timing in a principled manner and find its predictions consistent with behavioral data.

I-64. A framework for unifying and generalizing models of neural dynamics during decision-making

David Zoltowski¹
 Jacob Yates²
 Jonathan Pillow¹
 Scott Linderman³

ZOLTOWSKI@PRINCETON.EDU
 JACOBY8S@GMAIL.COM
 PILLOW@PRINCETON.EDU
 SCOTT.LINDERMAN@STANFORD.EDU

¹Princeton University
²University of Rochester
³Stanford University

An open question in systems and computational neuroscience is how neural circuits accumulate evidence towards a decision. fitting decision-making models directly to neural activity helps answer this question by providing a statistical test of how well the model explains neural responses, but current approaches limit the number of these models that we can fit. Here we propose a recurrent state-space framework for modeling population neural activity with decision-making dynamics. The framework unifies existing one-dimensional accumulator models such as the drift-diffusion and ramping models. It also enables generalizations to models with multi-dimensional accumulators, variable and collapsing bounds, non-constant boundary dynamics, and hybrid ramp-steps. To fit these models, we introduce an efficient variational Laplace-EM inference algorithm. We applied this approach to investigate neural dynamics in monkey parietal cortex during decision-making tasks. In a discrete-pulse task, we found that a two-dimensional accumulator model better captured the trial-averaged responses of a set of simultaneously recorded parietal neurons than a single accumulator model. Next, we identified a variable lower boundary in the responses of a parietal neuron during a random dot motion task. We expect that this framework will be useful for modeling neural dynamics in a variety of decision-making tasks.

I-65. Learning recurrent dynamic patterns in spiking neural networks

Christopher Kim
Carson Chow
National Institutes of Health

CHRISMKKIM@GMAIL.COM
CARSON.CHOW@NIH.GOV

Understanding the recurrent dynamics of cortical circuits engaged in complex tasks is one of the main questions in computational neuroscience. Most of recent studies train the output of recurrent models to perform cognitive or motor tasks and investigate if the recurrent dynamics emerging from task-driven learning can explain neuronal data. However, the possible range of recurrent dynamics that can be realized within a recurrent model after learning, particularly in a spiking neural network, is not well understood. In this study, we focus on investigating spiking network's capability to learn recurrent dynamics and characterize the learning capacity in terms of network size, intrinsic synaptic decay time and target decay time. We find that, by modifying recurrent synaptic weights, spiking networks can generate arbitrarily complex recurrent patterns if (1) the target patterns can be produced self-consistently, (2) the synaptic dynamics are fast enough to track the targets, and (3) the number of neurons in the network is large enough for noisy postsynaptic currents to approximate the targets. We examine spiking network's learning capacity analytically and corroborate the predictions by training spiking networks to learn arbitrary patterns and in-vivo cortical activity. Furthermore, we show that a trained network can operate in balanced state if the total excitatory and inhibitory synaptic weights to each neuron are constrained to preserve the balanced network structure. Under such synaptic constraints, the trained network generates spikes at the desired rate with large trial-to-trial variability and exhibits paradoxical features of inhibition-stabilized network. These results show that spiking neural networks with fast synapses and a large number of neurons can generate arbitrarily complex dynamics. When learning is not optimal, our findings can suggest potential sources of learning errors. Moreover, networks can be trained in dynamic regime relevant to cortical circuits.

I-66. Burst initiation in the preBotzinger Complex as synchronization spread- ing on a random graph

Valentin Slepukhin
Sufyan Ashhad
Jack Feldman
Alex Levine
University of California, Los Angeles

VALENTINSLEPUKHIN@PHYSICS.UCLA.EDU
SASHHAD@UCLA.EDU
FELDMAN@G.UCLA.EDU
ALEVINE@PHYSICS.UCLA.EDU

We show that experiments on the initiation of synchronized neuronal firing dynamics - that underlies the formation of rhythmogenic burstlets, i.e., rhythmic population output - in the preBotzinger Complex (preBotC) provide strong constraints on the connectivity of the underlying neuronal network. We modeled the neuronal firing dynamics on a variety of quenched random graphs. By comparing these results to experiments on the dynamics of burstlet formation in the preBotC induced by the synchronized excitation of ≤ 10 neurons, only Erdos-Renyi networks with possible addition of particular structures (experimentally observed triplet motifs) are quantitatively consistent with the data. Remarkably, other biologically common networks such as those with locally dominated connections or hierarchical networks with groups of central and peripheral neurons distinguished by connectivity cannot quantitatively account for the observed burstlet dynamics. In addition, burstlet dynamics provide constraints on the individual neuron model, which offer directly testable predictions at the single-neuron level. We conclude with a discussion of the dependence of the collective dynamics of the network as it relates to the number of neurons. These provide further quantitative tests of the validity of the underlying model, independent of network structure.

I-67. Evidence for multifaceted uncertainty-driven human exploration in a novel multi-armed bandit task

Angela Yu¹
 Dalin Guo¹
 florent Meyniel²

AJYU@UCSD.EDU
 DAG082@UCSD.EDU
 FLORENT.MEYNIEL@GMAIL.COM

¹University of California, San Diego
²NeuroSpin, CEA Saclay

While different forms of uncertainty have previously been suggested to explain how the brain weighs different sources of information in perception and learning, little is understood about their contribution in exploratory behavior. Here, we investigate the relative contributions of four factors – future (cumulative) reward potential, expected uncertainty, unexpected uncertainty, and random exploration (or unexplained variance), in addition to immediate reward – in driving exploration in a modified multi-armed bandit task. In contrast to a typical bandit task design, we introduce several novel components: (1) reward rates can undergo changes over time (at stochastic “change points”) with different frequency (high/low volatility conditions), (2) the timing of the change-points is coupled between two arms, and independent for the third arm, (3) the reward rates of the two temporally coupled arms are either highly correlated or uncorrelated in two separated conditions, (4) subjects occasionally report their beliefs about the reward rates and their confidence in this judgment. This particular design is motivated by our simulations showing that the contributions of the five hypothetical factors can be statistically distinguished based on the design. To model learning in the task, we utilize an ideal Bayesian observer model. Based on model fitting, comparison, and simulations, we report two major findings: (1) when choosing among the options, subjects significantly utilize all three of future reward potential, expected uncertainty, and unexpected uncertainty, in addition to a major effect of immediate reward and a minor one of random exploration, (2) a simpler policy that only takes immediate reward into account can only achieve slightly better performance (an improvement of .02 in probability of reward earned per trial) over the policy apparently utilized by subjects, indicating that humans pay a very small price in terms of reward in order to satisfy uncertainty-driven exploration.

I-68. The dynamic repertoire of neural activity across the forebrain: the view from the single neuron

Daniel Levenstein¹
 John Rinzel¹
 Gyorgy Buzsaki²

DL2820@NYU.EDU
 RINZELJM@GMAIL.COM
 GYORGY.BUZSAKI@NYULANGONE.ORG

¹New York University
²New York University Langone Medical Center

Heterogeneity is the rule, rather than the exception, in the brain. Functionally-relevant heterogeneity is seen, for example, in mean firing rates or between cell types. Further, neurons in different regions have different spiking properties depending on behavioral contexts (e.g. well-trained tasks vs sleep). Often experiments from single regions/contexts provide a picture of neural activity which is then extrapolated to a global picture of brain function. For example, in many experiments during which cortical cells are driven by tuned stimuli, the activity of individual neurons is considered to be a stochastic spike generation process following a time-varying rate, while in others precise timing is considered with respect to ongoing oscillations. A unifying framework is needed to make sense of these disparate pictures of neural activity. To characterize the repertoire of neuronal activity, we analyzed the statistics of inter-spike intervals (ISIs) from long duration (4-12hr) electrophysiological recordings in multiple cortical and limbic regions (500-1500 cells/region) during spontaneous behavior and sleep. We find that neurons across the forebrain have a limited repertoire of spiking modes, each with ISIs at distinct timescales and spiking (ir)regularity. Excitatory neurons spend the majority of time in a neural “ground state”: a low rate (~0.1-2Hz) mode characteristic to each neuron with Poisson-like spiking. On top of the ground state, neurons show region-, and state-specific “activated states”: higher rate modes (>1Hz) at specific timescales shared across neurons in the

local network. Activated states tend to have more regular spiking and can be coupled to oscillations or evoked by “in-field” behavior/stimuli. Inhibitory neurons, on the other hand, have consistent dynamics across regions: their rate follows that of the local population, and is coupled to local gamma oscillations during periods of high population rate. Together, these results present a unifying picture of distinct neural activity patterns across brain regions and contexts.

I-69. Deep reinforcement learning of extra-synaptic control in *C. elegans* chemosensory response

Jimin Kim
Eli Shlizerman
University of Washington

JIMINKIM9292@GMAIL.COM
SHLIZEE@U.WASHINGTON.EDU

C. elegans is known to adapt its course according to chemosensory stimuli. It was shown that extra-synaptic currents (neuromodulators) play an important role in such behavior. While the entire *C. elegans* connectome has been mapped, extra-synaptic currents are yet to be included in it. We thereby propose to discover extra-synaptic interactions through a combination of Deep Q-Network (DQN) learning along with *C. elegans* neuromechanical model simulating chemosensory responses. Specifically, we utilize a recently introduced computational model which simulates the entire nervous system with muscles and body and show that it can generate aversive chemosensory responses to stimulation of CO₂ and O₂ associated sensory-neurons similar to those observed in in-vivo experiments. We then set the DQN method to seek for extra-synaptic currents which modulate the responses from aversion to attraction, as described in experiments. The DQN method succeeds to find plausible controls of extra-synaptic currents in both chemosensory responses. Furthermore, DQN modulated network dynamics result in activity that is consistent with the activity measured in-vivo. Our results show that combining deep reinforcement learning in conjunction with a computational neuro-mechanical model can assist in inference of external control inputs associated with particular behaviors.

I-70. Inverse rational control in continuous problems

Minhae Kwon¹
Saurabh Daptardar²
Paul Schrater³
Xaq Pitkow⁴

MINHAE.KWON@RICE.EDU
SAURABH.DPTDR@GMAIL.COM
SCHRATER@UMN.EDU
XAQ@RICE.EDU

¹Rice University

²Google Inc

³University of Minnesota

⁴Baylor College of Medicine, Rice University

A fundamental question in neuroscience is how the brain processes sequences of ambiguous sensory information to plan and generate actions. This is naturally formulated as a control or reinforcement learning problem under partial observations, where an animal must estimate relevant latent variables in the world from its evidence, anticipate possible future states, and choose actions that optimize expected reward. This problem can be solved by the control theory which allows us to find the optimal actions for a given system dynamics. However, animals often appear to behave suboptimally. Where does this discrepancy come from? We hypothesize that animals have their own internal model of the world which may not be always correct, but they still behave rationally, i.e., choose actions with the highest expected subjective reward. Inverse Rational Control (IRC) [1] has been proposed to learn the agent’s internal model that best explains their behaviors in a task described as a Partially Observable Markov Decision Process (POMDP). In [1], the authors mainly focus on a foraging task in discrete belief and action spaces. In this abstract, We further extend the IRC to continuous problems in order to address

more naturalistic complex neuroscience tasks. In our approach, we first learn optimal policies generalized over the entire model parameter space. In order to address continuous state, action and parameter spaces, we use a deep reinforcement learning method called Deep Deterministic Policy Gradient (DDPG) [2] to optimize the policy. In order to maintain the interpretability of a deep learning-based approach, we provide the model parameters as additional inputs to actor-critic networks of DDPG such that the ensemble of optimal policies over the parameterized manifold of models can be trained. We then find the model parameter that maximizes the likelihood over observation and action trajectories computing gradient ascent. Our proposed method is able to recover the true model of simulated agents within theoretical error bounds given by limited data. This approach provides a foundation for interpreting the behavioral and neural dynamics of animal brains.

I-71. VIP interneuron contributions to state-dependent sensory processing

Katie Ferguson¹
Calvin Fang¹
Jessica Cardin²

KATIE.FERGUSON@YALE.EDU
CALVIN.FANG@YALE.EDU
JESS.CARDIN@YALE.EDU

¹Yale School of Medicine

²Yale University

Cortical activity patterns dynamically reflect not only changing sensory inputs, but also behavioral state (e.g., quiet wakefulness vs. active locomotion). Several populations of GABAergic interneurons (INs) are strongly regulated by neuromodulatory input during distinct behavioral states, and thus may be key contributors to flexible cortical function. One prominent model for state-dependent inhibitory interactions highlights the potential for vasoactive intestinal peptide-expressing INs (VIP-INs) to activate upon arousal, suppress the activity of downstream INs (SST-INs), and consequently disinhibit PNs (Fu et al., 2014). However, cortical VIP-INs have complex patterns of connectivity, including both synaptic connections to excitatory (E) pyramidal neurons (PNs) and reciprocal inhibitory (I) connections with other INs. Current models largely fail to account for the complex and context-dependent I-I interactions, leaving the role of distinct IN populations within the active cortical network largely unexplored.

To address these questions, we quantified the moment-to-moment statistics of activity across multiple IN populations and used both chronic and acute manipulations to examine the state- and context-dependent role of I-I interactions. We expressed the genetic calcium indicator gCaMP6 in VIP-INs and, using intersectional genetic tools, in their postsynaptic target SST-INs and PNs. We used large-scale two-photon dual population imaging and high-throughput analyses to assess simultaneous patterns of VIP-IN and SST-IN (or PN) activity in the primary visual cortex of awake mice across behavioral states and cortical layers.

Our data suggest a complex spatiotemporal pattern of I-I interactions within the local cortical circuit that is consistent with state- and context-dependent contributions of distinct IN populations in sensory processing. We find that VIP-INs contribute to state-dependent regulation of functional connectivity and make an unanticipated contribution to SST-IN visual response properties. Our findings suggest that the unidirectional linear disinhibitory circuit model is not sufficient to explain the impact of VIP-IN activity on the network during visual encoding.

I-72. Automatically detecting single cell ion channel events of neural activity using generative adversarial and recurrent neural networks

Numan Celik
Fiona O'Brien
Yalin Zheng
Frans Coenen
Richard Barrett-Jolley
University of Liverpool

NUMAN.CELIK@LIVERPOOL.AC.UK
FMOBB@LIVERPOOL.AC.UK
YZZHENG@LIVERPOOL.AC.UK
FSCOENEN@LIVERPOOL.AC.UK
RBJJ@LIVERPOOL.AC.UK

The paraventricular nucleus of the hypothalamus (PVN) is one of the most important autonomic control centres in the brain, with roles including (but not limited to) controlling cardiovascular, thermoregulatory and stress responses. Neurones in the PVN are home to a vast array of ion channels which control cell excitability, although the expression and interplay of PVN ion channels remains largely unknown. We have identified TRPV4 channels to play an important physiological role in the PVN, however, generating TRPV4 single channel recordings from PVN neurones is slow and the analysis is laborious. In addition, often large datasets of channel activity are required to build mathematical models. To overcome this, we developed a novel approach to generate semi-synthetic TRPV4 channel gating using deep generative models. First, we record pharmacologically identified TRPV4 channels from PVN neurons and build a kinetic profile of the gating. Single channels gate in a stochastic Markovian manner and therefore a large collection of semi-synthetic records was simulated using generative adversarial networks (GAN). We then propose a recurrent convolutional neural network (RCNN) model for training and validating against additional single channel recordings of TRPV4 channels from PVN neurones. The network model is able to automatically idealise single channel activity more accurately and faster than other traditional methods such as threshold crossing or segmental K-means (SKM). In our model, there are no pre-parameters that need to be set, including the baseline, channel amplitude and number of channels present. We believe this is the first use of deep learning to analyse patch-clamp single molecule recordings and such methods may revolutionise the unsupervised automatic detection of ion channel and other single-molecule transition events in the future.

I-73. Temporal compression as a mechanism for switching network operation modes

David Tingley
Gyorgy Buzsaki
New York University

DAVIDTINGLEY2@GMAIL.COM
GYORGY.BUZSAKI@NYUMC.ORG

We will present data that demonstrates a single neural circuit—hippocampal neurons sending convergent projects onto lateral septal neurons—which can alternate between operating modes. In one state, the network transmits theta-sequence content and spatial information is transformed into a rate-independent phase code for spatial position. In the second state, a temporally compressed sequence (i.e. replay or preplay) is selectively filtered out so that only the magnitude of the event is transmitted, in a manner that is independent from the sequence content. Thus, the data suggest that this network has two very distinct operating modes, where the 'content' being transmitted is different in different states. We will also provide a single spiking neural network of this brain circuit that is able to recapitulate all features of the data and make predictions as to the degree of convergence and compression required for this mode switching.

I-74. Local changes, global effects: how localised is the effect of deformed boundaries on grid cells?

Marino Krstulovic¹
 Pranoy Chaudhuri-Vayalambone*¹
 Marius Bauza²
 Michael Rule¹
 Pauline Kerekes¹
 Timothy O'Leary¹
 Julija Krupic³

MK823@CAM.AC.UK
 PC518@CAM.AC.UK
 M.BAUZA@UCL.AC.UK
 MER49@CAM.AC.UK
 PJK57@CAM.AC.UK
 TSO24@CAM.AC.UK
 JK727@CAM.AC.UK

¹University of Cambridge

²University College London

³Cambridge University

The hippocampal formation plays a key role in learning, spatial memory and navigation. Two major functional cell types in the hippocampal formation are: place cells, which are active in restricted portions of the environment; and grid cells, which show multiple firing fields arranged in a hexagonal lattice. Previous studies have shown that grid cell symmetry is shaped by the geometry of the enclosure. However, the exact nature of this dependence remained elusive, primarily due to the limited size of previously used enclosures. Moreover, the underlying neural mechanism is still not understood, partly due to the limited number of relevant neurons that can be recorded during experiments. The present experiment tackles these issues by simultaneously recording from multiple neurons in the medial entorhinal cortex (mEC), CA1 and CA3 using Neuropixels probes or multi-tetrode micro-drives in rats freely foraging within different large environments (~8 m²). Consistent with previous work, local changes to the configuration of the enclosure caused the largest shifts in grid fields near the reconfigured wall. Curiously, our pilot data suggest that patterns appeared most stable near the centre of the enclosure, and were not anchored to the stationary wall as previously thought. This indicates that grid anchoring points may depend on enclosure size and that, in larger environments, the grid may be stabilised by sensory cues and not by border cell inputs. Place cells' responses were more variable, with some place fields shifting, remapping or appearing/disappearing. We used a simple model of an agent running along a linear track to investigate how different strategies for integrating idiothetic and allothetic position estimates could distort a "one-dimensional" grid cell's firing pattern. Linking the behaviour of the model to the experimental results will allow us to determine what strategies the hippocampal cognitive map might use to integrate idiothetic and allothetic information.

I-75. flexible recruitment of memory-based choice representations by human medial-frontal cortex

Juri Minxha¹
 Ralph Adolphs²
 Stefano Fusi³
 Adam Mamelak¹
 Ueli Rutishauser¹

JM4833@COLUMBIA.EDU
 RADOLPHS@HSS.CALTECH.EDU
 SF2237@GMAIL.COM
 ADAM.MAMELAK@CSHS.ORG
 UELI.RUTISHAUSER@CSHS.ORG

¹Cedars Sinai Medical Center

²California Institute of Technology

³Columbia University

Behavior in complex environments requires decisions that flexibly combine stimulus representations with context, goals, and memory. Two key aspects of such cognitive flexibility are the retrieval of information from memory when needed, and the ability to selectively utilize relevant information depending on task demands. These aspects depend on the medial frontal cortex (MFC) and the medial temporal lobe (MTL), but it is yet unclear how these structures interact during memory retrieval. To study this, we recorded single neurons in MFC and MTL while human subjects alternated between making memory and categorization-based decisions. In our popula-

tion analysis, we find that persistently active MFC cells encode the current task settings, allowing the subject to produce the correct response on the upcoming trial. Furthermore, we show that the quality of these representations has direct consequences on the ability of the subject to make fast, contextually-correct decisions. We compared the neural representations of choice, familiarity, and image category across tasks using cross-condition generalization performance. We found a striking difference between brain areas: In the MTL, representations of stimulus familiarity and stimulus category were task-demand independent, whereas in the MFC, representations of stimulus familiarity, stimulus category, and choices were all highly sensitive to task demands. Lastly, choices requiring memory retrieval selectively engaged phase-locking between theta-frequency band oscillations in MTL and MFC neurons. This work reveals a mechanism for selectively engaging memory retrieval and shows that unlike perceptual decision-making, memory-related information is only represented in medial frontal cortex when choices require it.

I-76. Closing the loop: recurrent neural networks that interact with their environment through reward-modulated learning

Robert Rosenbaum¹
Ryan Pyle²

RROSEN1@ND.EDU
RYANPYLE1@GMAIL.COM

¹University of Notre Dame
²Rice University

Recurrent neural networks are widely used to model motor learning, but are typically trained using supervised learning rules on function approximation tasks: Produce a target time series in response to an input time series. Supervised learning rules are unrepresentative of biological motor learning because they assume the network has access to an explicit target output. In a reaching task, for example, the target is related to hand position, but network output models some representation of muscle activations. RNNs with supervised learning rules would need the mapping from hand position to muscle activation provided to them. In the parlance of control theory, supervised learning rules must be provided an inverse model.

Previous work proposes biologically parsimonious reward-modulated learning rules for RNNs, which can learn inverse kinematics, but we find that they are susceptible to catastrophic errors on non-trivial tasks unless the correct target output is provided during testing, known as teacher forcing. Supervised learning rules are also susceptible to these errors in the presence of noise.

We propose that errors in the absence of teacher forcing arise from a common, implicit assumption that RNNs do not receive sensory feedback. This is akin to solving a task blindfolded and numb. We demonstrate that, unlike previous models, reward modulated RNNs with closed sensorimotor loops can solve complex tasks like the well-known benchmark cart-pole without teacher forcing. Closed-loop, reward-modulated RNNs learn in a fundamentally different way than open-loop RNNs and can inform the development of more powerful and realistic approaches for understanding how animals learn to interact with their environment.

I-77. Which comes first: correlated pre- and post-synaptic activity or associative plasticity?

Aaron Milstein¹

Katie Bittner²

Yiding Li³

Christine Grienberger³

Ivan Soltesz¹

Jeffrey Magee³

Sandro Romani²

AARONMIL@STANFORD.EDU

BITKAT@GMAIL.COM

YIDING.LI@BCM.EDU

CHRISTINE.GRIENBERGER@BCM.EDU

ISOLTESZ@STANFORD.EDU

JEFFREY.MAGEE@BCM.EDU

ROMANIS@JANELIA.HHMI.ORG

¹Stanford University

²HHMI Janelia Research Campus

³Baylor College of Medicine

According to standard models of synaptic plasticity, correlated activity between connected neurons changes synaptic strengths to store associative memories. Here we directly tested this hypothesis in vivo by manipulating postsynaptic activity and measuring changes in stimulus selectivity. We found that dendritic calcium spikes rapidly reshape the spatial tuning of hippocampal place cells via bidirectional synaptic plasticity. To account for the magnitude and direction of plasticity at each synaptic input, we evaluated two models – one standard model that depended on pre- and post-synaptic correlation, and an alternative model that depended instead on the initial synaptic strength of each input prior to plasticity. While both models fit the data, they predicted opposite outcomes of a perturbation experiment, which ruled out the standard correlation-dependent model. Finally, network modeling suggested that this form of bidirectional synaptic plasticity enables population activity, rather than pairwise neuronal correlations, to drive plasticity in response to changes in the environment.

I-78. Angular and linear velocity coding in the parahippocampal circuit

Davide Spalla¹

Alessandro Treves¹

Charlotte Boccara²

DSPALLA@SISSA.IT

ALE@SISSA.IT

CHARLOTTE.BOCCARA@MEDISIN.UIO.NO

¹SISSA– Scuola Internazionale Superiore di Studi Avanzati –

²University of Oslo

Successful navigation depends on the integration of a wide breadth of information to construct an updated map of allocentric space at any given time. Information integration is implemented in interconnected subareas of the hippocampal region through neurons coding for specific navigational features such as position (place cells), direction (head direction cells), metrics (grid cells), boundaries (border cells) or linear velocity (speed cells). While self-motion signals are thought to be crucial for updating both position and direction, where and how such signals are integrated remains largely elusive. To explore this question, we analysed 1437 cells recorded from all layers of the presubicular (PrS, 605 cells), parasubicular (PaS, 436 cells) and medial entorhinal (MEC, 396 cells) cortices of rats freely exploring open fields. Here we report, for the first time, coding for angular velocity in a significant proportion of parahippocampal cells (MEC: 9.8%; Prs: 8.4%; PaS: 10.8%). This modulation followed a linear response to the instantaneous angular velocity of the animal, either in the clockwise or counterclockwise direction. Moreover, our analyses revealed that speed coding extends beyond MEC and that about one fifth of PrS and PaS cells responded to linear velocity (MEC: 16.7%; Prs: 20.6%; PaS: 19.8%). Additionally, while many cells were modulated by only one feature, some showed conjunctive tuning. For example, 14.4% of angular velocity cells were also coding for linear velocity. The discovery of an angular velocity signal downstream of the midbrain nuclei provides substantial support for the attractor network models of head direction coding. The presence of a widespread and distributed code for self-motion, that can sustain and update positional and orientational coding, will be essential for the theoretical understanding of spatial information integration.

I-79. Neural networks for sorting neurons

JinHyung Lee¹

Catalin Mitelut¹

Ian Kinsella¹

Shenghao Wu²

Eleanor Batty¹

Liam Paninski¹

Hooshmand Shokri¹

Ari Pakman¹

Yueqi Wang¹

Nishchal Dethe¹

Kevin Li¹

Eduardo Blancas Reyes¹

Alexandra Tikidji-Hamburyan³

Georges Goetz³

E.J. Chichilnisky³

David Carlson⁴

JL4303@COLUMBIA.EDU

MITELUTCO@GMAIL.COM

IAK2119@COLUMBIA.EDU

SHENGHAW@ANDREW.CMU.EDU

ERB2180@COLUMBIA.EDU

LIAMPANINSKI@GMAIL.COM

HOOSHMAND@CS.COLUMBIA.EDU

ARIPAKMAN@GMAIL.COM

YUEQI.WANG.PKU@GMAIL.COM

ND2506@COLUMBIA.EDU

KL2918@COLUMBIA.EDU

EB3079@COLUMBIA.EDU

ALEXTH.NN@GMAIL.COM

GGOETZ@STANFORD.EDU

EJ@STANFORD.EDU

DAVID.EDWIN.CARLSON@GMAIL.COM

¹Columbia University

²Carnegie Mellon University

³Stanford University

⁴Duke University

Spike sorting is a critical step in extracting neural signals from multielectrode array (MEA) data. Unfortunately, currently available packages still require significant manual oversight to achieve accurate spike activity extraction from MEA data. Here we present several new techniques that make spike sorting more robust and accurate, while preserving scalability for large MEAs.

We begin with a simple generative model of MEA voltage data: we sample spike templates, superimpose these templates at random times and locations, and add spatiotemporally correlated noise. Then we use an arbitrarily large number of samples from this generative model to train neural networks to detect, denoise, and cluster spikes. Importantly, this training procedure does not require users to perform any laborious hand-labeling of individual spikes; instead, we simply require a curated library of spike templates (which in many cases is already available from previous experiments, or may be generated de novo). The resulting clustering approach is Bayesian, and is able to infer a full posterior distribution over the cluster assignments and the number of clusters present in the data. finally, we recover collided waveforms via a novel iterative (non-greedy), temporally super-resolved deconvolution approach; this deconvolution step “cleans” the effects of collisions, significantly decreasing the effective noise level and unmasking units that previously could not be separated.

We apply the new pipeline to data recorded in the primate retina, where high firing rates and highly-overlapping axonal units provide a challenging testbed; in addition, the well-defined mosaic structure of receptive fields in this preparation provides a useful quality check on any spike sorting pipeline. We develop new methods for comparing the performance of multiple spike-sorters in the absence of ground truth, and find that the proposed methods improve over the state-of-the-art on real, simulated, and hybrid MEA data with > 500 electrodes.

I-80. Strongly correlated spatiotemporal encoding and simple decoding in the prefrontal cortex

Ehud Karpas¹
 Roozbeh Kiani²
 Ori Maoz¹
 Elad Schneidman¹

UDIKARPAS@GMAIL.COM
 ROOZBEH@NYU.EDU
 ORIMAOZ@GMAIL.COM
 ELAD.SCHNEIDMAN@WEIZMANN.AC.IL

¹Weizmann Institute of Science

²New York University

The analysis of neural codes has often focused on characterizing an encoding dictionary, mapping stimuli into neural responses, and a decoding one, mapping neural activity into perceptual or behavioral outcomes. While encoding and decoding are strongly coupled, their respective dictionaries may have very different structures and properties, since encoding needs to convey relevant information efficiently under constraints, whereas decoding needs to provide a concise representation of the stimulus to guide a decision or action. Here, we investigate population encoding and decoding in the prefrontal cortex of monkeys performing a visual discrimination task. We used statistical models to map the dynamics of population spike patterns of groups of up to 100 simultaneously recorded units. We found that in the different parts of the task (epochs), the population spike patterns (codebooks) were strongly shaped by both spatial and temporal correlations among neurons: the observed sequences of spike patterns over 300 ms were between 10^5 - 10^8 times more likely to appear than predicted by models that ignored these correlations and relied only on time-dependent firing rates (PSTH-based models). Further, temporal sequences of population spike patterns had strong history-dependence in all epochs, but differed in their relative strength and characteristic timescales, suggesting that different computational dynamics govern different epochs. Surprisingly, despite these strong correlations, decoding performance of models that ignored correlations was on par with spatiotemporal correlation-based models. The difference in the role of correlations in population encoding and decoding suggests that complex encoding in the prefrontal cortex is geared to simplify decoding such that downstream regions do not have to learn upstream correlations. Such encoding schemes create robust codes that can be learned easily and quickly.

I-81. Measuring human probabilistic segmentation maps

Jonathan Vacher¹
 Pascal Mamassian²
 Ruben Coen-Cagli³

JONATHAN.VACHER@EINSTEIN.YU.EDU
 PASCAL.MAMASSIAN@ENS.FR
 RUBEN.COEN-CAGLI@EINSTEIN.YU.EDU

¹Albert Einstein College of Medicine

²Ecole Normale Supérieure de Paris

³Albert Einstein College Of Medicine

Visual segmentation is a core function of biological vision, key to adaptive behavior in complex environments. Foundational work identified Gestalt principles of segmentation, e.g. grouping by similarity, proximity and good continuation and revealed that visual cortical neurons are sensitive to those cues. Early models inspired by the feedforward cortical architecture describe texture-based human segmentation as the process of comparing the summary statistics of low-level visual features across space. Indeed the summary statistics representation is the most prominent model of naturalistic texture perception, yet it has been challenged precisely because it does not fully capture the influence of segmentation. Here we consider the alternative view that, due to image ambiguity and sensory noise, perceptual segmentation requires probabilistic inference. This view is consistent with reports that humans combine multiple segmentation cues near-optimally in artificial displays, and that Gestalt laws reflect optimization to natural image statistics. The probabilistic approach is also widespread in computer vision algorithms for unsupervised segmentation but has not been used to model perceptual segmentation. We present new experiments that for the first time measure perceptual segmentation maps and their variability, allowing us to test the probabilistic inference hypothesis, and compare it quantitatively to summary statistics models. We

use composite textures, with segments characterized by different statistical relations between features. Optimal probabilistic inference assigns pixels to segments by evaluating which of those relations better explains the observed features (generative model), as opposed to comparing summary statistics at different locations (feature discrimination). We find the generative model best captures our data, and perceptual variability reflects image uncertainty beyond sensory noise. We also demonstrate the approach on natural images, which will allow testing more sophisticated segmentation algorithms. Our results provide a normative explanation of human perceptual segmentation as probabilistic inference and demonstrate a novel framework to study perceptual segmentation of natural images.

I-82. Task engagement selectively enhances population discrimination of behavior-relevant categories in primary auditory cortex

Charles Heller
Daniela Sadari
Mateo Lopez Espejo
Stephen David

HELLERC@OHSU.EDU
DANIELA.SADERI@GMAIL.COM
LOPEZESP@OHSU.EDU
DAVIDS@OHSU.EDU

Oregon Health and Science University

Previous work in primary auditory cortex (A1) has shown that responses of single neurons to sound are modulated by active task engagement, such that neural discriminability of task-relevant sound features is enhanced during behavior. While this work is suggestive of changes in the neural population code, no previous study has directly addressed how task engagement modulates population coding on single behavioral trials. Further, the effect of fluctuations in global arousal on coding of auditory stimuli has not been dissociated from task engagement effects. This dissociation is important, as changes in arousal themselves modulate neural activity and correlate with transitions between active and passive listening. We address these questions by recording from populations of neurons in A1 while head fixed ferrets perform a simple go / no-go tone versus noise discrimination task. We simultaneously monitor global arousal levels using measurements of pupil size. As expected, we find that task engagement and arousal both modulate neural responses to task stimuli. Using a population decoding approach, we demonstrate that increases in arousal, independent of task engagement, enhance discrimination of all sounds in A1. In contrast, task engagement selectively enhances the single-trial discrimination of target versus distractor sound categories but not discrimination between different distractors. Preliminary analysis suggests that the size of the task engagement effect is correlated with behavioral performance, suggesting that enhanced target representation in A1 may contribute to sound-driven behavior.

I-83. Non-normal amplification in the barrel cortex connectome

Gadi Naveh¹
Haim Sompolinsky²

GADIMINTZ@GMAIL.COM
HAIM@FIZ.HUJI.AC.IL

¹Hebrew University of Jerusalem

²The Hebrew University of Jerusalem and Harvard University

With the availability of connectomics datasets, it remains unclear how to understand their implications on dynamics and functionality. We study a previously suggested dynamical feature of biological connectivity matrices: their ability to support transient amplification of incoming signals under linear rate dynamics. A necessary condition for the existence of such amplification is a non-normal connectivity matrix i.e. one that possesses non-orthogonal eigenvectors. Non-normality has been previously suggested as a desirable mechanism for amplification of sensory signals [Murphy & Miller, 2009], where fast response times are required since it avoids dynamical slowing resulting from critical amplification (due to proximity of eigenvalues to instability). We quantify here for the first time the extent of non-normality found in a realistic biological network. We analyze anatomically constrained connec-

tivity data of the rat barrel cortex, containing the recurrent connectivity matrix within a cortical column as well as thalamic feedforward input, at the cell-type resolution. We quantify several dynamical amplification measures of the neuronal architecture implied by these connectivity data, corresponding to different input temporal and spatial structures. We explore the parameter space under two “budgets”: overall synaptic strength and maximum allowed dynamical slowing. Our findings are twofold: (1) For a wide parameter regime, under the two above constraints, all examined amplification measures were significantly larger in the anatomically constrained connectivity matrix relative to a comparable purely random E-I structure, implying that specific recurrent architecture found in the data contributes substantially to non-normal amplification on top of the division of neurons into E and I. (2) The thalamic input is particularly well-suited to exploit this capacity for amplification as it has a large overlap with the optimal input direction. Since the large scale architecture of cortical columns has similarities across species and cortical areas, it's likely that our findings hold for other cortical connectomes.

I-84. Discovering interpretable models of population dynamics from neural activity recordings

Mikhail Genkin
Tatiana Engel

MGENKIN@CSHL.EDU
ENGEL@CSHL.EDU

Cold Spring Harbor Laboratory

With advances in recording technologies, machine learning gains importance for translating neural activity data into theories about brain functions. Traditionally, data are fitted with ad hoc models representing a priori hypotheses. The best-fitting model is selected and interpreted as a biological mechanism, however none of the a priori hypotheses may be correct. Alternatively, data can be fitted with flexible models, which cover many hypotheses within a single model architecture (such as artificial neural networks). flexible models are usually optimized for predicting data (i.e. generalization) and then interpreted. However, it is unknown whether good data prediction guarantees accurate interpretation of the model. We develop a flexible and intrinsically interpretable framework for discovering neural population dynamics from data, and use this framework to test the connection between generalization and interpretation of flexible models. In our framework, population dynamics are governed by a non-linear dynamical system defined by a potential function. The activity of each neuron is related to the population dynamics through unique firing-rate functions, which account for heterogeneity of neural responses. The shape of the potential and firing-rate functions are simultaneously inferred from data via optimization over the space of continuous functions. The framework is flexible as it covers many hypotheses represented by non-linear dynamical systems. Also, the potential shape is intrinsically interpretable, e.g., the potential minima reveal attractors. The framework allows us to directly compare inferred models with the ground truth on synthetic data, thus testing the correctness of their interpretation. We find that many models discovered during optimization predict data equally well, yet they fail to match the correct hypothesis. We develop an alternative approach that identifies models with correct interpretation by comparing features discovered from independent data samples. We demonstrate the power of our approach by discovering metastable dynamics in spontaneous spiking activity in the primate area V4.

I-85. Heterogeneity of timescales in recurrent networks with clustered connectivity

Nicolae Istrate¹
Merav Stern²
Luca Mazzucato¹

ISTRATENICU@GMAIL.COM
MS4325@UW.EDU
LUCA.MAZZUCATO@GMAIL.COM

¹University of Oregon

²University of Washington

Recent experimental evidence suggests that neural circuits operate in a temporally irregular regime, where fluctuations in neural activity reveal the simultaneous presence of multiple timescales. Such heterogeneity of timescales, measured from single-cell autocorrelations, was observed both across different brain areas and, crucially, across neurons within the same circuit. This temporal variability has been confirmed during periods of ongoing activity, suggesting that it may be an intrinsic dynamical property of recurrent circuits. What are the neural mechanisms supporting heterogeneous distribution of timescales? Here, we investigate this question within the framework of recurrent rate networks with random connectivity. Random networks are known to exhibit a transition from fixed point to chaos, characterized by the emergence of an intrinsic dynamical timescale, which can be dialed from fast in the deeply chaotic regime, to slow near the transition to chaos. However, in the limit of large network size, all neurons in the network converge to the same timescale, inconsistent with the heterogeneity of timescales empirically observed across neurons within the same circuit. Here, we show that random networks with bistable units naturally exhibit large heterogeneity of timescales across neurons. Combining dynamical mean field theory with numerical simulations, we demonstrate that one can parametrically separate the timescales of different neuronal subpopulations within the same network, in the presence of a distribution of self-couplings. We provide a biophysical interpretation for the bistable firing rate units in terms of Hebbian assemblies, where self-couplings naturally represent potentiated synaptic couplings between neurons within the same assembly. We show that, in recurrent spiking networks where excitatory and inhibitory neurons are arranged in assemblies, neural activity generates heterogeneous distributions of timescales when assembly sizes are unequal. Our results establish a novel theoretical framework to investigate the heterogeneity of intrinsic neuronal timescales observed in behaving animals.

I-86. Premotor output to striatum carries action history information to support ongoing decision-making

Drew Schreiner
Christina Gremel

DCSCHREI@UCSD.EDU
CGREMEL@UCSD.EDU

University of California, San Diego

Decision-making is not done in isolated, discrete epochs but is continuous by nature and relies on historical information for adaptive control. Prior actions and associated consequences provide crucial feedback for the generation of self-initiated decision-making. Previous work investigating the role of history in decision-making often relies on constraining behavior via limiting choices (e.g., go/no-go, choose left/right), restraining animals, using cues to elicit choice, and instituting trials. However, decision-making is often non-binary, taking place in a continuous manner with one decision bleeding into the next. Animals must continually decide when and how (and how frequently and long) to act, often without explicit or consistent cues. To investigate how the brain performs continuous decision-making, we need to capture the integration of previous history into self-initiated behavior. We adapted a simple instrumental task; mice learn to hold down a lever for at least minimum duration to earn food. This task is self-initiated, self-paced, un-cued, and goal-directed. Furthermore, the duration of pressing serves as a continuous decision variable, allowing us to capture the influence of action history on ongoing decision-making. Corticostriatal circuits are hypothesized to underlie these ongoing processes, and rodent secondary motor cortex (M2; akin to primate pre-SMA/SMA) has particularly been implicated in associative integration necessary for action selection. However, it is unknown whether M2 utilizes action history to support continuous decision-making and if it sends this information to downstream striatal circuits. We measured calcium activity within M2 excitatory neurons and M2-striatal projection neurons during continuous decision-making. M2 calcium dynamics predicted ongoing decisions, and importantly, also reflected a history of prior actions. We find that lesion of these circuits directly impacts the use of history in decision-making. Thus, our findings suggest that continuous decision-making involves M2 processing and output to striatal circuits to govern the use of history for adaptive control.

I-87. Localized and sequential encoding of linguistic hierarchy in the human auditory cortex

Menoua Keshishian¹

Jose Herrero²

Nima Mesgarani¹

Bahar Khalighinejad¹

Hassan Akbari¹

Ashesh Mehta³

MK4011@COLUMBIA.EDU

JHERRERORU@NORTHWELL.EDU

NIMA@EE.COLUMBIA.EDU

BK2556@COLUMBIA.EDU

HA2436@COLUMBIA.EDU

AMEHTA@NSHS.EDU

¹Columbia University

²Feinstein Institute For Medical Research

³Feinstein Institute for Medical Research

Speech is generated by hierarchical combination of linguistic elements from phonemes, to syllables, words, and beyond. The processes of speech perception are also hypothesized to occur hierarchically in the brain, beginning with the mapping of a variable acoustic signal onto discrete sound categories (phonemes); combining these phonemes to form higher-level linguistic elements (according to phonotactic probabilities); ultimately giving rise to words which convey semantic information. The anatomical organization, temporal arrangement, and computational underpinnings of this processing, however, remain speculative. To determine where, when, and how different linguistic levels are processed, we measured the encoding of phonemic, phonotactic, and semantic information in neural responses recorded invasively from neurosurgical patients as they listened to natural speech. We leveraged hierarchically constructed linear regression models that predict the neural activity at each site from the acoustic and linguistic information present in the stimulus, to calculate the prediction power gained by adding increasingly higher-level features. We focused our analysis on three brain areas: Heschl's gyrus (HG), planum temporale (PT), and superior temporal gyrus (STG). Our results show phonemic information is represented across all three regions; phonotactic information is represented across PT and STG; and semantic information only shows up across STG. A further analysis of the timing of the effects of each feature showed increasing response delays going from phonemic to semantic information, together suggesting an anatomically localized and sequentially ordered emergence of phonemic, phonotactic, and semantic information that begins in the primary auditory cortex and extends to nonprimary auditory cortical areas. Furthermore, a nonlinear encoding model revealed increasingly more nonlinear stimulus-response relationship from primary to nonprimary auditory cortices and uncovered two distinct nonlinear neural transformations describing how linguistic features are extracted from the speech signal. These results provide direct evidence for a sequential and hierarchical organization of acoustic-to-linguistic mapping in the human auditory cortex.

I-88. Inferring neural population spiking rate from wide-field calcium imaging

Merav Stern

Eric Shea-Brown

Daniela Witten

University of Washington

MS4325@UW.EDU

ETSB@UW.EDU

DWITTEN@UW.EDU

Wide-field calcium imaging techniques allow recordings of high resolution neuronal activity across one or multiple brain regions. However, since the recordings capture light emission generated by the fluorescence of the calcium indicator, the neural activity that drives the calcium changes is masked by the calcium indicator dynamics. Here we develop and test novel methods to deconvolve the calcium traces and reveal the underlying neural spiking rate. Our methods take into account both the noise existent in the recordings and the temporal dynamics of the calcium indicator response. Our first method retrieves firing rates that are constant over discrete time bins. The size of each time bin depends on the data and is determined dynamically. Our second method retrieves the rate as a continuous function and is meant for studies that look for slow rate fluctuations rather than abrupt changes. We compare our results with those of two alternative deconvolution methods that were adjusted to fit wide-field

recordings deconvolution into spiking rate: direct deconvolution using a ‘first differences’ approach, and the ‘Lucy-Richardson’ image recovery method. We show that our methods outperform competitors on synthetic data as well as on recorded wide-field calcium for which parallel spiking recording exists.

I-89. Pattern completion and separation in the interplay between olfactory bulb and cortex

Gaia Tavoni
David Kersen
Vijay Balasubramanian
University of Pennsylvania

TAVONI@SAS.UPENN.EDU
ELDAVID3@GMAIL.COM
VIJAY@PHYSICS.UPENN.EDU

How does context change perception of sensory stimuli? In the olfactory system, experiments have shown that when animals make binary choices based on specific odor-context associations, cortical responses to similar odors are reshaped to favour odor generalization or odor acuity according to task demands. In particular, context can drive merging or separation of odor responses in the piriform cortex when odors are initially similar. By developing a model of the interplay between olfactory bulb and cortex that incorporates feedback from the central brain to represent context, we showed that these observed forms of cortical remodeling can be explained in terms of a statistical effect, partly similar to stochastic resonance. In particular, when odors share the same context, spatially random feedback inputs to the bulb push groups of mitral cells that are weakly tuned to either odor above the activity threshold required to transmit the sensory information to cortex. For realistic thresholds, such a change in the bulb activity amplifies the correlation in the cortical responses to the odors, in proportion to the initial odor similarity. In the opposite scenario in which odors are associated with alternative contexts, the model shows that spatially random, anticorrelated feedback inputs to the bulb produce pattern separation in the cortical responses, again in proportion to the initial odor similarity. Thus, the model provides the first theoretical interpretation of forms of pattern completion and separation that are observed in experiments, and links them to the measured response properties of neurons in the piriform cortex. The model also predicts that alternative properties (e.g. lower response thresholds) can produce opposite effects. This result suggests a way to further test this theory by manipulating neural excitability or by studying the consequences on olfactory perceptual learning of pathologies like Alzheimer’s disease that increase the excitability of piriform cortex neurons.

I-90. Nonlinear network activity regimes as a result of nonlinear neuronal dynamics

Pierre Ekelmans
Nataliya Kravnyukova
Tatjana Tchumatchenko
Max Planck Institute for Brain Research

PIERRE.EKELMANS@HOTMAIL.FR
NATALIYA.KRAYNYUKOVA@BRAIN.MPG.DE
TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

The firing rates of neurons that are part of a network depend on the strength of the recurrent and feedforward connections onto it, as well as the size of the network. The balanced state is a commonly used framework to study neural networks. It is exact in the limit of infinitely large networks, and is based on the assumption that excitatory and inhibitory inputs to a neuron exactly compensate each other. It is however unclear how accurate the balanced approximation is for finite networks. In our simulations of LIF networks, E-I balance is not always achieved for networks of biologically relevant size and is parameter dependent. Furthermore, many biological observations report a nonlinear relationship between the firing rate and the feedforward input, which cannot be achieved by balanced networks without synaptic plasticity. Here, we investigate how finite size effects can lead to significant nonlinearities, which can play a role in the computations performed by neural networks. An alternative rate model is the Stabilised Supralinear Network (SSN) which takes into account the neuronal transfer function. It

is based on the experimental observation that the neuronal input-output function can be approximated by a power law. Thanks to this built-in nonlinearity, the SSN is able to generate nonlinear network activity regimes without requiring synaptic plasticity. However, it is unclear whether SSN features can be observed in recurrent networks. Here, we show that spiking neural networks exhibit features predicted by the SSN. Specifically, we can recover regimes predicted by SSN in finite size spiking networks. The SSN can therefore be used to characterize the types of activity regimes which neural networks can exhibit, as well as the constraints on parameters to achieve them. This understanding of possible network dynamics and their requirements can facilitate the study of functional properties of neural circuits.

I-91. Mice detect increments of V1 spiking but not equivalent decrements

Jackson Cone
Morgan Bade
Nicolas Masse
Elizabeth Page
David Freedman
John Maunsell

The University of Chicago

JACKSON.J.CONE@GMAIL.COM
MORGAN5@UCHICAGO.EDU
NICOLAS.MASSE@GMAIL.COM
EAPAGE2@GMAIL.COM
DFREEDMAN@UCHICAGO.EDU
MAUNSELL@UCHICAGO.EDU

Changes in the retinal image evoke diverse responses in primary visual cortex (V1) neurons. How such changes in neuronal activity are decoded by other brain structures remains poorly understood. Although increments or decrements in neuronal activity can convey equivalent information, the emergence of ON and OFF channels early in the visual pathway suggests that spike rate increments might preferentially support visual perception. To explore this, we used optogenetic stimulation of V1 in trained, behaving mice to measure how increases and decreases in V1 spiking affect visual detection performance. Mice performed contrast change detection tasks while head-fixed. ChR2 was expressed in selected classes of neurons in the portion of the V1 retinotopic map that included the stimulus representation (Gabor). This enabled us to moderately decrease or increase visual responses in V1 and was confirmed with electrophysiology. Stimulation of excitatory neurons always enhanced detection of either contrast increments or decrements. Conversely, activation of inhibitory interneurons always impaired detection of either type of contrast change. We hypothesized that preferential decoding from spike rate increments may stem from the low basal firing rates typically observed in cerebral cortex, as this constrains the dynamic range available for decrements in firing to convey information. To examine this, we trained recurrent neural networks (RNNs) on the same task and implemented a spike cost. Costly spikes reduced neural activity in RNNs and forced them to preferentially weight spike rate increments. Motivated by these results, we trained mice to respond to ChR2-induced changes in V1 activity in the absence of visual stimuli. Sustained ChR2 input elevated V1 spike rates and provided a pedestal on which to present increments and decrements in spiking that were matched in absolute magnitude. Nevertheless, we found mice responded to optogenetically-induced increments in spiking but not to spiking decrements of comparable size.

I-92. Disambiguating memory from contrast in monkey inferotemporal cortex

Vahid Mehrpour¹
Travis Meyer¹
Eero Simoncelli^{2,3}
Nicole Rust¹

¹University of Pennsylvania

²New York University

³HHMI

V.MEHRPOUR@GMAIL.COM
TRMEYER@SAS.UPENN.EDU
EERO.SIMONCELLI@NYU.EDU
NRUST@PSYCH.UPENN.EDU

Humans are generally very good at remembering the images they have seen. Considerable evidence suggests that image memory is signaled by a reduction in response for familiar as compared to novel images ('repetition suppression') in high-level visual areas such as inferotemporal cortex (IT). However, IT neural responses are modulated not only by familiarity, but also by other factors such as contrast and attention, and it is unclear how the brain disambiguates visual memory signals from other types of firing rate modulations. To address this question, we analyzed behavioral and neural data collected from IT as two rhesus monkeys performed a single-exposure visual memory task in which they viewed images and indicated whether they were novel (never seen before) or familiar (seen exactly once before) while disregarding changes in image contrast across image repetitions. As expected, the monkeys were largely able to detect familiarity in the presence of contrast changes. In comparison, linear decoders applied to the early stimulus-evoked response (200 ms) confused memory and contrast and were inconsistent with the mapping from IT neural signals to behavioral reports. However, we found that a linear decoder applied to the end of the 500 ms viewing period successfully disambiguated memory from contrast and produced a much more accurate behavioral prediction. We also found that the temporal evolution of IT population responses from confusion to disambiguation could not be attributed to the simple disappearance of contrast modulation, but was well-described as a non-linear reformatting of this information across the viewing period.

I-93. VTA stimulation induces remapping in multiple hippocampal regions

Dorgham Khatib¹

Gilad Tocker²

Jonathan Gross³

Genela Morris⁴

Dori Derdikman¹

DKHATIB@CAMPUS.TECHNION.AC.IL

GILADTOCKER@GMAIL.COM

IJDM10@GMAIL.COM

GENELA.MORRIS@GMAIL.COM

DERDIK@TECHNION.AC.IL

¹Technion

²Northwestern University

³Rappaport Faculty of Medicine and Research Institute, Technion – Israel Institute of Technology

⁴Technion – Israel Institute of Technology

The hippocampus and associated structures contain spatially selective cell types such as place cells in CA1 and CA3, and grid cells in the medial entorhinal cortex (MEC). These are believed to be part of the brain's cognitive map of space. According to a theory proposed by Lisman & Grace in 2005, projections from the Ventral Tegmental Area (VTA) to the hippocampal formation are related to the signaling of novelty, and the formation of new memories. In the current study, we aim to understand how the VTA signals novelty at the level of the hippocampal network. One could use the phenomenon of remapping as an indication for novelty. In remapping place cells form a new map of space for different context. This leads us to the major question of this study, which is how the VTA conveys novelty to the cognitive map. We hypothesize that the VTA induces remapping, through which is related to signaling novelty. To validate this, TH-Cre mice were injected with AAV-hM3D(Gq) virus into the VTA and then implanted with a microdrive either in CA1 or the MEC. Each animal was trained in the same environment for 4-6 days until it was familiar with the it and had stable place cells. After activating the VTA in CA1 implanted mice, place cells in the 'CNO' condition remapped when compared to the 'Saline' recording that was done 10 minutes earlier, although the environment or the cues did not change (figure 1Bii, 1Cii), this effect was not seen in control mice (figure 1Bi, 1Ci). In the MEC, Grid cells expanded in scale after the activation of the VTA (figure 2B, 2C), which could have led to the remapping in CA1. These results indicate that the VTA might indeed cue novelty to the hippocampal formation and evoke encoding of a new space.

I-94. Sex-modulated norepinephrine function in mediating exploration-exploitation tradeoff

Cathy Chen¹
Becket Ebitz¹
Sylvia Bindas²
Evan Knep¹
Nicola Grissom¹

CHEN5388@UMN.EDU
REBITZ@GMAIL.COM
BINDAS@OHSU.EDU
KNEPX001@UMN.EDU
NGRISSOM@UMN.EDU

¹University of Minnesota

²Oregon Health and Science University

In an uncertain world, we balance two goals: exploiting rewarding options when they are available, and exploring potential better alternatives. Often, there are a range of equivalently good individual solutions to the explore-exploit dilemma, but we know little about the neural bases of this interindividual variability. One major axis of interindividual variability in executive functions and decision-making is sex. Sex is a profound modulator of the locus coeruleus-norepinephrine (LC-NE) system that has long been thought to underpin exploration and learning. However, it remains unclear how the balance between exploration and exploitation differs across sexes, much less the role of LC-NE in these sex differences. Here, we compared male and female mice in a restless two-arm bandit task, which encourages both exploration and exploitation. The reward probabilities of two arms changed independently and stochastically over trials so that the animals could only infer values through sampling and integrating reward and choice history. We found that that females were more likely to repeat behaviors that produced reward, indicating that they learned more from wins than males. fitting reinforcement learning models revealed that the learning rate of females was higher than males overall and, unexpectedly, increased over time. This suggests that that females, but not males, “learned to learn”. To determine whether increased learning in the females was due to increased LC-NE tone, we systemically administered propranolol (5mg/kg), a non-selective beta-adrenergic receptor antagonist. Surprisingly, this did not reduce the learning rate in females. Instead, blocking norepinephrine increased the stickiness of choice behaviors and reduced the level of decision noise. Together, these results suggest that males and females solve the explore/exploit dilemma differently and suggest a novel role for the LC-NE system in regulating exploration via regulating behavioral stickiness, rather than the rate of learning.

I-95. Oxytocin enables social transmission of maternal behavior

Ioana Carcea¹
Robert Froemke²

IOANA.CARCEA@RUTGERS.EDU
ROBERT.FROEMKE@MED.NYU.EDU

¹Rutgers University

²New York University

Maternal care is profoundly important for mammalian survival, and maternal behaviors can also be expressed by non-biological parents after experience with infants. One critical molecular signal for maternal behavior is oxytocin, released by hypothalamic paraventricular nucleus (PVN) and enabling plasticity within auditory cortex for recognizing infant vocalizations. To determine how these changes occur during natural experience, we continuously monitored homecage behavior of female virgin mice co-housed for days with an experienced mother and litter, synchronized with in vivo recordings from virgin PVN cells, in particular from oxytocin neurons. Mothers engaged virgins in maternal care by ensuring that virgins were in the nest, and demonstrated maternal behavior in spontaneous pup retrieval episodes. These social interactions activated virgin PVN and gated behaviorally-relevant cortical plasticity for pup distress calls. Thus rodent maternal behavior can be learned by social transmission, and our results describe a mechanism for adapting the brains of new parents to infant needs via endogenous oxytocin.

I-96. Solutions to the assignment problem balance tradeoffs between local and catastrophic errors

W Jeffrey Johnston
David Freedman
The University of Chicago

WJEFFREYJOHNSTON@GMAIL.COM
DFREEDMAN@UCHICAGO.EDU

An observer watching a barking dog and purring cat together in a field has distinct pairs of representations of the two animals in their visual and auditory systems. Without prior knowledge, how does the observer infer that the dog barks and the cat purrs? This binding of disparate representations is called the assignment problem, and it must be solved to integrate distinct representations across but also within sensory modalities. Here, we identify and analyze a solution to the assignment problem: the representation of one or more common stimulus features in pairs of relevant brain regions – for example, estimates of the spatial position of both the cat and the dog represented in both the visual and auditory systems. Using a statistical model, we derive the assignment error rate as a function of the number of stimuli, the number of features that define each stimulus, the number of common features, and the mean squared-error, or distortion, of each stimulus feature when decoded. Using rate-distortion theory, we derive the number of bits of information necessary to achieve a particular assignment error rate and level of distortion. The model reveals two tradeoffs for a fixed number of bits: first, increasing the number of common features decreases the assignment error rate but increases distortion – that is, makes the estimate of each stimulus noisier; second, representing the common features with different distortion across the two brain regions decreases distortion across all features but increases the assignment error rate. These tradeoffs balance redundancy, which reduces the assignment error rate, and efficiency, which reduces distortion. Assignment errors reported in humans are broadly consistent with our model; we also make further, specific predictions that are yet to be tested. Overall, this work provides key insight into an important computational problem that arises from distributed neural representations.

I-97. Non-linear inhibitory plasticity can homeostatically regulate excitatory weight dynamics

Christoph Miehl
Julijana Gjorgjieva
Max Planck Institute for Brain Research

CHRISTOPH.MIEHL@GMAIL.COM
GJORGJIEVA@BRAIN.MPG.DE

Synaptic changes underlying learning and memory are believed to be implemented by Hebbian plasticity based on pre- and postsynaptic activity. But this mechanism on its own is unstable, leading to runaway weight dynamics in networks without an additional fast homeostatic mechanism. We propose a novel experimentally-inspired non-linear inhibitory plasticity mechanism which can homeostatically regulate excitatory weight dynamics. We study this rule in both a rate- and a spike-based plastic network. first, in the rate-based case, we show that non-linear inhibitory plasticity is sufficient to counteract excitatory runaway weight dynamics, to form receptive fields in excitatory neurons and to generate bidirectional connections between neurons with similar receptive fields. No additional homeostatic mechanisms, no upper bounds and no fine-tuning of parameters are required. We also show that our inhibitory plasticity rule can be mapped to the sliding threshold of the established Bienenstock-Cooper-Munro (BCM) rule, which has successfully been used to model e.g. orientation selectivity. Second, we show why in the spike-based case, besides non-linear inhibitory plasticity, an additional homeostatic mechanism is needed to counteract excitatory runaway weights. Most studies assume some form of global homeostasis, i.e. a modification of all synaptic weights onto a neuron. However, such global mechanisms lead to hyperselectivity where a single synapse wins even within input patterns. Motivated by recent experimental findings on local forms of homeostasis, we introduce a synapse-specific metaplastic mechanism, where the induction threshold of long-term depression and potentiation at each excitatory synapse changes dynamically depending on the recent induction of plasticity. We show that this form of local homeostasis successfully stabilizes excitatory weight dynamics and allows for emergence of selectivity in spiking networks.

I-98. Hierarchical inference guides perceptual decision-making in a dynamic environment

Julie Charlton¹
 Wiktor Mlynarski
 Yoon Bai¹
 Ann Hermundstad²
 Robbe Goris¹

JCHARLTON@UTEXAS.EDU
 WIKTOR.MLYNARSKI@IST.AC.AT
 YOONBAI@UTEXAS.EDU
 HERMUNDSTADA@JANELIA.HHMI.ORG
 ROBBE.GORIS@UTEXAS.EDU

¹University of Texas at Austin

²HHMI Janelia Research Campus

Organisms perform perceptual tasks across a wide variety of contexts. This necessitates sensory coding strategies that seek to jointly interpret incoming sensory signals and keep track of contextual changes in the environment. Here, we investigate how humans integrate sensory and contextual information while making perceptual decisions in a dynamic environment. We first developed a Bayesian observer model that implements the optimal strategy for interpreting noisy sensory stimuli under contexts that specify a dynamically-switching stimulus distribution. On each trial, the model observer performs two stages of inference. It first estimates the most likely task context, and then uses this information to optimally decode the stimulus identity on the basis of noisy sensory information. Consistent with Bayesian inference, the stimulus estimate depends on the continually-evolving belief about task context. Immediately following a context-switch, the model observer's uncertainty about the context is maximal, and thus the influence of contextual information on the perceptual report is minimal. When the context is stable across trials, contextual uncertainty gradually diminishes. As a consequence, the prior becomes stronger and exerts greater influence on the perceptual decision. To test whether humans exhibit similar dynamics, we measured perceptual judgments of six naive observers in a dynamic orientation discrimination task. Prior to stimulus onset, task-context was communicated to the observer via a visual cue. Subjects were not told what this cue signified. Their choices depended on stimulus orientation and task context. We used a generalized linear model to characterize the cross-trial dynamics of the decision-making process. We found that contextual influence was minimal upon a context-switch and increased with context stability, closely resembling the ideal observer. We conclude that humans rely on a hierarchical inference strategy and make use of contextual uncertainty when performing perceptual tasks in dynamic environments.

I-99. Unsupervised depth estimation via visual-motor integration in recurrent neural networks

Elliott Abe
 Philip Parker
 Cristopher Niell
 Yashar Ahmadian
 University of Oregon

ELLIOTTABE@GMAIL.COM
 PPARKER@UOREGON.EDU
 CNIELL@UOREGON.EDU
 YASHAR@UOREGON.EDU

Vision allows the brain to form internal models of distant surroundings. These representations are traditionally thought to arise from bottom-up processing of retinal input. However, naturally behaving animals move through their environment, and self-motion information is sent, via efference copies, into visual cortex where it can be combined with visual input to calculate objective information about the surroundings. To probe such computations, and inspired by ongoing experiments in mouse vision, we focused on monocular depth estimation through motion parallax, in which animals combine self-motion and visual signals to calculate absolute distances to environmental objects. We simulated a camera-agent performing random walks in a 3D environment, and fed the recorded 2D camera frames to a multi-layer recurrent neural network (RNN), which was trained to predict future frames. Accurate prediction of future sensory inputs requires learning an internal model for the dynamics of the agent and environment. These dynamics are best modeled in terms of natural dynamical variables (such as 3D coordinates of objects), which may not be explicit in the sensory inputs (such as 2D retinotopic inputs), but can be extracted

from them. We therefore hypothesized that unsupervised predictive learning results in an explicit representation of depth, and that providing self-motion signals to the RNN will further improve this representation. To test our hypothesis, we trained linear readouts from neural activations in different layers of the trained RNN to match ground-truth depth maps. We found that the RNN does indeed form depth representations, without being explicitly tasked to do so. Moreover, the depth representation becomes more explicit and accurate in deeper layers of the network, an effect that is amplified when the RNN also receives self-motion input. These results suggest that internal representations of depth can arise from learning to predict future monocular visual inputs by integrating visual and self-motion signals.

I-100. Signatures of criticality observed in efficient coding networks

Anna Levina¹
Shervin Safavi²
Nikos Logothetis²
Matthew Chalk³

ANNA.LEVINA@UNI-TUEBINGEN.DE
RESEARCH@SHERVINSAFAVI.ORG
NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE
MATTHEW.CHALK@INSERM.FR

¹University of Tübingen

²Max Planck Institute for Biological Cybernetics

³Sorbonne Université

Over the last decades, multiple studies have reported signatures of criticality observed in various neuronal recordings. Moreover, theoretical investigations demonstrate that multiple aspects of information processing are optimized at the second-order phase transition. These studies motivated the hypothesis that the brain operates close to a critical state. To evaluate how distance from criticality may influence neural computations, researchers typically considered neural models that can attain various states (including critical and non-critical) depending on control parameters (e.g. connection-strength) and quantified how general information processing capabilities such as sensitivity to input, dynamic range, or information-transmission depend on these parameters. Certainly, being in a state with such optimized capabilities are relevant for the computations in the brain, but they are too abstract to provide a concrete implementation. For instance, all the mentioned capabilities are relevant for coding sensory information, but mere adjusting for the closeness to criticality cannot provide a neural implementation for coding given resource constraints. Whereas, frameworks like efficient coding both provide the objective to maximize and the implementation. Therefore, we introduce a novel complementary approach. We study a network that implements efficient coding and we investigate the presence of the scale-free neuronal avalanches in an optimized network. We consider a network of LIF neurons with synaptic transmission delays whose connectivity and dynamics are optimized for efficient coding. Previously, it was shown that the performance of such networks varies non-monotonically with the noise amplitude. We consider networks with different noise amplitudes and evaluate how close they are to a critical state by measuring deviations from the nearest power-law of avalanche size distributions. Interestingly, only in the optimized network the distribution of avalanche sizes truly follow a power-law. This result has important implications, as it shows how two influential, and previously disparate fields - efficient coding, and criticality - might be intimately related.

I-101. Learning structure from the environment using a recurrent spiking network

Amadeus Maes
Mauricio Barahona
Claudia Clopath
Imperial College London

AMADEUS.MAES@GMAIL.COM
M.BARAHONA@IMPERIAL.AC.UK
C.CLOPATH@IMPERIAL.AC.UK

The brain is able to adapt to the statistics of the environment. Information about the structure of the world

around us is extracted and informs future behaviours and decision-making. Long-term plasticity is thought to underlie the formation of internal models that match the external world. However, it is unclear how the brain learns representations as most of the current approaches use non biologically-plausible ways to set or "train" the synaptic weights. For example, FORCE training uses a non-local optimization to learn a temporal target sequence. Back-propagation through time can be used to learn a target distribution. Here, we study a general model architecture, where read-out neurons are connected to a recurrent network, and use Hebbian learning in two different tasks. first, we investigate learning temporal spiking sequences. To this end, we train the recurrent network to generate a sequential dynamics. The read-out synapses then learn the mapping of this sequential dynamics to the target signal. All synapses are plastic under typical Hebbian learning rules. We demonstrate that the model is able to learn temporal spiking patterns and is robust to synaptic failure. Second, we study learning a probability distribution. The read-out neurons now encode discrete states and are trained to spike proportionally to a target probability. We assume here that learning is a product of a presynaptic trace and an error term, computed at the postsynaptic neuron. After learning, we show that the model generates fast samples from the probability distribution during spontaneous dynamics. Importantly, the two tasks demonstrate a clear link between network connectivity, spiking dynamics and learning on behavioural time scales.

I-102. Structure-function relationships in neural networks from geometric features of error landscapes

Tirthabir Biswas
James fitzgerald

TIRTHABIR@GMAIL.COM
FITZGERALDJ@HHMI.ORG

HHMI Janelia Research Campus

Neural computation in biological and artificial networks relies on nonlinear synaptic integration. The synaptic connectivity matrix of weights between neurons in these networks is a critical determinant of overall network function. However, what patterns of connectivity are specifically required to generate measurable network-level computations are largely unknown. Here we introduce a geometric framework for calculating and analyzing connectivity constraints in recurrent neural networks of rectified linear units. We focus on steady-state functional responses to a finite set of input conditions, each corresponding to the neural activity pattern that provides feedforward drive to the network. Our formalism allows us to analytically calculate all feedforward and recurrent connectivity matrices that can generate the measured response patterns. It thus allows us to comprehensively analyze rigorous links between neural network structure and function. We illustrate this point by deriving conditions for when synapse signs are certain across the ensemble of model solutions.

I-103. Complementary encoding pathways build a memory hierarchy in a model of hippocampus

Louis Kang¹
Taro Toyozumi²

LOUIS.KANG@BERKELEY.EDU
TARO.TOYOIZUMI@RIKEN.JP

¹University of California, Berkeley
²RIKEN Center for Brain Science

The hippocampus is crucial for our ability to remember personally experienced events. Its central structure, CA3, is believed to perform pattern completion, which is the recall of stored memories from incomplete cues. While many models of pattern completion have been developed and analyzed, they have not implemented a key feature of the hippocampal architecture that determines how information converges at CA3. When information enters the hippocampus, it is split into two pathways. One of them leads directly to CA3, whereas the other projects to dentate gyrus (DG) before arriving at CA3. Experiments suggest that the two pathways encode information in different ways. Activity through DG is sparser and less correlated, which is thought to accentuate differences

between similar experiences. Yet, the computational capabilities of these two pathways have not been firmly established by experiments or simulations.

We demonstrate with a Hopfield-like model how complementary encoding pathways can enable CA3 to perform unsupervised categorization through the accumulation of correlated patterns. By adjusting the threshold for neural activation, CA3 can recall either an individual pattern or the category to which it belongs. This capability is made possible by complementary encodings of the same information. The sparse, uncorrelated pathway through DG prevents the mixing of different patterns, even if they originally contained overlaps. Meanwhile, the dense, correlated pathway promotes the network to recover the features common to many overlapping patterns, thus constructing categories. By interchanging between dense and sparse encodings, our model can traverse a memory hierarchy of categories and individual examples. It explains how representing information at different resolutions, which is considered a key feature of memory, can be implemented by neural circuits in the hippocampus.

I-104. A GLM approach for characterizing orbitofrontal cortical responses to multiple reward attributes

David Hocker
Cristina Savin
Christine Constantinople
New York University

DH148@NYU.EDU
CS5360@NYU.EDU
CMC472@NYU.EDU

The orbitofrontal cortex (OFC) is a part of the prefrontal cortex that is thought to be critical for goal-directed behaviors and value-based decision-making. Studies of value-based decision-making in animal models often use sensory cues (e.g., static images, odors) to convey information about one or more reward attributes simultaneously (e.g., identity, amount, and/or probability). Here, we trained rats on a task in which distinct sensory cues occurring independently and variably in time conveyed information about dissociable reward attributes (reward probability and amount). We exploited the temporal variability of task events to fit a generalized linear model (GLM) to OFC firing rates, to generate a rich description of the encoding of various task variables, including distinct reward attributes, choice, and trial history, over time in individual neurons. This approach, which has successfully been applied in sensory, motor, and premotor areas during perceptual decision-making, provides a powerful platform for characterizing neural coding in areas subserving value-based decision-making, including the OFC.

I-105. The liminal state: Neural signatures of perceptual awareness

Meenakshi Asokan¹
Kenneth Hancock²
Ross Williamson²
Daniel Polley¹

MASOKAN@G.HARVARD.EDU
KEN_HANCOCK@MEEI.HARVARD.EDU
ROSS.S.WILLIAMSON@GMAIL.COM
DANIEL_POLLEY@MEEI.HARVARD.EDU

¹Harvard University

²Massachusetts Eye and Ear Infirmary

Perception involves using the sensory information available to us to form a meaningful representation of the world around us. In most real world scenarios, that involves detecting and tracking the signals of interest amidst potentially distracting stimuli, like hearing out one speaker in a crowded cocktail party. While even the state-of-the-art speech recognition software do not perform well in such scenarios, how does our brain solve this problem? Here, we explore this fundamental question by examining neural circuits and codes that underlie auditory perceptual awareness in a crowded scene. We train mice to perform a ‘real world’ auditory challenge where they get rewarded when they report hearing out a stream of regularly repeating target tones embedded in an ongoing

stochastic multi-tone distractor. We perform a systematic electrophysiological survey of neural activity by simultaneously recording from hundreds of neurons in widely distributed brain regions spanning the auditory neural axis from midbrain to thalamus to cortex, while mice are engaged in the task. The results reveal a graded transformation of the neural code along the neural axis with the auditory midbrain adhering faithfully to the physical stimulus irrespective of the animal's perceptual state. The perceptual awareness is reflected most prominently in the auditory cortex and to a lesser extent in the auditory thalamus, as an enhanced representation of the target tone in the liminal state of perception. We further test our hypothesis that the cortico-thalamic feedback projections play a key role in propagating this perceptual information from cortical to subcortical regions. In summary, we have developed a framework to address how the neural codes are reformatted along the neural axis to reflect perceptual awareness. These findings contribute to the emerging view on how perceptual representations are distilled out of raw sensory representations.

I-106. Using task-optimized neural networks to understand why brains have specialized processing for faces

Katharina Dobs¹

Alex Kell²

Julio Martinez¹

Michael Cohen¹

Nancy Kanwisher¹

KDOBS@MIT.EDU

ALEX.KELL@COLUMBIA.EDU

JULIOM@MIT.EDU

MICHAELTHECOHEN@GMAIL.COM

NGK@MIT.EDU

¹Massachusetts Institute of Technology

²Columbia University

Previous research has identified multiple functionally specialized regions of the human visual cortex and has started to characterize the precise function of these regions. But why do brains have functional specialization in the first place, and why do they have the particular specializations they do (e.g., for faces or scenes, but apparently not for food or cars)? Here, we address these questions using the well-studied case of face selectivity. We used deep convolutional neural networks (CNNs) to test the hypothesis that face-specific regions are segregated from object cortex in the primate visual system because the computations and feature spaces supporting face and object perception differ from each other. We trained two separate CNNs with the AlexNet architecture to categorize either faces or objects. The face-trained CNN performed worse on object categorization than the object-trained CNN and vice versa, demonstrating that the features optimized for each task differ from one another. To determine whether a CNN could learn to share features across these tasks, we trained dual-task CNNs with a branched architecture, varying the number of layers that were shared between tasks (Kell et al., 2018). The fully-shared network performed worse than the separate CNNs indicating a cost for sharing both tasks in one system. However, a network that shared early processing stages, like the primate visual system, exhibited unimpaired performance. Do these results generalize to architectures with larger capacity? We trained three networks with a VGG16 architecture: one on faces, one on objects, and one on both. In contrast to the AlexNet results, the dual-task CNN performed as well as the separate networks. However, lesion experiments showed that segregation of face and object processing emerged spontaneously in the network optimized for both face and object recognition, suggesting a possible reason for the specialization of face processing in the human brain.

I-107. A model for oscillatory gating of information flow between neural circuits as a function of local recurrence

Mikail Khona
Qianli Xu
Ila Fiete

MIKAIL@MIT.EDU
QLXU@MIT.EDU
FIETE@MIT.EDU

Massachusetts Institute of Technology

Physical connections between neurons are hardwired on the timescale of seconds, but information and signals must be routed flexibly for computation on the same timescales. At the same time, species from insects to mammals exhibit circuit-level oscillations with systematic phase relationships between spikes and the regional oscillation. These observations have led to the hypothesis that oscillations regulate communication between areas, but our mechanistic understanding of how phase and frequency effects could affect inter-areal communication is weak. We present a simple dynamical network model of communication between circuits: A receiver network with tunable recurrent weight strengths that exhibits attractor dynamics in the strong weight regime and receives spatially untuned oscillatory input, together with direct input from one or more transmitter networks with oscillating, spatially tuned outputs. We show that for same-frequency oscillations, there is an optimal phase for information transfer that depends on the strength of local recurrence in the receiver network: While a weakly recurrent receiver responds with a phase difference of 0, in a network with strong recurrent weights, external inputs are drowned out by local contributions when cells are maximally depolarized, and a phase offset of 90 degrees is optimal for communication. With multiple transmitters and a strongly recurrent receiver, only the transmitter with closest phase to optimal drives the receiver. Our results generalize to transmitters and receivers with different input frequencies, to inputs with asynchronous modulation and to receivers with discrete as well as continuous fixed points. These results can provide a mechanistic basis for understanding experimental results from the hippocampus, where the relative influence of presynaptic inputs depends on oscillatory coupling, and from the visual cortex, where oscillatory phase relationships determine the influence of one subregion on another and make predictions about the relative phases and frequencies of oscillations between brain areas engaged in information exchange.

I-108. Unraveling the neural circuitry of predictive coding in mouse visual cortex

Hannah Choi¹
Marina Garrett²
Rylan Larsen²
Nicholas Cain^{2,3}

HANNAHCH@UW.EDU
MARINAG@ALLENINSTITUTE.ORG
RYLANL@ALLENINSTITUTE.ORG
NICAIN.SEATTLE@GMAIL.COM

¹University of Washington

²Allen Institute for Brain Science

³Google AI

Predictive coding has been proposed as a theory of hierarchical cortical computation to create efficient neural codes. It postulates that cortical circuits learn and represent information about the world through hierarchical interactions across cortical areas and layers. The internal representation of the world is actively compared and updated via feedback predictions about neural activities of the lower cortical areas, and feedforward signals representing the residual errors between the feedback predictions and the encoding expectation. Hierarchical predictive coding models have been widely used to explain a variety of neural responses in sensory systems, including center-surround antagonism in the retina (Srinivasan et al., 1982) and end-stopping effects in V1 (Rao & Ballard, 1999). Yet, how algorithmic nodes of predictive coding models connect to detailed cortical circuits has been highly speculative (Keller & Mrsic-flogel, 2018). In this work, we directly test key assumptions of predictive coding in mouse visual cortex by leveraging in vivo 2-photon imaging. We measured neural activity in response to expected and unexpected sequences of natural images across three hierarchically-related areas in mouse visual cortex: the primary visual cortex (VISp), the posterior medial higher order visual area (VISpm), and retrosplenial

cortex (RSP). These measurements were made across a range of depths, enabling us to identify layer specificity of neuronal signals. We used two different types of expectation violation stimuli: sequence violation with oddball images and partial occlusion. Our analysis shows that prediction error signals are present in all three areas across the range of cortical layers. However, the error signals are not indistinguishably encoded in the neuronal populations. We found that: 1) the sensitivity of spatial prediction errors changes across regions and layers, and 2) error signals are image-specific. Our study thus provides us experimental insights into underlying neural mechanisms of predictive coding in mammalian visual cortex.

I-109. Action context influences a sensorimotor transformation across cortical areas

Yi-Ting Chang¹
Eric A Finkle¹
Jae Hun Lee¹
Emily E Lubin²
Daniel O'Connor¹

YTCHANG@JHMI.EDU
EFINKEL88@GMAIL.COM
JLEE527@JHU.EDU
EMILYLUBIN29@GMAIL.COM
DAN.OCONNOR@JHMI.EDU

¹Johns Hopkins University

²University of Pennsylvania

The brain receives sensory inputs from different modalities, but only a subset of them triggers action. In other words, the relevance of a sensory stimulus to current action goals (action context) shapes its representations and affects how it is transformed to behavioral outputs. It remains unclear how action context modulates a sensorimotor transformation. Here, we developed a cross-modal selection task for head-fixed mice. In one context, mice made a specific movement in response to tactile stimuli while ignoring visual stimuli but then switched to the other context where they made a separate movement to visual stimuli while ignoring tactile stimuli. First, we characterized the effects of action context on processing of tactile stimuli. We found that tactile-evoked activity was enhanced and sustained in both somatosensory and motor cortex when a tactile stimulus was behaviorally relevant, compared to when it was not. Moreover, the context-dependent change was stronger in motor cortex than in somatosensory cortex. This result reveals that tactile inputs were processed in sensory and motor areas in a context-dependent manner. Surprisingly, action context influenced not only sensory-evoked activity but also baseline activity. A higher proportion of motor cortex single units showed significant baseline difference between action contexts compared with somatosensory cortex. This suggests that the sensorimotor circuits related to current action context are primed to process relevant sensory inputs and motor networks can be important for context-dependent gating. Finally, to address the mechanisms underlying the context-dependent changes, we have begun to optogenetically inhibit frontal and parietal areas. The preliminary results suggest that parietal cortex may contribute to setting up modulation of neural activity by context. Overall, our findings show tactile-to-behavior processing depends on action contexts and adjusting network states may be the key to achieve sensorimotor flexibility.

I-110. Transient working memory is controlled by slow synaptic dynamics

Sophia Becker
Tatjana Tchumatchenko
Andreas Nold

SOPHIA.BECKER@BRAIN.MPG.DE
TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE
ANDREAS.NOLD@BRAIN.MPG.DE

Max Planck Institute for Brain Research

Working memory plays a crucial role in the perception and processing of sensory stimuli in both animals and humans. Several psychiatric diseases such as schizophrenia and autism are characterized by deficits in working memory. Nevertheless, it is still unclear how working memory is implemented in the brain and how its physiological

function is impaired in disease. We want to shed light on these questions by investigating the effect of slow synaptic modulation on the synaptic working memory model presented by Mongillo et al. (Science, 2008).

The multitude of activity patterns observed in experimental working memory studies has fueled a discussion about whether short-term memory is represented by transient or persistent activity during the delay period (Lundqvist et al., J Neurosci, 2018; Constantinidis et al., J Neurosci, 2018). While other modeling studies implement working memory as quiescent synaptic traces or persistent population spiking, we find that networks with short-term plastic (STP) synapses can additionally produce transient working memory activity. This not only reconciles different experimental observations – it also touches on the question which role forgetting, i.e. the fading of memories, could have for short-term memory.

To understand how persistent, transient and quiescent working memory representations are modulated by slow synaptic timescales, we integrate an experimentally observed, STP-modulating signaling mechanism based on the postsynaptic molecule PRG1 into the synapses of our network. While persistent and quiescent network regimes are nearly unaffected by PRG1, the duration of transient working memory activity can be controlled by PRG1 signaling. This is consistent with the links between the breakdown of PRG1 signaling, STP impairments and schizophrenia.

To conclude, we show that transient working memory representations exist in the synaptic working memory framework (Mongillo et al., Science, 2008). Its duration can be controlled by parameters of the PRG1 signaling mechanism.

I-111. fine-scale organization of the functional diversity in mouse superior colliculus

Hymavathy Balasubramanian¹
Jeremie Sibille²
Jens Kremkow²

HYMAVATHY.BALASUBRAMANIAN@BCCN-BERLIN.DE
JEREMIE.SIBILLE@CHARITE.DE
JENS.KREMKOW@CHARITE.DE

¹Bernstein Centre for Computational Neuroscience
²Charite-Universitätsmedizin Berlin

The variety of retinal ganglion cells (RGC) in the mouse eye suggests that vision is encoded with up to 30 different parallel streams of information (Baden et al. 2016). The mouse superior colliculus (SC) receives up to 80% of RGC's axons, while only ~30% of them project to the lateral geniculate nucleus (LGN) (Ellis et al. 2016). Recently it was shown that the functional diversity in the LGN can be explained by a combination of few RGC response types (Roson et al. 2019). However, the extent of retinal convergence leading to the corresponding diversity in the SC remains under-examined. As a first step to address this question, we performed extracellular recordings in mouse SC using Neuropixel probes (Jun et al. 2017). The probe was tangentially inserted into the visual layers of the SC thus maximizing the spatial resolution resulting in several hundred responsive channels. To characterize the functional diversity we presented a full-field chirp stimulus and identified classes of responses using unsupervised clustering (Non-Negative Matrix Factorization (NNMF) - Roson et al.2019). The majority of the obtained classes exhibit clear features of ON, OFF and ON-OFF responses with both transient and sustained firing. In these classes, we observed a higher representation of OFF-transient and ON-OFF-transient responses. The spatial resolution of our recordings further revealed defined regions exhibiting response of the same class suggesting a fine-scale organization of the functional diversity in the SC both within and across SC layers. We further aim to model the SC responses as a function of the identified retinal input classes (Baden et al.2016, Roson et al.2019) to determine the combination of RGC types leading to the diversity observed in the SC.

I-112. Hippocampal population representations in reinforced continual learning

Samia Mohinta¹
Stephane Cioocchi²
Rui Ponte Costa¹

MOHINTA1234@GMAIL.COM
CIOCCHI@PYL.UNIBE.CH
RUI.COSTA@BRISTOL.AC.UK

¹University of Bristol

²University of Bern

The hippocampus has long been associated with spatial memory and goal-directed spatial navigation. However, how hippocampal neurons jointly encode spatial and behaviourally relevant features as animals learn multiple tasks is largely unclear. Here, we analyse population-level activity of hippocampal CA1 neurons as animals learn to continually adapt to different spatial navigation strategies. Demixed Principal Component Analysis (dPCA) is applied on neuronal recordings from 612 hippocampal CA1 neurons of rodents learning to perform allocentric and egocentric spatial tasks. The components uncovered using dPCA on the firing activity reveal that hippocampal neurons encode relevant task variables such as decisions, navigational strategies and reward location. In particular, we found components that (i) separate correct and incorrect trial outcomes, (ii) are akin to an eligibility trace of a directional choice, (iii) dissociate allocentric and egocentric navigational rules and (iv) have traces of reward prediction errors. Next, we compare our experimental observations with standard reinforcement learning algorithms, highlighting similarities and differences. Deep reinforcement learning model achieves similar average performance when compared to animal learning, but fails to capture behaviour during task switching. Overall, our results give insights into how the hippocampus is involved in reward and multi-task encoding, and put forward a framework to explicitly compare biological and machine learning during spatial continual learning.

I-113. Where can a place cell put its fields? Let us count the ways

Thibaud TAILLEFUMIER¹
Man Yi YIM¹
Lorenzo SADUN¹
Ila fiete²

TTAILLEF@AUSTIN.UTEXAS.EDU
MANYI.YIM@UTEXAS.EDU
SADUN@MATH.UTEXAS.EDU
FIETE@MIT.EDU

¹University of Texas at Austin

²Massachusetts Institute of Technology

A place cell in the hippocampus can exhibit multiple firing fields within and across environments. What determines the configurations of these fields? Could they be set down in arbitrary locations to form arbitrary maps? To answer these questions, we conceptualize place cells as performing high-level evidence combination with various inputs, thus formulating them as perceptrons. We characterize and count which place-field arrangements are realizable in the interiors of environments that lack external sensory cues based on multiscale periodic grid-cell drive. All possible field arrangements are realizable over environments achieving the “separating capacity”, set by the rank of the input matrix. This rank scales as the sum of distinct periods, which is tiny relative to the coding range of the input, which scales as the product of periods. Beyond the separating capacity, not all field arrangements are realizable and the fraction of realizable field arrangements goes down as the size of the environment increases. Over the input coding range, the realizable arrangements make up a vanishing fraction of potential arrangements. Compared to random inputs, grid-cell-structured inputs permit a modestly smaller fraction of realizable arrangements but substantially larger separating margins, which confer stability to selected place-field arrangements. To conclude, grid-cell-driven place cells can only exhibit highly restricted field arrangements in large environments.

I-114. Deep neuroethology of a virtual rodent

Josh Merel¹
 Diego Aldarondo²
 Jesse Marshall³
 Yuval Tassa¹
 Greg Wayne¹
 Bence Olveczky³

JSMEREL@GOOGLE.COM
 DIEGOALDARONDO@G.HARVARD.EDU
 JESSE_D_MARSHALL@FAS.HARVARD.EDU
 TASSA@GOOGLE.COM
 GREGWAYNE@GOOGLE.COM
 OLVECZKY@FAS.HARVARD.EDU

¹DeepMind

²Harvard

³Harvard University

Parallel developments in neuroscience and deep learning have led to mutually productive exchanges, pushing our understanding of real and artificial neural networks in sensory and cognitive systems. However, this interaction between fields is less developed in the study of motor control. In this work, we develop a virtual rodent as a platform for the grounded study of motor activity in artificial models of embodied control. We then use this platform to study how motor activity differs across contexts by training a model to solve four complex tasks. Using methods familiar to neuroscientists, we describe the behavioral representations and algorithms employed by different layers of the network using a neuroethological approach to characterize motor activity relative to the rodent's behavior and goals. We find that the model uses two classes of behavioral representations which respectively encode the task-specific behavioral strategies and task-invariant behavioral kinematics. These representations are reflected in the sequential activity and population dynamics of neural subpopulations. Overall, the virtual rodent facilitates grounded collaborations between deep reinforcement learning and motor neuroscience.

I-115. Emergence of functional and structural properties of head direction system by optimization of RNNs

Christopher Cueva
 Peter Wang
 Matthew Chin
 Xue-Xin Wei

CCUEVA@GMAIL.COM
 PETERWANG724@GMAIL.COM
 MATTCHIN35@GMAIL.COM
 WEIXXPKU@GMAIL.COM

Columbia University

Recent work suggests that goal-driven training of neural networks can be used to model the neural activity in the brain. Here we ask if an artificial neural network can recover both neural representations and the anatomical properties of biological circuits. We addressed this question in a system where the connectivity and the functional organization have been characterized, namely, the head direction circuits of the rodent and fruit fly. Specifically, we trained recurrent neural networks (RNNs) to estimate head direction (HD) through the integration of angular velocity inputs (AV). first, we found that the trained network could accurately track head direction. We then examined the firing properties of model units in the trained RNN by plotting joint HD*AV tuning. This revealed two distinct classes of units whose functional architecture can be mapped onto the fly HD system. The first class of units exhibited HD tuning with minimal AV tuning, and we refer to these as Ring Units. The second class of neurons were tuned to both HD and AV and could be further subdivided into two populations that favor clockwise/counterclockwise rotation, referred to as CW/CCW Shifters, respectively. Connectivity of trained RNNs exhibit characteristics that are consistent with fly physiology, including local excitation and global inhibition in Ring Units, and asymmetric excitation from Shifters to Ring units. We then performed a series of perturbation experiments by 'lesioning' subsets of connections between different classes of neurons. Results suggest that Ring Units are responsible for maintaining a stable bump, while CW/CCW Shifters are involved in shifting the activity bump CW/CCW, respectively. Overall, our results suggest that optimization-based RNN models can recapitulate the structure and function of biological circuits, suggesting that they can be used to study the brain at the level of both neural activity and anatomical organization.

I-116. State space oscillator models to identify and parameterize oscillatory signals in EEG

Amanda Beck¹
 Hugo Soulat²
 Emily Stephen³
 Patrick Purdon^{4,5}

AMBECK@MIT.EDU
 HSOULAT@MIT.EDU
 EMILYPS14@GMAIL.COM
 PATRICK.PURDON@MGH.HARVARD.EDU

¹Massachusetts Institute of Technology

²University College London

³University of California, San Francisco

⁴Massachusetts General Hospital

⁵Harvard University

Oscillatory neuronal activity reflects the functional coupling of brain networks and can be measured with local field potentials, or in the aggregate at the scalp with electroencephalogram (EEG). Many well-established oscillatory patterns have been characterized in response to tasks or states of arousal. These patterns are often composed of multiple oscillations, such as slow waves and spindles during sleep, slow oscillations and alpha waves during anesthesia, or harmonic series representing nonlinear oscillations such as mu-rhythms. In some scenarios, the data analysis problem focuses on determining whether or not an oscillation is present; in others, the problem focuses on separating component oscillations so that their distinct features or interactions may be studied, revealing new information [1].

A widely used technique for isolating these oscillatory signals is to bandpass filter the signal with cutoff frequencies that vary throughout neuroscience literature. Bandpass filtering requires arbitrary decisions about frequency band cutoffs and has inherent issues with faithfully re-creating non-sinusoidal waveform shapes. We propose an alternative method to avoid making these arbitrary decisions and use a data-driven approach to retrieve the complex signal with a more flexible model that allows for non-sinusoidal features and varying oscillatory characteristics. We use a low-dimensional ARMA model, reminiscent of the second order differential equations used to model physically oscillating systems, to define a single oscillatory component and then naively add additional components to the multivariate state equation as identified by iterating through a greedy algorithm on the innovations of the previous model until the innovations are not significantly different from white noise. We estimate these parameters using an expectation maximization algorithm and incorporate Von Mises prior probability distributions on the center frequency. This procedure allows for a more faithful characterization of oscillatory power and waveforms that can reveal information not visible because of distortion in traditional bandpass filtering.

I-117. Dissection of low-dimensional models of V1 circuitry with multiple inhibitory types

Kenneth Miller¹
 Agostina Palmigiano¹
 Nataliya Kravnyukova²
 Daniel Mossing³
 Hillel Adesnik³
 Tatjana Tchumatchenko²

KDM2103@COLUMBIA.EDU
 AP3676@COLUMBIA.EDU
 NATALIYA.KRAYNYUKOVA@BRAIN.MPG.DE
 DAN.MOSSING@GMAIL.COM
 HADESNIK@BERKELEY.EDU
 TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

¹Columbia University

²Max Planck Institute for Brain Research

³University of California, Berkeley

Advances in experimental research in the last several decades revealed previously inconceivable detail in the inhibitory microcircuitry of the cerebral cortex. In V1, three main inhibitory classes PV, Sst, and Vip compose 80% of GABAergic interneurons and belong to a precisely sculpted circuit but their differential contribution to cortical

computations remains poorly understood. Here, we exploit the biological microcircuit's properties to mathematically dissect low-dimensional models of interneuronal circuits and study the response properties of such networks to stimuli of varying size and contrast.

Specifically, we analyze three- and four-dimensional models of population rates, which correspond to different possible circuit motifs. We assume that the input-output relation of network populations is rectified power-law. We map the steady-state solutions to the zero crossings of a polynomial in one variable and derive population responses to perturbations and stability conditions for these systems to find that i) every second polynomial zero-crossing corresponds to an unstable steady state ii) conditions on top of those that have to be satisfied by a 2D system are automatically satisfied if the excitatory population is itself stable and that iii) these conditions can be rewritten in terms of the responses to weak perturbations, linking stability and response properties in a single equation.

To investigate connectivity regimes consistent with observed V1 responses, we use cell-type-specific two-photon calcium-imaging recordings obtained while locomoting mice are shown gratings of different contrast levels. We construct surrogate contrast-response curves for each cell type by sampling from a Gaussian distribution with mean and standard deviation taken from the data and its standard error. This ensemble allows to select several data-compatible model parameters for which the network steady-states provide best non-negative least-squares fit. We then analyze how increases in contrast modify the structure of the polynomial, and how complex stimulus-size manipulations modulate the circuit's effective connectivity.

I-118. Predicting optogenetic responses through full-dimensional models of interneuronal circuitry

Agostina Palmigiano¹

Daniel Mossing²

Hillel Adesnik²

Kenneth Miller¹

Francesco Fumarola³

AP3676@COLUMBIA.EDU

DAN.MOSSING@GMAIL.COM

HADESNIK@BERKELEY.EDU

KDM2103@COLUMBIA.EDU

FRANCESCO.FUMAROLA@RIKEN.JP

¹Columbia University

²University of California, Berkeley

³RIKEN Center for Brain Science

Optogenetic perturbation of in-vivo cortical circuits offers a novel angle to characterize the operating regime of cortex. Experiments in mouse show that weak optogenetic perturbation to inhibitory populations decreases both the mean excitatory and, paradoxically, the mean inhibitory activity, consistently with inhibition-stabilized-network paradigms. Surprisingly, single-cell stimulation produces a spatially structured inhibitory halo around a local excitatory blob. There is currently no understanding of these results given the presence of multiple inhibitory cell-types, and no model for the structure of inhibitory stabilization within multiple inhibitory populations. Also lacking is the fundamental link between low-dimensional (LD) population-rate models yielding those predictions, and the high-dimensional (HD) models incorporating the full population of each cell-type.

Here we bridge the gap between LD models of V1 circuitry and HD networks of multiple interneuronal types. Incorporating the precisely sculpted architecture of GABAergic connectivity, we find that LD circuit response to perturbation possesses several parameter-independent invariants; this is moreover tightly linked to the stability of network sub-circuits. To investigate the robustness of these results, we map an LD circuit to an HD one with log-normally distributed synaptic weights whose means are the LD inter-population weights. Relying on field-theoretical methods for the analysis of such matrix ensembles, we calculate the response of the HD circuit to optogenetic perturbations affecting arbitrary fractions of any cell-type population. The final results reveal that: i) LD and HD circuit responses are intuitively linked only for full-population perturbations unlikely to occur in current experimental setups; ii) partial perturbations, surprisingly, give rise to a bimodal distribution of responses; iii) the fraction of paradoxical responses is in fact a nonlinear function of the fraction of optogenetically stimulated cells. We then proceed to study the extension of the theory to space-dependent synapses, and recover the specific profile of the inhibitory response halo found in the data.

I-119. Recurrent neural circuits overcome partial inactivation by compensation and re-learning

Colin Bredenber
 Roozbeh Kiani
 Cristina Savin
 New York University

CJB617@NYU.EDU
 ROOZBEH@NYU.EDU
 CS5360@NYU.EDU

Technical advances in artificial manipulation of neural responses have precipitated a surge in studies of causal contributions of brain circuits to cognition and behavior. However, since underlying neural computations are often distributed across multiple regions, the interpretation of these experimental results is fraught with difficulty. A prime example is perceptual decision-making based on integration of sensory evidence, which recruits multiple frontoparietal cortices and subcortical areas, and has been a target for many causal studies, with contradictory outcomes. Making sense of inactivation results requires a theoretical framework for the systematic exploration of causal manipulations. Here we take a step in this direction, using recurrent neural networks (RNNs) trained to emulate the computations that underlie discrimination of random dots motion direction. Whereas inactivation of a sensory bottleneck in the network has dramatic effects on choice accuracy and reaction times, the inactivation of the sub-circuits that implement integration can have diverse effects. In an architecture where a single pool of neurons is responsible for integration, network dynamics implement a shallow bistable attractor, where the network state moves proportional to the evidence supporting the competing choices, effectively implementing linear evidence accumulation. Inactivation of neurons in this circuit causes a behavioral deficit only if the attractor structure is disrupted, with the magnitude of loss-of-function strongly correlated with qualitative phase plane changes. Critically, we also show that networks that learn during inactivation can recover function quickly, often much faster than the original training time. Our results demonstrate that challenges of interpreting "causal" studies can be overcome only through models that account for the complexity of the manipulated system. Using our framework, we explain past empirical observations by clarifying how complex circuits compensate for manipulations and could develop new functional interactions substantially different from the original.

I-120. Multiscale low-dimensional neural dynamics explain naturalistic 3D movements

Hamidreza Abbaspourazad¹
 Mahdi Choudhary²
 Yan Wong³
 Bijan Pesaran²
 Maryam Shanechi¹

¹University of Southern California

²New York University

³Monash University

HABBASPO@USC.EDU
 MC5368@NYU.EDU
 YAN.WONG@MONASH.EDU
 BIJAN@NYU.EDU
 SHANECHI@USC.EDU

Studies of motor cortical dynamics have largely focused on analyses of low-dimensional dynamics in the spiking activity of a population of neurons. However, motor function relies on computations implemented across multiple spatiotemporal scales of brain activity ranging from small-scale processes measured by spiking activity to integrated action of larger-scale brain networks reflected in local field potentials (LFP). Spiking and LFP activity measure different biological processes, have different time-scales and signal characteristics, and may differentially encode aspects of a sensorimotor behavior. How low-dimensional dynamics at the different spatiotemporal scales reflected in spiking and LFP activity relate to each other and to naturalistic behavior remains largely unexplored. Here, we build novel multiscale latent state-space models of spiking, LFP and combined spike-LFP activities recorded from non-human primates during naturalistic 3D reach-and-grasp movements. We partition the dynamics into modes, each characterized with a unique pair of time-decay and frequency. Each mode specifies how one component of neural response decays and rings in time in response to excitations. We compute

the modes and how well each mode is predictive of behavior. We find that both spiking and LFP network activity exhibited several principal modes, among which one was dominant in predicting behavior. Interestingly, despite the existence of several distinct modes in spiking and LFP, this predictive mode was shared between them, across experimental sessions and monkeys. Moreover, the predictive mode did not simply replicate behavior modes and was present even in the low-frequency bands of LFP. The predictive mode's time-decay and frequency explained behavior prediction accuracy. The predictive mode time-decay was longer than all other modes, suggesting that naturalistic behavior relies on integrating multiscale motor cortical dynamics over longer time-periods. We propose that multiscale, low-dimensional motor cortical state dynamics underlie the control of naturalistic reach-grasp movements due to the integrated action of large-scale brain networks.

I-121. The PHATE of learning networks

Scott Gigante¹
Adam Charles²
Smita Krishnaswamy¹
Gal Mishne³

SCOTT.GIGANTE@YALE.EDU
ADAMSC@PRINCETON.EDU
SMITA.KRISHNASWAMY@YALE.EDU
GMISHNE@UCSD.EDU

¹Yale University

²Princeton University

³University of California, San Diego

Understanding the evolution of neuronal networks is a key challenge in the analysis of learning, memory and other neurological phenomena. To this end, we introduce a novel visualization algorithm that reveals the internal geometry of dynamically evolving networks: Multislice PHATE (M-PHATE), the first method designed explicitly to visualize how a network's response to stimuli changes over time. M-PHATE leverages a novel multi-slice kernel to build a time-aware graph over high-dimensional data, visualizing this graph in two dimensions using techniques from manifold learning. Using artificial neural networks as a test case, we demonstrate that our visualization provides intuitive, detailed summaries of the learning dynamics. Furthermore, M-PHATE better captures both the dynamics and community structure of the test networks as compared to visualization based on standard dimensionality reduction methods (e.g., ISOMAP, t-SNE). We demonstrate M-PHATE with two vignettes: continual learning and generalization. In the former, the M-PHATE visualizations display the mechanism of "catastrophic forgetting" which is a major challenge for learning in task-switching contexts. In the latter, our visualizations reveal how increased heterogeneity among hidden units correlates with improved generalization performance. Much as in common measurements of biological neural networks, our visualization requires only to observe the activations of the artificial neurons with respect to a representative subset of input data; in this way, our work can be readily extended to develop biological understanding of learning and the evolution of neural pathways.

I-122. Learning and stabilizing manifold attractors in heterogeneous networks

Alexander Rivkind¹
Ran Darshan²
¹The Weizmann Institute of Science
²HHMI Janelia Research Campus

ALEXANDER.RIVKIND@WEIZMANN.AC.IL
DARSHANR@HHMI.ORG

Continuous attractor dynamics are believed to constitute the neural representations of various computations, such as spatial working memory and path integration. Within this framework, a set of continuous variables at a specific time point is represented as a point on a stable low-dimensional manifold. Elegant models were constructed to embed such manifolds. Unfortunately, the modeled dynamics collapses to a small number of isolated attractors even under slightest alterations in the network connectivity. Inevitable in biological systems, such alterations are attributed to any synaptic heterogeneities or irregularities in the environment. Here we show how synaptic

heterogeneities can be rather utilized to stabilize manifolds than to destabilize them. To this end, we train random recurrent networks to produce manifold attractors. Starting with the canonical example of a low-dimensional manifold - the ring attractor - we use mean field approach to derive the conditions for the stability of a learned heterogeneous ring manifold. We then extend the perfectly symmetric ring attractor to arbitrary closed curve attractors. We find that the geometry of the attractor strongly affects its stability, with some curves being learnable only in presence of heterogeneities. Heterogeneities contribute to robustness by both facilitating the contraction toward the manifold and by limiting drift on its surface. By examining the network as an encoder-decoder system, we show how robustness is achieved by adding roughness to the originally smooth manifold. This can either be a result of learning from a finite number of points on the manifold or emerge by learning a manifold on top of autapse connectivity, with the latter being introduced as an example of extreme roughness. Our work emphasizes how synaptic heterogeneities, which are catastrophic for standard continuous attractor models, can be used during learning to stabilize the attractor and making it robust to, unforeseen, conditions.

I-123. Spatiotemporal neural correlations and network dynamics

Yanliang Shi¹
 Roxana Zeraati²
 Anna Levina³
 Tatiana Engel¹

SHI@CSHL.EDU
 ROXANA.ZERAATI@TUEBINGEN.MPG.DE
 ANNA.LEVINA@UNI-TUEBINGEN.DE
 ENGEL@CSHL.EDU

¹Cold Spring Harbor Laboratory

²University of Tübingen

³University of Tübingen

With fast development of recording techniques, simultaneous recordings of large groups of neurons reveal widely distributed spatiotemporal neural correlations in the cortex. Pairwise neural correlations are related to functional properties of neurons. They affect sensory information processing, learning and plasticity, and cognitive functions such as attention. At the population level, the spatial and temporal modes of correlations intermingle, which possibly reflects underlying anatomical circuit structure, network dynamics and operating regimes of neural activity. However, a systematic approach to disentangle the mixed patterns of spatial and temporal modes in correlations has not been fully developed. Here we develop a theoretical framework that relates the spatial and temporal modes of pairwise neural correlations to the network connectivity structure and the operating regime of dynamics in interacting neurons. We analyze spatiotemporal correlations in network models of binary units with different connectivity structures and dimensions. We derive analytical expressions for spatial and temporal correlations and verify them with numerical simulations. Our theory demonstrates how multiple timescales in auto- and cross-correlations arise from spatial interactions between units. We find that because of spatial dependence of interactions, each timescale is associated with fluctuations of a particular spatial frequency and makes hierarchical contributions to the correlations. We then study how local versus distributed spatial connectivity shapes the timescales and spatial patterns of neural correlations. Finally, we evaluate the influence of external inputs on the operating regime of the global network activity and show how it affects the timescales of correlations. Our work reveals the relationship between spatial and temporal patterns of correlations, which is determined by the network structure, dynamics and the operating regime of population activity. Analytical methods developed here can be used to extract and interpret spatiotemporal features of neural dynamics during sensory and cognitive processing, to advance understanding of neural circuit functions.

I-124. Neural distribution alignment via hierarchical optimal transport

Max Dabagia
 Eva Dyer

MAXDABAGIA@GATECH.EDU
 EVADYER@GATECH.EDU

Georgia Institute of Technology

As the scale of neural recording grows, the prospect of matching individual neurons across diverse datasets becomes daunting. Fortunately, experimental evidence shows low-dimensional latent dynamics drive the activity of many neural circuits, with robust, stable structure across recording sessions and even subjects. However, making comparisons between the latent spaces of different datasets requires that they are first aligned, to account for transformations that may have been induced by the process of spike sorting, dimensionality reduction, and other errors. Automating the process of alignment, especially when the latent space has more than a few dimensions, is challenging, and leveraging additional structure is necessary to make the problem tractable. Here, we present an approach to do so by identifying subspaces within the latent state spaces to be aligned. Using a metric from the study of optimal transport, we quantify the distance between distributions lying on subspaces as the work required to morph one into the other, and identify similar subspaces based on this notion of distance. The transformation to align two low-dimensional, linear subspaces has a closed form, and so by aggregating the transformations between pairs, we can efficiently find a global transformation that aligns the entire state space. We apply this technique to macaque motor cortex and rat striatum, and show that neural activity can be aligned across recordings, as well as to external task-related covariates. By coupling dimensionality reduction and distribution alignment, neural datasets can be compared robustly across time and even across different subjects. Our unsupervised strategy paves the way for more robust brain decoders, and will facilitate the study of how the latent spaces of neural activity are tied to processes like attention and learning.

II-1. Optimal population response for robust linear readout

Blake Bordelon
Cengiz Pehlevan
Harvard University

BLAKE_BORDELON@G.HARVARD.EDU
CPEHLEVAN@SEAS.HARVARD.EDU

Neural population codes face a fundamental tradeoff between discriminability and robustness. High dimensional codes are capable of encoding large amounts of information per neuron, allowing precise discrimination between stimuli, but are often brittle and generalize poorly when exposed to noise. Conversely, correlations between the mean activities of neurons introduce redundancy in the code, reducing brittleness at the cost of information content. We explore the tradeoff between these two coding desiderata, discrimination and robustness, in the context of a linear readout from population responses. Calculation of the classification error due to small stimulus perturbations allows a derivation of constraints on the spectral properties of robust codes. Generally, we show that robust perceptual predictions of a readout neuron can arise even if the population code is not smooth in the large neuron limit. Our analysis suggests that optimal codes are approximately white for low frequencies but filter out high frequency content to allow for smooth predictions. Such a code, while higher dimensional than smooth codes, still generalizes better. We compare our theoretical predictions to calcium imaging recordings from Mouse V1 in response to grating stimuli.

II-2. Slow coordinated activity across hemispheres in PFC predicts pupil size

Akash Umakantha
Ryan Williamson
Byron Yu
Matthew Smith
Carnegie Mellon University

AUMAKANT@ANDREW.CMU.EDU
RCW1@ANDREW.CMU.EDU
BYRONYU@CMU.EDU
MATTSMITH@CMU.EDU

Neural activity observed within a brain area may reflect local computations within the area, inputs from another area, or shared computations/inputs across many areas. Ideally, these distinct mechanisms could be studied separately. However, because multiple processes can influence groups of neurons, it is not obvious how to separate the neural signals that should be attributed to each process. Here we investigate the different behavioral

roles of neural variability shared across hemispheres and neural variability local to each hemisphere. To do this, we applied an extension of probabilistic canonical correlation analysis (called pCCA- FA) to bilateral prefrontal cortex (PFC) array recordings. This allowed us to identify latent variables representing activity shared across hemispheres, as well as latent variables representing activity local to each hemisphere. We found that variability shared across hemispheres was dominated by a process that slowly evolved across trials. This process was highly correlated with trial-to-trial fluctuations in mean pupil diameter, suggesting that the variability shared across hemispheres represents a global cognitive state signal (e.g., arousal). We also found that activity local to each hemisphere did not exhibit slow time-scale dynamics, nor was it correlated with mean pupil diameter. Taken together, our approach allowed us to isolate a signal that is shared across both hemispheres of PFC and is consistent with a slow global modulatory process related to cognitive state.

II-3. Mixed representations in retrosplenial cortex mediate flexible navigation decisions in mice

Shinichiro Kira¹
Giuseppe Pica²
Stefano Panzeri²
Christopher Harvey¹

SHIN_KIRA@HMS.HARVARD.EDU
GPICA89@GMAIL.COM
STEFANO.PANZERI@IIT.IT
CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

¹Harvard University

²Istituto Italiano di Tecnologia

Decision-making in a naturalistic setting often occurs as animals navigate through an environment while flexibly combining internal signals, such as past experience, with sensory signals to guide locomotor actions. The neural mechanisms for such decisions are not fully understood because navigation and decision-making have often been studied in separate experimental paradigms. To study flexible navigation decisions, we designed a delayed match-to-sample task for mice based on movement through a virtual reality T-maze. In the maze, two cues were presented sequentially separated by a delay. The mouse was required to combine memory of the first cue (sample cue) and visual information about the second cue (test cue) to make an appropriate choice (right or left turn). This task separates various types of processing that are often entangled in tasks with fixed sensory-motor associations. To identify dorsal cortical areas necessary for accurate performance of the task, we used unbiased methods based on optogenetic inhibition. Inhibition of posterior parietal cortex (PPC), retrosplenial cortex (RSC), or visual cortices reduced behavioral choice accuracy. The effect was most prominent when these areas were inhibited during the trial phase in which mice had to combine visual and memory information to make a choice. We recorded neural activity from these areas using two-photon calcium imaging during the task and found that each area had cells with heterogeneous activity patterns spanning selectivity to all aspects of the task. Notably, relative to PPC and V1, RSC contained a large fraction of neurons with mixed representations of the sample and test cues. These representations were stronger on correct trials and were largely absent on error trials. Our findings reveal an essential role for a visual-parietal-retrosplenial network in navigation decisions. Further, we propose a novel function of RSC in flexibly combining memory signals with visual signals to inform navigation decisions.

II-4. Blockage of cerebellar outflow changes the spatiotemporal organization of muscle activity

Nirvik Sinha¹
Abdulraheem Nashef²
Sharon Israely²
Yifat Prut²

NIRVIKSINHA2024@U.NORTHWESTERN.EDU
ABDULRAH.NASHEF@MAIL.HUJI.AC.IL
SHARON.ISRAELI1@MAIL.HUJI.AC.IL
YIFATPR@EKMD.HUJI.AC.IL

¹Northwestern University

²The Hebrew University of Jerusalem

The cerebello-thalamo-cortical (CTC) pathway is the dominant mediator of cerebellar fine tuning of motor outputs in primates. This study aimed to determine the impact of disruption of this pathway on the spatiotemporal organization of muscle activation patterns in a behaving primate. One adult monkey (*Macaca fascicularis*) was trained to perform a planar, center-out reaching task (to 8 evenly distributed peripheral targets) following a “GO” signal using its left upper limb. The CTC pathway was disrupted by high frequency stimulation (HFS) of the superior cerebellar peduncle (SCP) during the task performance. EMG was recorded from 16 muscles of the same limb during 128 control and 112 HFS trials. The muscle activity of the period -1000ms to +2000ms (relative to the movement onset) was extracted from each trial for further processing. Sets of four muscle synergies were obtained from the control and HFS EMG trials using the formulation of d’Avella et al. The temporal profile of one synergy showed increased activation of the wrist extensors in HFS as compared to control. Another synergy showed decreased activation of the wrist flexors in HFS as compared to control. Additionally, in three out of the four synergies, the onset delay relative to the start of the extracted period was prolonged in HFS as compared to control. Finally, using a multivariate measure, the amount of unstructured variance i.e. noise was observed to be higher in HFS EMG data as compared to that of control. These results indicate that the blockage of the CTC pathway by HFS in a behaving primate disrupts the proper timing and composition of muscle synergies and increases the amount of noise in its encoding.

II-5. Inferring function from structure with connectome and task constrained neural networks

Janne Lappalainen¹
Sridhama Prakhya¹
Mason McGill²
Fabian Tschopp³
Srini Turaga¹

LAPPALAINENJ@JANELIA.HHMI.ORG
PRAKHYAS@JANELIA.HHMI.ORG
MASON.B.MCGILL@GMAIL.COM
TSCHOPFA@STUDENT.ETHZ.CH
STURAGA@GMAIL.COM

¹HHMI Janelia Research Campus

²California Institute of Technology

³ETH Zurich

The structure of a neural network cannot be uniquely inferred from the function computed by the circuit. However, is it possible to infer function from the structure? In this study, we show how for sufficiently structured and sparsely connected neural networks, it is possible to predict the tuning properties of individual neurons in a network, with knowledge of the neural connectivity and a high-level understanding of the overall computational function of the circuit.

We trained sparse feedforward neural networks (teacher network) to classify handwritten digits. We then asked if training a student network with the same connectivity pattern to perform the same task would yield a network with the same tuning properties in its hidden units. We found that for sparsely connected networks, the tuning properties of the hidden units of the student network strongly correlated with the tuning of the same unit in the teacher network, while for densely connected networks there was no correlation. This provides hope that a connectome constrained network trained to perform the same task as the original circuit in the brain might provide a faithful simulation.

To test our hypothesis, we constructed a simplified connectome-based computational model of the first two stages of the fly visual system, the lamina and medulla. The resulting hexagonal lattice point neuron convolutional network with threshold linear dynamics was trained using backpropagation through time to compute optic flow, ego motion and depth estimation in natural scene videos. Such networks discovered the well-known direction selectivity and on/off tuning properties in T4 and T5 neurons.

Our results suggest that for sufficiently sparsely connected networks, connectome and task constrained networks can lead to an understanding of circuit function from structure.

II-6. Representation of uncertainty in macaque visual cortex

Zoe Boudy-Singer¹

Olivier Henaff²

Kristof Meding³

Corey Ziemba¹

Robbe Goris¹

ZOEBSINGER@UTEXAS.EDU

OLIVIER.HENAFF@GMAIL.COM

KRISTOF.MEDING@UNI-TUEBINGEN.DE

ZIEMBA@UTEXAS.EDU

ROBBE.GORIS@UTEXAS.EDU

¹University of Texas at Austin

²DeepMind, London

³University of Tubingen

Sensory systems must represent a world that cannot be known perfectly. Uncertainty about the world can arise externally, when sensory cues are unreliable, or internally, when their neural representation is unreliable. Because perception is fundamentally uncertain, perceptual tasks are often formalized as statistical inference problems which require observers to take the reliability of sensory information into account. This implies that the neural circuits which mediate perception encode uncertainty about the stimulus. How they do so is a topic of debate. We propose a view of visual cortex in which average neural response strength encodes stimulus features, while cross-neuron variability in response gain encodes the uncertainty of these features. Thus, visual neurons behave as if they have not one but two receptive fields: the first one governs the mapping of stimulus features onto response mean, and the second one governs the mapping of stimulus uncertainty onto gain variability (“The uncertainty receptive field”). To test our theory, we studied spiking activity of neurons in macaque V1 and V2 elicited by repeated presentations of stimuli whose uncertainty was manipulated in distinct ways. We show that gain variability of individual neurons is tuned to stimulus uncertainty, that this tuning is invariant to the source of uncertainty, and that it is specific to the features encoded by these neurons. We demonstrate that this behavior naturally arises from known gain-control mechanisms. Analysis of the temporal dynamics of gain fluctuations revealed that they are slow. Thus, if downstream circuits use gain variability to assess stimulus uncertainty, they cannot decode this from individual neurons over time. We derived how downstream circuits can jointly decode stimulus features and their uncertainty from sensory population activity. Together, these results establish cross-neuron variability in response gain as a candidate currency of uncertainty in sensory cortex.

II-7. Theory of gating in recurrent neural networks

Kamesh Krishnamurthy¹

Tankut Can²

David Schwab²

KAMESHK@PRINCETON.EDU

TCAN@GC.CUNY.EDU

DSCHWAB@GC.CUNY.EDU

¹Princeton University

²Initiative for Theoretical Sciences, CUNY

Understanding the collective dynamics of networks of neurons is a major research direction in theoretical neuroscience that can inform the analysis of recently available high-dimensional neurophysiological data. Most of the theoretical work in understanding the dynamics of these networks has focused on models with ‘additive’ interactions, where the input to a neuron is a weighted sum of the output of the rest of the network. However, there is ample evidence from neurophysiology that neurons can have gating – i.e. multiplicative – interactions; for example, the output of one neuron can shunt the output of another neuron thereby effectively silencing it. Such gating interactions lead to qualitatively different behavior at the level of individual neurons, and are likely to have more dramatic effects at the network level. Furthermore, researchers in machine learning have independently and empirically found that model neural networks which include gating interactions are capable of learning significantly more complex tasks than neural network models with only additive interactions. Thus, gating interactions have significant implications for collective dynamics which appear to be useful for information processing and learning tasks. Here we leverage tools from random matrix theory and the field theory of disordered systems to develop a theory of gating in a canonical neural network model which captures the essential elements of gating between

its units. The theory allows us to probe how various aspects of gating shape the collective dynamics. Specifically, we characterize the appearance of long timescales and marginal stability due to gating. The theory also allows us to study asymptotic behavior by means of the maximum Lyapunov exponent. Furthermore, we use our theory to elucidate dynamical aspects of learning, in particular to understand the dynamics of gradient descent based learning.

II-8. Dynamical mechanisms of flexible timing by temporal scaling

Manuel Beiran¹
Mehrdad Jazayeri²
Srdjan Ostojic¹

MANUEL.BEIRAN@ENS.FR
MJAZ@MIT.EDU
SRDJAN.OSTOJIC@ENS.FR

¹Ecole Normale Supérieure

²Massachusetts Institute of Technology

Animals can flexibly perform a given action at vastly different speeds. Recent analyses of neural recordings in monkeys performing flexible timing tasks have shown that neural activity is temporally stretched when performing the same action at different time intervals. At the level of the neural population, flexible timing relies on neural activity evolving at different speeds along a low-dimensional invariant trajectory in neural state-space [Wang et al. 2018]. In this work, we investigated the dynamical mechanisms that allow a network of recurrently connected units to implement flexible timing through such temporal scaling along an invariant manifold. We relied on the framework of recurrent networks with low-rank connectivity: we first trained recurrent networks with constrained rank to perform flexible timing, and next reverse-engineered them to isolate the mechanisms used to solve the task. Applying mean field theory, we then analytically determined the core dynamical phenomenon and identified classes of networks that implement it. We found that recurrent networks with rank-two connectivities are able to precisely and robustly perform a flexible timing task spanning a very broad range of timescales. Similarly to experimental findings, such networks relied on an invariant, low-dimensional manifold along which the speed is controlled by external cues. Mean-field theory allowed us to analytically determine the origin of this manifold and infer a class of rank-two networks that give rise to it: an asymmetry between the two rank-one patterns in the connectivity generate a slow manifold that corresponds to a ring attractor closeby in parameter space, while an overlap between the speed input command and the connectivity structure adjusts the speed on this manifold. Altogether, we identified a novel dynamical mechanism for temporal flexibility that is robust, simple to implement, and can be fully understood analytically.

II-9. A learning-forgetting tradeoff from the perspective of control: a *Drosophila* case study

James Bennett¹
Luca Manneschi²
Eleni Vasilaki²
Thomas Nowotny¹

JAMES.BENNETT@SUSSEX.AC.UK
LMANNESCHI1@SHEFFIELD.AC.UK
E.VASILAKI@SHEFFIELD.AC.UK
T.NOWOTNY@SUSSEX.AC.UK

¹University of Sussex

²University of Sheffield

Learning accurate reward predictions is important for making meaningful comparisons between multiple options and, thus, effective decision-making. Recent studies also argue that forgetting may benefit decision-making, allowing it to quickly adapt when reward contingencies are liable to change, and that neural circuits exhibit mechanisms to actively degrade memories stored in synapses. Here, we show how active forgetting also results in the loss of accurate reward predictions, and ask how neural circuits may balance a tradeoff between accurate reward predictions and adaptive decision-making. We investigate an augmentation to delta-rule learning, inspired by

control theory, that helps to re-establish accurate reward predictions while maintaining the benefits of forgetting. Specifically, we frame learning as a control problem in which the available reward provides a target state and the reward prediction is the system's current state. We present a neural circuit model that implements an approximation to Proportional-Integral-Derivative (PID) control, find a mapping between our approximation and true PID control, and illustrate how a subcircuit in the *Drosophila* mushroom body – a structure key to reward learning – exhibits the necessary connectivity to implement our PID-inspired circuit model.

II-10. Sequential and efficient neural-population coding of complex task information

Sue Ann Koay¹
Stephan Thiberge²
Carlos D Brody¹
David Tank¹

KOAY@PRINCETON.EDU
THIBERGE@PRINCETON.EDU
BRODY@PRINCETON.EDU
DWTANK@PRINCETON.EDU

¹Princeton University

²Princeton Neuroscience Institute, Princeton University

A crucial component of cortical computation is context, which may be described by multiple variables, both external and internal. However, the need to simultaneously represent many pieces of information in neural activity can also pose computational challenges for the brain. How are multiple variables represented together and read out, considering e.g. crosstalk? Do these representations include temporal context? We recorded from large neural populations in posterior cortices as mice performed a complex, dynamic task involving multiple interrelated visual, motor, cognitive, and memory variables. Our analysis conceptualizes the collective activity of neurons as a point in a high-dimensional neural state space, where each coordinate is the activity level of one neuron. We investigated the geometry of the neural representation by examining state-space directions defined by the neural encoding and decoding of task variables. The neural encoding was such that correlated task variables were represented by uncorrelated modes in an information-coding subspace. This can enable optimal decoding directions to be insensitive to neural noise levels. Across posterior cortex, principles of efficient coding thus applied to task-specific information with neural-population modes as the encoding unit, while individual neurons exhibited a spectrum of (prevalently positive) pairwise signal correlations. Remarkably, this encoding function was multiplexed with rapidly changing, sequential neural dynamics, yet reliably followed slow changes in task-variable correlations through time, so as to keep modes uncorrelated in the coding subspace. If fast sequential dynamics are not coordinated across neurons, changes in neural activity due to circuit dynamics could potentially be confounded by the read-out as being due to changes in task conditions. However, we show that neural time-modulations can be random and still implement stable encoding of task information. We explain this as due to a mathematical property of high-dimensional spaces, which the brain might exploit as a temporal scaffold.

II-11. Convolutional dictionary learning of stimulus from spiking data

Andrew Song¹
Bahareh Tolooshams²
Simona Temereanca³
Demba Ba²

ANDREW90@MIT.EDU
BTOLOOSHAMS@SEAS.HARVARD.EDU
SIMONA_TEMERANCA_IBANESCU@BROWN.EDU
DEMBA@SEAS.HARVARD.EDU

¹Massachusetts Institute of Technology

²Harvard University

³Brown University

The recording of neural activity in response to repeated presentations of an external stimulus is an established experimental paradigm in the field of neuroscience. Generalized Linear Models (GLMs) are commonly used to

describe how neurons encode external stimuli, by statistically characterizing the relationship between the covariates (the stimuli or their derived features) and neural activity. An important question becomes: how to choose appropriate covariates? We propose a data-driven answer to this question that learns the covariates from the neural spiking data, i.e. in an unsupervised manner, and requires minimal user intervention. Specifically, we cast the problem of learning the covariates (or templates) from the data as a Convolutional Dictionary Learning (CDL) problem, where the goal is to learn shift-invariant templates and the times when they occur. Our contribution is two-fold. First, we formulate an optimization objective with sparsity constraints, which accounts for the binary nature of the spikes. This aspect of the data poses an additional challenge as the observations cannot be assumed Gaussian, a common assumption in the CDL literature. Second, we propose iterative algorithms to solve this objective, with our key insight being that the observations need to be modified in a specific manner at each iteration. We apply our framework to neural spiking data recorded from the Barrel cortex of mice in response to periodic whisker deflections. Classical GLM analyses suggest that whisker velocity—obtained from the ideal whisker position programmed into the piezoelectrode used to move the whiskers—strongly modulates neural spiking. Our method obtains an estimate of whisker velocity, which suggests that during a deflection, whisker motion with respect to the piezoelectrode is highly variable. Moreover, when used as a covariate, this data-driven estimate of whisker velocity yields better goodness-of-fit, in terms of the Kolmogorov-Smirnov test, than GLMs.

II-12. Reinforcement learning models reveal quantitative reward prediction errors in ventral pallidum

David Ottenheimer¹

Bilal Bari¹

Elissa Sutlief¹

Kurt Fraser¹

Tabitha Kim¹

Jocelyn Richard²

Jeremiah Cohen¹

Patricia Janak¹

DOTT92@GMAIL.COM

BBARI1@JHMI.EDU

ESUTLIE1@JHMI.EDU

KFRASER3@JHU.EDU

TKIM87@JHU.EDU

RICHARDJ@UMN.EDU

JEREMIAH.COHEN@JHMI.EDU

PATRICIA.JANAK@JHU.EDU

¹Johns Hopkins University

²University of Minnesota

Reinforcement learning provides a rich framework for characterizing how individuals learn from their environment. In particular, the observation that ventral tegmental area (VTA) dopamine neurons signal a key feature of reinforcement learning—reward prediction errors (RPEs), or the difference between actual and expected reward outcomes—has led to the hypothesis that the brain employs reinforcement learning-like algorithms to organize reward-seeking behavior. A key area of inquiry is how dopamine neuron RPEs are computed from their inputs; previous work has identified components of RPEs in upstream regions but rarely a full RPE signal. Here, in a departure from these observations, we identify RPE encoding in a major input to the VTA, the ventral pallidum (VP). We used a computational approach to determine the impact of recent reward history on the activity of single VP neurons ($n=436$) recorded during a task where rats receive random presentations of two differentially preferred rewards. For each neuron, we fit a reinforcement learning model, as well as two simpler control models, to classify neurons as encoding a history-derived RPE (17%), the current outcome only (29%), or no outcome-specific information (54%). This population of VP RPE cells integrated expected value over multiple previous trials and demonstrated high fidelity with model-predicted firing rates on a trial-by-trial basis. Notably, there were more RPE cells in VP than in another input to VTA, the nucleus accumbens (8%, $n=183$), and the model described more of the variance in VP. We further found that VP RPE neuron activity correlated with task engagement, and optogenetic inhibition of VP, mimicking a negative prediction error, led to a corresponding decrease in task engagement. Our results demonstrate a robust RPE signal upstream of dopamine neurons that mediates engagement in reward seeking, providing novel insight into the calculation of RPEs and their contribution to adaptive behavior.

II-13. Discrete latent states underlie sensory decision-making behavior in mice

Zoe Ashwood¹
 Anne Urai²
 Valeria Aguilon²
 International Brain Laboratory³
 Anne Churchland²
 Alexandre Pouget⁴
 Jonathan Pillow¹

ZASHWOOD@PRINCETON.EDU
 ANNE.URAI@GMAIL.COM
 VALERIA.AGUILLON@INTERNATIONALBRAINLAB.ORG
 INFO@INTERNATIONALBRAINLAB.COM
 CHURCHLAND@CSHL.EDU
 ALEX.POUGET@GMAIL.COM
 PILLOW@PRINCETON.EDU

¹Princeton University

²Cold Spring Harbor Laboratory

³University College London

⁴University of Geneva

Classical models of perceptual decision-making assume that a single strategy or policy describes how an animal integrates sensory evidence and uses it to form a decision on each trial. Here we provide analyses showing that this common view is incorrect. Specifically, we use a latent variable modeling framework to show that decision-making behavior in mice reflects a complex interplay between different strategies, which alternate on a timescale of tens of trials. This provides a powerful alternate explanation for “lapses” commonly observed during psychophysical experiments. Formally, our approach consists of a Hidden Markov Model (HMM), where each state corresponds to a different decision-making strategy. Each strategy is parameterized by a distinct Bernoulli generalized linear model (GLM), whose weights describe the factors driving decisions under that strategy. Traditional models of lapse behavior assume that mistakes on “easy” trials arise from changes in the animal’s attentional or motivational state that arise independently on each trial; lapse events are, thus, assumed to last for at most one or two trials. The traditional lapse model corresponds to a special case of the 2-state GLM-HMM. We fit the GLM-HMM to data from a cohort of mice trained to perform a contrast detection task and found that their behavior was far better described by a 3-state model in which states persist for many trials in a row. Thus, lapses do not arise independently, but reflect state dynamics that encourage the animal to remain in a relatively engaged or disengaged state for extended periods of time. The GLM-HMM framework therefore provides a powerful lens for identifying the discrete set of strategies underlying decision-making behavior; and to ask if different animals solve the task in the same way.

II-14. Human visual motion perception shows hallmarks of Bayesian structural inference

Sichao Yang¹
 Johannes Bill²
 Jan Drugowitsch²
 Samuel J Gershman²

SICHAO@CS.WISC.EDU
 JOHANNES_BILL@HMS.HARVARD.EDU
 JAN_DRUGOWITSCH@HMS.HARVARD.EDU
 GERSHMAN@FAS.HARVARD.EDU

¹University of Wisconsin-Madison

²Harvard University

Bayesian inference has emerged as a successful description of elementary human visual motion perception, but little is known about how we make sense of the complex, nested motion relations found in real-world dynamic scenes. Here we expand the computational understanding of human motion perception to the domain of structural inference by virtue of two theory-driven experiments. To do so, we leveraged a recently proposed framework for generating analytically tractable motion-structured stimuli. A first experiment asked participants to categorize the (often hierarchically nested) motion structure of short dynamic scenes. Their psychometric function, when analyzed in terms of the statistically correct Bayesian posterior—a key facility of our tractable stimulus design—, revealed the signature logistic shape of correct structural inference. We noticed that participants exhibited dis-

tinct perceptual error patterns between certain structures. A Bayesian observer model linking human responses to the statistical evidence in individual trials was able to account for these patterns and, even more, explained participants' responses with single trial resolution. We therefore hypothesized that the Bayesian description could provide a general account for human motion structure perception. We tested this hypothesis in a second experiment: a two-alternative forced choice task targeting human perception of highly ambiguous scenes, which we generated by morphing continuously between two prototypical motion structures. Strikingly, the Bayesian observer model from the first experiment qualitatively predicted human choices for the new, ambiguous scenes, thereby confirming the Bayesian model's generality. As an additional hallmark of probabilistic reasoning, participants' reported decision confidence correlated with the certainty in the correct posterior distribution. Taken together, our results suggest that humans perceive hierarchically arranged motion scenes in close similarity to artificial observers of Bayesian structural inference. Our behavioral task should readily transfer to neuroscientific animal experiments to shed light on the neural representations of motion structure perception.

II-15. Mechanistic model of a neural circuit for ZCA-type whitening in olfaction

Nikolai Chapochnikov¹
Cengiz Pehlevan²
Dmitri Chklovskii^{3,4}

NCHAPOCHNIKOV@FLATIRONINSTITUTE.ORG
CPEHLEVAN@SEAS.HARVARD.EDU
DCHKLOVSKII@FLATIRONINSTITUTE.ORG

¹flatiron Institute

²Harvard University

³Flatiron Institute

⁴New York University

Sensory systems and machine learning pipelines often perform data preprocessing steps before further analysis. However, it is unclear what preprocessing sensory systems actually implement and how exactly a biological neural circuit would perform such computations. To advance our understanding of sensory preprocessing, we studied the peripheral olfactory circuit of the drosophila larva, composed of 21 olfactory receptor neurons (ORNs) connected through feedback loops with several interconnected inhibitory local neurons (LNs). We analyzed the circuit connectome and ORN responses to a set of odors and found that the subspace S_W of ORNs \rightarrow LNs connection weight vectors significantly aligns with the subspace S_A of top PCA vectors of ORN activity. To understand the circuit's computation, we postulated an objective function, for which the online algorithm, consisting of neural dynamics and Hebbian synaptic updates, maps onto the olfactory circuit architecture. We analyzed this objective function and the associated circuit and found: 1. as a result of learning, the connectivity subspace S_W converges to the activity subspace S_A , as found in our data analysis; 2. the circuit performs a ZCA-type whitening of the ORN activity; 3. ORNs – LNs and LNs – LNs synaptic weights self-organize so that the LNs encode the top PCA directions of ORN activity and dampen these directions in ORNs via inhibitory feedback; 4. this computation leads to ORN and pattern decorrelation, thus allowing more efficient coding and better downstream odor discrimination. In summary, using a normative approach we can explain the experimentally observed relationship between synaptic weight vectors and ORN activity and link it to the computational task and algorithmic implementation by the neural circuit. Our study thus supports the idea that the olfactory circuit learns and dampens the most statistically prevalent direction of ORN activity as a preprocessing step before downstream analysis.

II-16. Meta-learning biologically plausible semi-supervised learning rules

Niru Maheswaranathan
 Keren Gu-Lamberg
 Sam Greydanus
 Luke Metz
 Jascha Sohl-Dickstein
 Google Brain

NIRUM@GOOGLE.COM
 KERENGU@GOOGLE.COM
 GREYDANUS.17@GMAIL.COM
 LMETZ@GOOGLE.COM
 JASCHASD@GOOGLE.COM

The question of how neurons embedded in a network update their synaptic weights to collectively achieve behavioral goals is a longstanding problem in systems neuroscience. Since Hebb's hypothesis that cells that fire together strengthen their connections, cellular studies have shed light on potential synaptic mechanisms underlying learning. These mechanisms have directly driven the careful hand design of biologically plausible models of learning and memory in computational neuroscience. However, these hand designed rules have yet to achieve satisfying success training large neural networks, and are dramatically outperformed by biologically implausible approaches such as backprop. We propose an alternative paradigm for designing biologically plausible learning rules: using meta-learning to Learn a parametric synaptic update rule which is capable of training deep networks. We demonstrate this approach by meta-learning an update rule for semi-supervised tasks, where sparse labels are provided to a deep network but the majority of inputs are unlabeled. The meta-learned plasticity rule integrates bottom-up, top-down, and recurrent inputs to each neuron, and generates weight updates as the product of pre- and post- synaptic neuronal outputs. The way in which the inputs to each neuron are combined to produce a learning signal, however, is itself a Meta-learned function, parameterized by a neural network. Critically, the meta-learned update rule integrates only neuron-local information when proposing updates—that is, our learning rule is spatially localized to individual neurons. After meta-learning, the resulting synaptic update rule is capable of driving task-relevant learning for semi-supervised tasks. We demonstrate this capability on two simple classification problems, where it outperforms backprop. In general, we believe meta-learning to be a powerful approach to finding more effective synaptic plasticity rules, which will motivate new hypotheses for biological neural networks, and new algorithms for artificial neural networks.

II-17. Revealing human brain dynamics underlying ECoG power via Directed Information in frequency

Sudha Yellapantula¹
 Kiefer Forseth²
 Nitin Tandon²
 Behnaam Aazhang¹

SUDHA@RICE.EDU
 KIEFER.FORSETH@UTH.TMC.EDU
 NITIN.TANDON@UTH.TMC.EDU
 AAZ@RICE.EDU

¹Rice University

²University of Texas Health Science Center

Historically, multiple invasive human electrocorticography (ECoG) studies have confirmed the existence of strong relationships between the task and the elicited high gamma ($h\gamma$) power responses ($>60\text{Hz}$) of ECoG data. However, the underlying information flow patterns among brain regions that cause these power responses remains unknown. This is problematic, as critical decisions like brain resections or stimulation frequencies solely depend on local activity, without considering connections from other areas. In this work, we uncover spatio-temporal and spectral brain dynamics from intracranial ECoG data from epileptic patients performing a picture-naming task. Our results identify the most likely set of brain regions responsible for eliciting each power response, along with the spectral basis of information flow among the brain regions, that correlates with the $h\gamma$ response. To obtain temporal dynamics, we divided the task into multiple sliding time windows. In each window, a data-driven measure of causality called Directed Information (DI) was used to measure directional information flow among all pairs of brain regions. The causality between brain regions was represented as a graph, with electrodes as the nodes, and normalized DI as the weighted directed edge metric. Analysis of the resulting time-varying DI graphs revealed

significant temporal correlations between $h\gamma$ power responses, and the total information entering a brain region. 43% of electrodes (310 electrodes in 7 patients) had these strong correlations. Moreover, both increases and decreases in $h\gamma$ power correlated with the increase in total information flow, allowing the identification of the brain regions responsible for power activations and suppressions. We developed a new metric, “Directed Information in frequency” (Dif). Dif revealed that the information flow associated with $h\gamma$ power was concentrated predominantly in theta or beta bands, depending on the specific brain region, thus providing a spectral view of connectivity underlying power, crucial for brain stimulation strategies.

II-18. Probing song motor circuit architecture and function in *Drosophila melanogaster*

Frederic Roemschied
Diego Pacheco Pinedo
Xinping Li
Elise Ireland
Mala Murthy
Princeton University

FREDERIC.ROEMSCHIED@GMAIL.COM
DPACHECO@PRINCETON.EDU
XINPING.LI@PRINCETON.EDU
EIRELAND@PRINCETON.EDU
MMURTHY@PRINCETON.EDU

A major open question in neuroscience is how activity within motor circuits gives rise to the dynamics of natural behavior. Answering this question can be difficult when motor circuits are incompletely mapped for most behaviors in most organisms. We address this problem through a novel approach - exploring a wide range of activity patterns via optogenetics in a small set of known neurons within the putative song production motor circuits of *Drosophila* (male flies sing to advertise themselves to the female by vibrating a single wing). For each activity pattern in each identified cell type, we record song production in freely behaving flies and analyze both song patterns and locomotor dynamics using custom software - from this large dataset we generate models (using Stimberg et al., eLife 2019) of the song motor circuits that include known cell types in addition to hypothesized new cell types. We then test these models using targeted optogenetic experiments and whole-brain calcium imaging to identify new neurons within the pathway. This framework for comprehensive circuit identification and characterization enabled identification of previously unknown neurons of the male song circuit, as well as a revision of the role of a known descending neuron in shaping song dynamics. In particular, we find that male song choice is shaped by a mutually inhibitory brain circuit with post-inhibitory rebound dynamics. In sum, our analysis framework, encompassing experiments informing circuit models, and models generating testable predictions, paves the way towards a comprehensive understanding of the motor circuits underlying complex behavior.

II-19. Attractor dynamics gate cortical information flow during decision-making

Lorenzo Fontolan¹
Arseny finkelstein¹
Michael Economo²
Nuo Li³
Sandro Romani¹
Karel Svoboda¹

FONTOLANL@JANELIA.HHMI.ORG
FINKELSTEINA@JANELIA.HHMI.ORG
MNE@BU.EDU
NUO.LI@BCM.EDU
ROMANIS@JANELIA.HHMI.ORG
SVOBODAK@JANELIA.HHMI.ORG

¹HHMI Janelia Research Campus

²Boston University

³Baylor College of Medicine

Sensory-based decisions involve information flow from sensory to motor areas. Decisions about future actions are held in memory until enacted, making them vulnerable to distractors. What neural mechanisms control decision robustness to distractors remains unknown. We trained mice to report detection of sensory stimuli with directional licking, following a brief delay epoch separating sensation and action. To precisely control activity in the sen-

sory cortex, we used direct photostimulation of genetically defined neurons in the vibrissal somatosensory cortex (vS1) as a 'sensory stimulus'. We found that distracting stimuli biased future actions only when presented during the early, rather than late, segment of the delay period – demonstrating temporal gating of sensory information flow. Gating occurred even though distractor-evoked activity percolated through the cortex without attenuation. Instead, choice-encoding dynamics in anterior lateral motor cortex (ALM) became progressively robust to distractors as time passed. We trained recurrent neural networks (RNN) to reproduce the heterogeneous activity of ALM neurons during trials without distractors. We showed that RNNs were able to reproduce ALM activity in trials with distractors, which they had not experienced during training, and to predict temporal gating as observed in the data. Reverse-engineering of trained RNN revealed that chosen actions were stabilized via attractor dynamics, which gated out distracting stimuli. Our results demonstrate a dynamic gating mechanism that operates by controlling the degree of commitment to a chosen course of action.

II-20. Representing relational knowledge in cognitive maps for generalization

Mona Garvert¹
Nicolas Schuck²
Christian Doeller¹

GARVERT@CBS.MPG.DE
SCHUCK@MPIB-BERLIN.MPG.DE
DOELLER@CBS.MPG.DE

¹Max Planck Institute for Human Cognitive and Brain Sciences

²Max Planck Institute for Human Development

The brain organises knowledge about spatial and non-spatial relationships in cognitive maps. Such a representation of relationships may facilitate goal-directed behavior by enabling generalization of information across related states. Yet, how this generalization works mechanistically, i.e. how knowledge about relationships between features in one domain guides inference about features in other domains is not known. Here, we combine a virtual reality task with computational modeling and functional magnetic resonance imaging (fMRI) to investigate whether humans generalize across spatially related states to infer reward values that were never directly experienced. In this task, spatial relationships between stimuli learned on day 1 predict reward relationships in a non-spatial choice task on day 2. We find that subjects not only update the stimulus-reward associations they experience directly, but they also use their knowledge about the spatial relationships between stimuli to predict values of stimuli which were not directly sampled. This behaviour can be captured by a Bayesian model of the believed value distributions which is updated according to subjects' choices. By including stimuli in the choice task whose spatial positions are unknown, we further demonstrate that novel stimuli can be integrated into existing cognitive maps if their reward statistics are consistent with experienced regularities. Relational knowledge organized in cognitive maps can thus be used to extrapolate across related states and thereby facilitate novel inference. Using fMRI, we investigate the neural dynamics underlying the spread of values across cognitive maps in hippocampal-medial prefrontal networks. We find that a hippocampal cognitive map together with a prefrontal/striatal value representation are combined to enable value inference. Together, our approach opens up the possibility to connect seemingly disparate fields of spatial coding, learning and decision behaviour.

II-21. Neural basis of active odour trail tracking in rodents

Siddharth Jayakumar
Alexander Mathis
Ventkatesh Murthy
Harvard University

SJAYAKUMAR@FAS.HARVARD.EDU
AMATHIS@FAS.HARVARD.EDU
VNMURTHY@FAS.HARVARD.EDU

Active sensing, where sensing is under voluntary control, is extensively used for scanning the environment to extract features of interest from relevant stimuli. Odour-guided navigation is important for rodents as they rely heavily on olfactory cues for finding food, mates and avoiding predators. Navigating odour trails involves bilateral

comparison from the two nares across multiple sniffs, followed by subsequent motor actions. The neural basis for such flexible yet precise behaviour remains poorly understood. Among the olfactory cortical areas, the anterior olfactory nucleus (AON) has privileged access to information coming from both nostrils and sends feedback projections to the olfactory bulb. This circuit architecture suggests that the AON is a key hub, relevant for olfactory guided behaviours, where bilateral odour concentrations are continuously estimated and integrated for correct motor outputs. Here we use an open-loop treadmill with dynamic odour trails that continuously challenge the mice to navigate with high precision to get rewards. We use DeepLabCut to estimate positions of the snout and other body parts with high accuracy (< 0.7 mm on test set). Mice learn this behavior quickly ($>70\%$ rewards collected after 3 days). We observed that while mice predominantly cast while tracking, there exist other modes of following a trail that have not been previously characterized. We corroborated previous work showing that naris occlusion biases the trail tracking behavior. To begin to address the neural basis of how cues are integrated in the brain to modulate motor output, we chose to study the AON as well as the antero-lateral motor cortex (ALM), a hub for motor planning. We find that the AON and the ALM are critical for this behavior, since targeted chemogenetic perturbations led to perturbed trail tracking. We will discuss the relevance of AON for tracking that possibly extend beyond simple bilateral integration of olfactory information.

II-22. EASE: EM-Assisted Source Extraction from calcium imaging data

Pengcheng Zhou¹

Jacob Reimer²

Ding Zhou¹

Amol Pasarkar¹

Ian Kinsella¹

Emmanouil Froudarakis²

Dimitri Yatsenko²

Paul Fahey²

Agnes Bodor³

JoAnn Buchanan³

Dan Bumbarger³

Gayathri Mahalingam³

Russel Torres³

Sven Dorkenwald⁴

Dodam Ih⁴

Kisuk Lee⁴

Ran Lu⁴

Thomas Macrina⁴

Jingpeng Wu⁴

Nuno Costa³

Clay Reid³

Andreas Tolias²

Liam Paninski¹

ZHOUPC1988@GMAIL.COM

REIMER@BCM.EDU

DZ2336@COLUMBIA.EDU

APP2139@COLUMBIA.EDU

IAK2119@COLUMBIA.EDU

FROUDARA@BCM.EDU

DVYATSEN@BCM.EDU

PAUL.FAHEY@BCM.EDU

AGNESB@ALLENINSTITUTE.ORG

JOANNB@ALLENINSTITUTE.ORG

DANB@ALLENINSTITUTE.ORG

GAYATHRIM@ALLENINSTITUTE.ORG

RUSSELL@ALLENINSTITUTE.ORG

SVENMD@PRINCETON.EDU

DODAM.A.IH@GMAIL.COM

KISUK@PRINCETON.EDU

RANL@PRINCETON.EDU

TMACRINA@PRINCETON.EDU

JINGPENG@PRINCETON.EDU

NUNOD@ALLENINSTITUTE.ORG

CLAYR@ALLENINSTITUTE.ORG

ASTOLIAS@BCM.EDU

LIAM@STAT.COLUMBIA.EDU

¹Columbia University

²Baylor College of Medicine

³Allen Institute for Brain Science

⁴Princeton University

Combining two-photon calcium imaging (2PCI) and electron microscopy (EM) provides arguably the most powerful current approach for connecting function to structure in neural circuits. Recent years have seen dramatic advances in obtaining and processing CI and EM data separately. In addition, several joint CI-EM datasets (in which CI was performed in intact tissue, followed by EM reconstruction of the same volume) have been collected. However, no automated analysis tools yet exist that can match each signal extracted from the CI data to a cell segment extracted from EM; previous efforts have been largely manual and focused on analyzing calcium activity

in cell bodies, neglecting potentially rich functional information from axons and dendrites. There are two major roadblocks to solving this matching problem: first, dense EM reconstruction extracts orders of magnitude more segments than are visible in the corresponding CI field of view, and second, due to optical constraints and noisy expression of the calcium indicator in each cell, direct matching of EM and CI spatial components is nontrivial.

In this work we develop a pipeline for fusing CI and densely-reconstructed EM data. We model the observed CI data using a constrained nonnegative matrix factorization (CNMF) framework, in which segments extracted from the EM reconstruction serve to initialize and constrain the spatial components of the matrix factorization. We apply this model to joint CI-EM data from mouse visual cortex and recover hundreds of dendritic components from the CI data (outnumbering somatic components by about 5x), visible across multiple functional scans at different depths, matched with densely-reconstructed three-dimensional neural segments recovered from the EM volume. We use the resulting database to train a stand-alone neural network module to denoise spatial components estimated from 2PCI data.

II-23. Inferring network motifs from neural activity using analytic input-output relation of LIF neurons

Safura Rashid Shomali¹
Majid Nili Ahmadabadi²
Seyyed Nader Rasuli³
Hideaki Shimazaki⁴

SAFURA.RASHID.SHOMALI@GMAIL.COM
MNILI@UT.AC.IR
RASULI@IPM.IR
HIDEAKI.SHIMAZAKI@GMAIL.COM

¹Institute for Research in Fundamental Sciences (IPM)

²University of Tehran

³University of Guilan

⁴Kyoto University

One of the challenges in Neuroscience is to identify network motifs in cortical microcircuitries from recorded neural activity. Having a nonlinear relation between synaptic inputs and output spikes helps us to narrow down possible architecture that can produce the observed activity. Nevertheless, theoretical approaches to this inverse problem remain challenging due to the lack of an analytic input-output relation of a neuron that operates under in vivo conditions. Here we present a method to uncover the network motifs from correlated population activity, using the recently proposed analytical solution of spike timing for a leaky integrate-and-fire neuron model that responds to strong synaptic inputs under noisy, balanced background inputs constituted by many weak synapses. Using this method, we found unique motifs behind the sparse population activity ubiquitously reported for neighboring cortical neurons, which is characterized by their positive pairwise and strong negative triple-wise interactions. We show that two major motifs can express the sparse population activity: (i) inhibitory common inputs to three neurons and (ii) excitatory common inputs to pairs of neurons. By quantitatively comparing model prediction with the activity of macaque V1 neurons [Ohiorhenuan et. al. Nature 2010], we conclude that common inhibition cannot explain the observed sparse activity, but excitatory inputs to pairs well explain it. We quantitatively verify that when the spontaneous activity of postsynaptic neurons is low, as observed in the data, excitatory inputs to pairs can induce strong negative interactions while inhibitory inputs to trios cannot. Accordingly, by considering all possible directional and reciprocal connections among three neurons in 16 motifs, we demonstrate that the motifs of excitatory inputs to pairs are required to induce the strong negative triple-wise interactions. These results indicate that local, excitatory inputs to neuron pairs constitute major cortical motifs that are involved in signal processing of in-vivo cortical microcircuitries.

II-24. Dynamics of natural odor plumes structure the activity of neural populations at the first stage of odor processing in mice

Suzanne Lewis¹
Lai Xu²
Mohammad Tariq¹
Venkatesh Gopal³
Agnese Seminara⁴
Merav Stern¹
David Gire¹

SLEWIS4@UW.EDU
XLCATHERINE@126.COM
TARIQMF@UW.EDU
VGOPAL@ELMHURST.EDU
AGNESE.SEMINARA@UNICE.FR
MS4325@UW.EDU
DHGIRE@UW.EDU

¹University of Washington

²University of Southern California

³Elmhurst College

⁴Universite Cote d'Azur

Although mice can navigate and locate resources using turbulent airborne odor plumes, how they achieve this is poorly understood. One of the main challenges for animals navigating in these regimes is plume dynamics, the stochasticity and intermittency of the fluctuating plume. Population activity at the glomerular level is thought to help process this complex spatial and temporal information, but how plume dynamics impact odor representation in this early stage of the mouse olfactory system is not known. In the past, limitations in odor detection technology have made it impossible to measure plume fluctuations while simultaneously imaging from the mouse's brain. Thus, previous studies have measured the olfactory bulb's response to subsets of features found within olfactory plumes via controlled odor presentations such as odor pulses. Here we measured the olfactory bulb's response to natural olfactory scenes using a miniaturized odor sensor combined with wide-field GCaMP signaling from mitral and tufted cells imaged in olfactory glomeruli in head-fixed mice. We precisely tracked plume dynamics and investigated glomerular responses to this fluctuating input. Manipulating the airspeed of the wind tunnel created characteristically different plumes where increased flow level resulted in increased intermittency. Our results suggest that populations of glomeruli encode odor concentration across plume encounters as demonstrated by a significant correlation between the first principal component of the population mitral and tufted cell response and ethanol sensor amplitude across odor fluctuations (Low airflow: $R = 0.34 \pm 0.06$, Med: 0.43 ± 0.03 , High: 0.48 ± 0.04 , $n=3$ mice). In addition, glomeruli with the most variable responses across plume presentations were best at tracking odor concentration dynamics, suggesting that this response variance is structured by stochastic plume dynamics. Taken together, these data demonstrate that high fidelity encoding of plume dynamics is present at the first stage of odor processing in the mouse olfactory system.

II-25. Activation of distinct cortical circuitries by expected and unexpected stimuli

Farzaneh Najafi¹
 Natalia Orlova¹
 Dmitri Tsybouski²
 Sam Seid¹
 India Kato¹
 Sara Kivikas¹
 Marina Garrett¹
 Peter Groblewski¹
 Rylan Larsen¹
 Doug Ollerenshaw¹
 David Sullivan¹
 Wayne Wakeman¹
 Quinn L'Heureux¹
 Shawn Olsen¹
 Jerome Lecoq¹

FARZNAJ@GMAIL.COM
 NATALIAO@ALLENINSTITUTE.ORG
 TSYBOULSKID@JANELIA.HHMI.ORG
 SAMS@ALLENINSTITUTE.ORG
 INDIA.KATO@ALLENINSTITUTE.ORG
 SARA.KIVIKAS@ALLENINSTITUTE.ORG
 MARINAG@ALLENINSTITUTE.ORG
 PETERG@ALLENINSTITUTE.ORG
 RYLANL@ALLENINSTITUTE.ORG
 DOUGO@ALLENINSTITUTE.ORG
 DAVIDSU@ALLENINSTITUTE.ORG
 WAYNEW@ALLENINSTITUTE.ORG
 QUINN.LHEUREUX@ALLENINSTITUTE.ORG
 SHAWN@ALLENINSTITUTE.ORG
 JEROMEL@ALLENINSTITUTE.ORG

¹Allen Institute for Brain Science
²HHMI Janelia Research Campus

The predictive coding framework for understanding the brain function assumes that the brain contains an internal representation of the world, which is updated by comparing predictions about the external world with the actual sensory. Whether and how the cortex might be involved in computing and updating predictions are still open questions in the field. In particular, it has remained unclear how cortical layers and areas may interact to encode predictions and error signals, and whether distinct neural pathways are recruited following expected vs. unexpected stimuli. Here, we attempted to shed light on these questions by leveraging our recent development of a novel two-photon microscope that allows simultaneous imaging of 8 cortical planes at single cell resolution at 10Hz. We imaged, in behaving mice, 4 cortical layers in areas V1, and LM, a higher order visual area in mouse which is strongly connected to V1. To study the neuronal circuitries that underlie predictive coding, we imaged 3 mouse lines that were tagged for excitatory and inhibitory subpopulations VIP and SST. The behavioral task involved an image-change detection: an image is repeatedly flashed, and after a varying number of presentations, it changes to a different image. The animal is trained to detect the change and report it by licking a waterspout. To investigate the cortical representation of prediction signals, we studied neural responses to images (“expected” events), as well as the following two “unexpected” events: 1) infrequent and random omissions of the visual stimulus; 2) presentation of a novel image set. Here, we demonstrate how different cell types (excitatory neurons and inhibitory neurons of VIP and SST subtypes) in 4 layers of V1 and LM represent prediction signals during visual behavior. Additionally, we have leveraged our simultaneous recordings to study the interactions between cortical columns during expected and unexpected events.

II-26. Precise excitation-inhibition balance - computational implications for feedforward circuits

Aanchal Bhatia¹
 Sahil Moza¹
 Upinder Bhalla²

AANCHALJB@GMAIL.COM
 SAHIL.MOZA@GMAIL.COM
 BHALLA@NCBS.RES.IN

¹National Centre for Biological Sciences
²National Centre for Biological Sciences, TIFR

Excitation-Inhibition (EI) balance is known to be an organizational principle of the neural code. Average proportionality between excitation and inhibition has been observed at the level of the whole brain, circuits, and even at

single neurons. However, it is unknown if EI balance emerges only in response to specific stimuli, or if arbitrary presynaptic input combinations are balanced in a network. To address this, we used optogenetic stimulation of CA3 cells to precisely and combinatorially control the inputs and recorded from EI responses from CA1 neurons in mouse hippocampal slices. Using the first-ever combinatorial analysis of EI balance in the brain, we discovered that hundreds of randomly chosen combinations of presynaptic neurons delivered synaptic inputs with identical I/E ratios at a given postsynaptic neuron, demonstrating a striking degree of precision in EI balance in the hippocampal CA3-CA1 network.

Additionally, we discovered that the EI delay, generally taken to be invariable, decreased with increasing input strength, resulting in a dynamic millisecond-range window for postsynaptic excitation. We show, both with experiments and a biophysical model, that the combination of precise EI balance and dynamic EI delays results in divisive normalization at a subthreshold level. Hence, we show two fundamental control parameters by which the presynaptic network can tune the subthreshold gain and in turn spike probability and timing of a postsynaptic neuron.

Our proposed rules are generally applicable to any precisely balanced feedforward network. Further, precise balance for all inputs is well-positioned to be an input gating mechanism, specifically relevant for CA1 silent cell to place cell conversion. In summary, this study ties together key and apparently disparate network properties of balance, timing, and input gating in a single underlying mechanism, which we characterize experimentally and through models.

II-27. Reward expectancy restructures and enhances spatial encoding in the hippocampus

Seetha Krishnan
Chery Cherian
Mark Sheffield

University of Chicago

SEETHAKRIS@GMAIL.COM
CHERIANCHERY@UCHICAGO.EDU
SHEFFIELD@UCHICAGO.EDU

Place cells in the hippocampus provide an internal representation of the external world. Reward contingencies affect place cell firing; there is evidence that place cells overrepresent rewarded locations, increase their firing during goal approach and encode reward probability. However, in a familiar environment with learned rewarded locations, it is unclear if place cell firing is related to reward per se or the learned reward expectancy of the animal. To determine how reward and reward expectancy might independently influence spatial encoding in the CA1 at the population level, we designed an approach that allowed us to omit or insert rewards and continuously measure reward expectation while maintaining periods of spatial navigation behaviors that were comparable in the presence and absence of reward. By imaging from large populations of CA1 pyramidal neurons in behaving mice traversing virtual linear environments, we show that reward expectancy, but not reward per se, restructures and improves spatial encoding in the hippocampus. Sudden removal of a water reward from a familiar environment caused pre-emptive licking (reward expectancy) to gradually disappear lap-by-lap. Spatial maps initially remained stable during this non-rewarded period until pre-emptive licking completely stopped, at which point spatial maps abruptly restructured into impoverished new maps, with fewer place cells, decreased spatial precision and increased out-of-field firing. Reintroduction of reward re-established pre-emptive licking which also led to an abrupt restructuring of spatial maps, with improved spatial encoding that was distinct from the original rewarded spatial map. These findings reveal that the same physical space can be encoded by distinct spatial maps, based on the animal's internal state of expectation, enabling the memory of distinct contextual episodes. This supports the idea that motivational and emotional states can modulate episodic memory by shaping how external events are encoded by hippocampal populations.

II-28. Contribution of sensory-motor gain to signal and noise in motor control

Seth Egger
 Stephen Lisberger
 Duke University School of Medicine

SETH.EGGER@DUKE.EDU
 LISBERGER@NEURO.DUKE.EDU

Variability in animal behavior is classically attributed to noise in sensory and/or motor systems. However, processing downstream of sensation and upstream of motor coordination such as motor preparation or the control of sensory-motor gain is also subject to noise. Here we combine computational modeling and strategic experimentation to show that noise from the computation of sensory-motor gain contributes significantly to variability of behavior. We start by showing that a model with variation in the trial-by-trial gain shows signal-dependent motor noise, consistent with classical models. However, the gain-noise model further predicts that the rate of increase in variation with signal depends on the mean gain. We test this prediction by measuring smooth pursuit eye movements in response to moving dot patches, a sensory-motor behavior where gain plays a seminal role determining eye speed. We control the mean gain by varying the size of the moving patches. Consistent with the gain-noise model, variation in eye speed increases with its mean at a rate depending on the mean gain. To provide a biologically-relevant account of this finding, we generalize our computational model to a biomimetic circuit constructed from model MT neurons that drive the downstream computation of the gain as well as pursuit behavior. By titrating the noise in the module responsible for computing the gain, we reproduce not only our behavioral results, but also neuron-behavior correlations consistent with observed physiology. These results challenge the typical assumptions of motor noise, suggesting instead that variation downstream of sensation results from the computation of sensory-motor gain.

II-29. Hippocampal replay slows down with experience, as greater detail is incorporated

Alice Berners-Lee^{1,2}
 Ting Feng³
 Delia Silva³
 Xiaojing Wu⁴
 Ellen Ambrose³
 Brad Pfeiffer⁵
 David Foster²

ABERNERSLEE@BERKELEY.EDU
 FTINGSJTU@GMAIL.COM
 S.DELIA.1@GMAIL.COM
 WUXJ03@GMAIL.COM
 AMBROSE.ELLEN@GMAIL.COM
 BRAD.PFEIFFER@UTSOUTHWESTERN.EDU
 DAVIDFOSTER@BERKELEY.EDU

¹Johns Hopkins University

²University of California, Berkeley

³Johns Hopkins University School of Medicine

⁴New York University

⁵UT Southwestern Medical Center

When animals explore a new space, it may be advantageous to quickly get the gist of the spatial layout first, while only incorporating details later. Hippocampal place cells represent an animal's current location as it traverses space. During pauses in running, place cells replay spatial trajectories on a speeded-up timescale, and with a "hover-and-jump" dynamics such that replays jump discontinuously between discrete hover locations, on a timescale consistent with slow-gamma (20-50Hz). By recording simultaneously from large ensembles of place cells as rats explored novel linear tracks, we investigated how the structure of replay changed with experience. Surprisingly, for novel tracks, the duration of replay events increased, and their trajectories slowed down, across laps. Specifically, in early laps, replay events depicted spatial paths traveling rapidly across the track, whereas in later laps, the events took more time to traverse the same physical length. This is the opposite of predictions from several models of replay in which increases in synaptic weight with experience drive faster sequences; however, it is consistent with our speculation that greater details are "filled in" during later laps. Indeed, we determined that the principal change over laps was an increase in the number of hover locations. This raises the ques-

tion of how intermediary locations can be inserted into an already learned sequence. We found evidence for a novel mechanism for this insertion, whereby increased inhibition during later laps masked early-lap associations between spatially distant hover locations (the "gist"), without obscuring stronger associations supporting the updated sequence with the newly inserted locations (the "detail"). finally, we found that populations of prefrontal neurons recorded simultaneously exhibited longer-duration modulation to slower replays, possibly enabling increased computation for non-spatial aspects of experience. These changes in the hippocampal-cortical network during replays suggest that fast and slow replays meet different computational requirements across experience.

II-30. Cellular and synaptic mechanisms of familiarity-evoked theta oscillations in the visual cortex

Alexander Chubykin
Mang Gao
Purdue University

CHUBYKIN@PURDUE.EDU
GAO227@PURDUE.EDU

Recognizing familiarity and novelty in the environment is critical for animal survival. Surprisingly, very little is known about how animals distinguish familiar from novel stimuli. We have recently discovered persistent stimulus-triggered theta oscillations in the visual cortex of mice, which are specific to the familiarity and spatial frequency features of the stimulus. We have been studying these theta oscillations using a comprehensive approach involving *in vivo* extracellular recordings using 64-channel silicon probes, patch-clamp recordings in awake head-fixed mice and pupillometry. Using extracellular and patch-clamp recordings, we have discovered distinct neuronal ensembles active at different cycles of the oscillation. Interestingly, familiarity-evoked oscillations lead to the decrease in suprathreshold responses to the non-preferred directions of the drifting sinusoidal gratings leading to the improved direction selectivity in V1 neurons. Using a special pipette holder Optopatcher to deliver blue laser light directly into the patch pipette, we were able to use optogenetics to directly measure the synaptic strength of the light-triggered EPSCs in V1 neurons in head-fixed mice. Using this method, we have discovered the weakening of the thalamocortical projections coinciding with the strengthening of the intracortical projections in V1 following visual experience and emergence of the familiarity-evoked theta oscillations. finally, we performed simultaneous recordings from V1 and two visual thalamic nuclei dLGN and LP, as well as the retrosplenial cortex (RSC) and did not detect theta oscillations in those areas, suggesting that the theta oscillations may originate in V1.

II-31. The dentate gyrus classifies cortical representations of learned stimuli

Fabio Stefanini¹
Nicholas Woods²
Daniel Apodaca-Montano²
Jeremy S Biane²
Mazen Kheirbek²

FABIOEDOARDOLUIGIALBERTO@GMAIL.COM
NICHOLAS.WOODS@UCSF.EDU
DANIEL.APODACA@UCSF.EDU
JEREMY.BIANE@UCSF.EDU
MAZEN.KHEIRBEK@UCSF.EDU

¹Columbia University

²University of California San Francisco

Animals must discern important stimuli and place them onto their cognitive map of their environment. The neocortex conveys general representations of sensory events to the hippocampus, however, how the hippocampus classifies and sharpens the distinctions between important stimuli remains unclear. Here, we monitored the activity of dentate gyrus granule cells (DG GCs) and lateral entorhinal cortex (LEC) neurons across days to understand how sensory representations are modified by experience. We found neural representations of olfactory stimuli in DG GCs that were directly related to future olfactory learning. In addition, with learning, LEC neurons maintained stable representations of the odors as opposed to DG GCs, which strongly changed their responses to odor stimuli and responded more to the conditioned and less to the unconditioned odorant. Thus, with learning, DG GCs

amplify the cortical representations of important stimuli, which may facilitate information storage and recall and guide appropriate behaviors.

II-32. Accurate extraction of membrane potential for in vivo voltage imaging

Michael Xie¹
Yoav Adam²
Linlin Fan¹
Urs Boehm¹
Ian Kinsella³
Ding Zhou³
Liam Paninski³
Adam Cohen¹

M_XIE@COLLEGE.HARVARD.EDU
YOAV.ADAM@MAIL.HUJI.AC.IL
LINLINFAN@G.HARVARD.EDU
ULBOEHM@FAS.HARVARD.EDU
IAN.KINSELLA@COLUMBIA.EDU
DZ2336@COLUMBIA.EDU
LIAM@STAT.COLUMBIA.EDU
COHEN@CHEMISTRY.HARVARD.EDU

¹Harvard University

²Hebrew University of Jerusalem

³Columbia University

The ability to probe simultaneously membrane potential of multiple neurons in the brain would have a profound impact on neuroscience research. Genetically encoded voltage indicators (GEVIs) are a promising tool for this purpose, and recent developments have achieved high signal to noise ratio in vivo with 1-photon imaging. These recordings present new challenges for analysis. In the live brain, out-of-focus fluorescence and motion artifacts likely contaminate low-amplitude subthreshold signals from neurons in the imaging plane. The activity of background cells is often correlated with the true voltage signals of cells in focus, for example due to shared inputs. These correlated dynamics confound signal extraction techniques that assume statistical independence between signal and background.

Here we present a novel signal extraction methodology, Spike-Guided PMD-NMF, which resolves both supra- and sub-threshold behaviors with high accuracy, preserves inter-neuron correlations, and rejects correlated background. The method uses the facts that neuron spiking is less correlated than subthreshold voltages and that out-of-focus signals are spatially smoother than in-focus signals.

The pipeline was validated using simulated datasets and synthetic movies composed by combining overlapping single-cell recordings. Under both conditions, the methods produced signals that were highly correlated with the ground truth data and showed similar correlation matrices to the known inputs. We applied the pipeline to in vivo recordings from zebrafish spinal cord, mouse hippocampus, and mouse cortex. In all cases, we observed that the pipeline automatically identified spiking cells and separated their signals from correlated background. Our tool enables the use of voltage imaging technology for exploration of subthreshold membrane potential dynamics in live animals.

II-33. Anterior and posterior striatum play distinct roles in evidence accumulation

Adrian Bondy
Diksha Gupta
Carlos D Brody
Thomas Luo
Princeton University

ADRIAN.BONDY@GMAIL.COM
DIKSHAG@PRINCETON.EDU
BRODY@PRINCETON.EDU
THOMAS.ZHIHAO.LUO@GMAIL.COM

The ability to use sensory information to guide decisions—and to accumulate that information over time—is a critical cognitive ability. In rats, various studies show that inactivation of two distinct subregions of the dorsal

striatum, the anterior dorsal striatum (ADS) and the posterior tail of the striatum (TS), impair performance in tasks requiring accumulation of auditory evidence, suggesting that multiple striatal pathways may be involved. ADS neurons, and those in a homologous region in primates, have “ramping” firing rates that increase or decrease at a rate determined by the strength of evidence. However, analogous recordings have not been performed elsewhere in striatum. We performed multi-electrode array recordings at multiple striatal sites in rats performing a pulsatile auditory evidence accumulation task. To describe neuronal encoding, we used a generalized linear model (GLM), which describes neuronal firing rates as a combination of temporal kernels corresponding to different task events (e.g. individual evidence pulses, choice, reward, etc.). This revealed a pronounced difference in encoding along the striatal anteroposterior axis. First, it showed strong choice-modulation both before and during the execution of the choice in ADS neurons but not TS neurons. Furthermore, the kernels for evidence pulses were significantly weaker, but more long-lasting, in ADS neurons. These longer-lasting responses can sum over time, offering a mechanism for the “ramping” activity previously reported in ADS. By contrast, the results suggest TS neuron activity is more consistent with a representation of the instantaneous sensory evidence. These results establish a diversity of timescales in the encoding of sensory evidence along the anteroposterior axis of striatum and clarify how information flows through parallel basal ganglia circuits during decision making.

II-34. Hebbian STDP protects against assembly fusion in stochastic networks

Xinruo Yang
Brent Doiron

XIY84@PITT.EDU
BDOIRON@PITT.EDU

University of Pittsburgh

Hebbian spike time dependent plasticity (STDP) is often associated with the reinforcement of causal synaptic interactions within a network. Most famously this occurs in the development of feedforward architectures arising from sequential activation of neuron chains (Gerstner et al. 1993; fiete et al. 2010). However, how STDP rules support the development and stability of recurrent assembly structure (clustered wiring) is less understood, despite such wiring being commonly associated with cortical and hippocampal circuits (Harris and Mirsc-flogels, 2013). Ocker and Doiron (2018) developed a mean field theory of synaptic learning in spiking networks and showed how STDP rules with a balance between potentiation and depression allow the correlated spiking activity occurring post-training to reinforce any learned assembly structure. However, they did not consider the case where a subset of neurons can be members of both assemblies, i.e. so that the assembly structure has overlap. It is easy to imagine that with sufficient overlap the dynamics of the overlap neurons could merge initially segregated assemblies into a single fused assembly, effectively erasing any learned structure. We extend the theory in Ocker and Doiron (2018) to include assembly overlap and investigate how the shape of the learning rule determines whether assemblies will remain segregated or fuse together after training. We show that a Hebbian (anti-symmetric) STDP rule suppresses the fusion dynamics due to overlap projections. However, adding a strong symmetric component to the STDP rule, effectively diluting its causal structure, allows overlap projections to fuse assemblies. In sum, we show that the casual structure of Hebbian STDP protects against circuit dynamics that would otherwise conspire to degrade learned assembly structure.

II-35. The role of hippocampal sequences in deliberative control: a computational and algorithmic theory

David Theurel¹
Matthew Wilson²
Zhe Chen³

THEUREL@MIT.EDU
MWILSON@MIT.EDU
ZHE.CHEN@NYULANGONE.ORG

¹Massachusetts Institute of Technology

²Massachusetts Institute of Tehcnology

³New York University

Building upon current theories, we advance a multi-level quantitative theory of HPC function, embedded in a working model of the whole deliberative decision-making system. In the system-level model: mPFC and HPC participate in a hierarchical mapping and planning scheme, in which mPFC flexibly computes high-level plans that are relevant to the task at hand, and then prompts dHPC to refine each step of the plan as needed. The task of dHPC is formalized as an optimal control problem. By asking how to solve this task efficiently and flexibly, we arrive at an algorithmic theory in terms of Pontryagin's maximum principle. We find specific roles for SWR sequences and theta sequences in solving the ensuing system of differential equations. This framework parsimoniously incorporates a wide and diverse array of observations: the complementary roles of PFC and HPC in enabling flexible behavior and multiple forms of memory; the various forms of coherent neural activity between PFC and HPC; the transition from vicarious trial and error to ballistic behavior over the course of learning; and the effects of learning, of varying reward contingencies, and of fear conditioning on HPC sequences. It explains the functional role of the theta rhythm for the deliberative system; the complementary roles of forward and reverse SWR sequences and theta sequences; and the puzzle of mixed SWR sequences. The framework also makes a number of novel predictions. Its strongest prediction so far consists in a simple equation which must be satisfied by the individual trajectories of all HPC sequences. On the linear track this equation suffices to completely determine each trajectory up to a single number. We are currently testing this prediction against experimental data. The success of this framework would mean a novel, unifying and quantitative paradigm for the deliberative system.

II-36. Adaptive integration behavior from static integration windows

Richard McWalter
Josh McDermott

MCWALTER@MIT.EDU
JHM@MIT.EDU

Massachusetts Institute of Technology

Adaptive behavior is common in biological systems, which often adjust performance characteristics to the local environment. Such adaptive characteristics are evident in temporal integration, which in several domains scales with stimulus variability, plausibly to average over longer extents when necessary to obtain stable estimates. For instance, sound texture statistics appear to be averaged with an integration window that scales with the input, with more variable textures averaged over longer extents. However, many different statistics are needed to account for texture perception, raising the possibility that different statistics might have different integration time scales, potentially combining to yield apparent adaptation to input variability. To test this idea, we measured integration windows for individual classes of statistics from a standard auditory texture model. In a psychophysical experiment, listeners judged which of two textures was most similar to a reference texture, with textures synthesized to vary in one class of statistic. We measured performance for different stimulus durations. In all cases, discrimination improved with stimulus duration but then leveled off, presumably signaling the extent of the averaging window used to estimate the statistics mediating the decision. However, the performance leveled off at different durations for different statistics, ranging from 200 milliseconds (for the cochlear envelope mean) to a few seconds (slow modulation power). We then incorporated fixed but class-specific integration windows (derived from the psychophysical results) into an observer model that operated on the experiment stimuli. When tested with stimuli in which all statistics varied, model performance, like that of humans, leveled off at longer durations for more variable textures (because statistics with longer integration times limited discrimination). The results suggest the presence of multiple static integration windows that collectively produce different performance characteristics in different stimulus regimes. Apparent adaptive behavior can thus occur without actual adaptation within a sensory system.

II-37. Resonating neurons stabilize heterogeneous grid-cell networks

Divyansh Mittal¹
Rishikesh Narayanan²

DIVYANSH@IISC.AC.IN
RISHI@IISC.AC.IN

¹Indian Institute of Science, Bangalore

²Indian Institute of Science

Grid cells in the medial entorhinal cortex have multiple firing fields organized in a hexagonal pattern. The continuous attractor network (CAN) model is one of several models that have been employed to understand grid-cell firing (Burak and Fiete, 2009). CAN models for grid-cell firing typically consider homogeneous neuronal models that act as integrators of information, which neither account for the several biological heterogeneities nor consider individual neurons as resonators exhibiting theta-frequency (3–10 Hz) resonance. In this rate-based modeling study, we assessed grid-cell emergence within the CAN framework in the presence of three distinct forms of neural heterogeneities and in scenarios where individual neurons act as resonators. First, in an integrator network, we demonstrate that introduction of heterogeneities in the integration time constant (intrinsic property) or in velocity coupling (afferent input) or in local connectivity (synaptic strength) hampered the emergence of grid cell firing. Quantitatively, we show that the grid score associated with the emergent firing pattern reduced with progressive increase in the degree of each type of heterogeneity, with synaptic heterogeneity playing a dominant role. Second, in a homogeneous network of neurons, we replaced all neurons by resonators using an abstraction that introduced an additional high pass filter to together yield resonance in the theta-frequency range. We found that an increase in neuronal resonance frequency resulted in reductions in grid spacing and grid-field size. Finally, we demonstrate that networks endowed with resonating neurons were resilient to the incorporation of heterogeneities, and exhibited well-defined grid fields even in the combined presence of the highest degree of all three heterogeneities. Together, our study unveils the ability of resonance in individual neurons to regulate grid-cell spacing and to stabilize grid-cell firing in the face of distinct forms of network heterogeneities.

II-38. Non-perturbative renormalization group analysis of strongly-coupled spiking networks

Braden Brinkman
Stony Brook University

BRADEN.BRINKMAN@STONYBROOK.EDU

To fully explain how neural dynamics transmit information and perform computations, we need to understand the structure of the coordinated activity of neurons and their responses to external inputs. Given a model of neural dynamics and their synaptic connections, we would in principle achieve this goal by calculating the statistical and response functions of the network—a notoriously intractable task for all but the simplest models. While several approximations (e.g., mean-field theories, linear response theories, and diagrammatic series) have been successfully applied to many network models, they are “perturbative” in that they are often limited to networks with relatively weak synaptic connections or fluctuations. Networks with strong nonlinear behavior often invalidate these approaches’ predictions, demanding new approximation methods.

The goal of this work is to develop such methods and apply them to strongly coupled spiking networks. We adapt “non-perturbative renormalization group” methods from statistical physics and apply them to soft-threshold leaky integrate-and-fire networks. We show that the true mean firing rates of the network satisfy a nonlinear system of equations formally similar to the mean-field system but with a different effective nonlinearity. We explicitly derive a differential equation for this nonlinearity and solve it numerically. The structure of the equation itself reveals that synaptic connections with identical eigenvalue distributions yield identical nonlinearities. Our results predict the distribution of firing rates in simulated networks of neurons reasonably well, even in strong coupling regimes in which perturbative calculations begin to break down.

Most importantly, the non-perturbative character of these approximation methods renders it possible to study critical phenomena and identify phases of collective behavior in strongly coupled networks. Continued work with

this methodology will address these questions, and in particular focus on elucidating the relationship between the structure of synaptic connections and phases of collective activity.

II-39. Interneuron diversity emerges in spiking networks optimized for a compartment-specific excitation-inhibition balance

Joram Keijsers
Sprekeler Henning
TU Berlin

JORAMKEIJSER@GMAIL.COM
H.SPREKELER@TU-BERLIN.DE

Neural circuits consist of excitatory projection neurons and inhibitory interneurons (INs). In recent years, it has become clear that these INs are highly diverse in their morphological, electrophysiological, and molecular properties (Tremblay et al. Neuron 2016). However, the functional role of this diversity is not well understood. Here we show that several aspects of IN diversity can be explained from the objective of balancing excitation (E) and inhibition (I) across different cellular compartments of pyramidal cells.

To this end, we simulated a network consisting of interneurons and 2-compartment pyramidal neurons. These pyramidal neurons use spikes and bursts to represent input signals targeting their somatic and dendritic compartments, as recently suggested (Naud et al. PNAS 2018). We leveraged recent advances in training spiking networks (Neftci et al. arXiv 2019) to optimise the network connectivity and short-term plasticity for a compartment-specific E/I balance. Before the optimisation, the interneurons form a homogeneous population. After the optimisation, the majority of the interneurons falls within one of two classes. One class receives short-term depressing inputs from pyramidal neurons and targets their soma, consistent with the properties of parvalbumin (PV) expressing INs. In contrast, the other class receives short-term facilitating inputs from pyramidal neurons and targets their dendrites, consistent with somatostatin (SOM) expressing INs.

This suggests that the division of labour between PVs and SOMs, and their connectivity and short term plasticity emerges from the need to balance excitation and inhibition across cellular compartments. We suggest that the optimisation approach we employed may provide a general framework to link interneuron diversity to functional requirements.

II-40. The influence of noise on the efficient coding of natural scenes

Na Young Jun
Greg field
John Pearson
Duke University

NAYOUNG.JUN@DUKE.EDU
FIELD@NEURO.DUKE.EDU
JOHN.PEARSON@DUKE.EDU

In the field of neuroscience, an increasing number of neuron types are discovered every year in various regions in the brain, including the retina. How should the nervous system spatially arrange receptive fields across diverse cell types? The mammalian retina processes and encodes an enormous amount of information from light entering the eye, converting visual scenes into electrical signals that are sent to the brain. To accomplish this encoding, the retina consists of diverse retinal ganglion cell (RGC) types that extract and process distinct visual features in parallel, such as contrast, motion, and color. For each RGC type, the receptive fields tile visual space to form a mosaic-like pattern. A recent study[1] indicates that mosaics of receptive fields are organized with respect to each other. For example, ON and OFF brisk transient cells are anti-aligned, meaning that the receptive fields of ON mosaics are situated between, rather than on top of, receptive fields for the OFF mosaics, and these relationships are conserved across diverse species (e.g. mice, rats, and primates)

What determines this inter-mosaic relationship? We suggest that the efficient coding principle[2] can predict these statistical regularities in inter-mosaic organization. Using a previously proposed model of efficient coding by

RGCs, we calculate optimal spatial arrangements of receptive fields and find that inter-mosaic patterns (aligned vs. anti-aligned) depend on the levels of input noise (to the retina) and output noise (from RGCs). In addition, we show that there exists a phase transformation between the two mosaic alignments as a function of the distribution of stimuli and these noise parameters. We compare these predictions to data and show they are in agreement. These results generate a theoretical understanding of the highly conserved organizations of retinal cell types and their mosaic patterns across species.

II-41. Visual pattern analysis by motor neurons

Antimo Buonocore

Matthias Baumann

Ziad Hafed

University of Tuebingen

ANTIMO.BUONOCORE@CIN.UNI-TUEBINGEN.DE

MATTHIAS-PHILIPP.BAUMANN@STUDENT.UNI-TUEBINGEN.DE

ZIAD.M.HAFED@CIN.UNI-TUEBINGEN.DE

Coordination between specialized brain areas is a necessary prerequisite for optimal adaptive behavior. Such coordination entails multiple interconnected pathways not only for transducing sensory inputs, but also for manipulating internal state and driving motor outputs. Perhaps counterintuitively, supporting such coordination may require endowing certain brain areas with access to information that appears otherwise outside of their expected functional repertoire. We show a striking example of this in the final oculomotor common path of the primate brainstem. We recorded omnipause neurons (OPN's) in the nucleus raphe interpositus (rip), constituting the very final gateway for allowing or preventing saccades (Sparks, 2002). We hypothesized that behavioral robustness means granting OPN's privileged access to early visual information, therefore maintaining flexibility to sense and react to the outside world right at the very last stage of motor control. We presented stimuli of different features (spatial frequency, contrast, orientation, and motion) during fixation. OPN's, normally tonic and only pausing to allow saccade execution, showed robust early (<50 ms) phasic visual responses, which were also feature-tuned. Consistent with another motor structure, superior colliculus (SC) (Chen et al., 2018), OPN's preferred low spatial frequencies. To directly compare OPN and SC responses, we also replicated (Chen et al., 2018) and found that OPN visual activity was as early as, or even earlier, than in SC. What is the functional implication of early visual flow to such two late motor areas? In this case, it is to differentiate the possible motor outcomes. When we injected short microstimulation pulse trains in OPN's, to "simulate" brief phasic visual responses, we completely inhibited saccades. Contrarily, the same trains in SC triggered saccades. Therefore, we uncovered a sensory race between motor areas, providing a highly mechanistic description of exactly what to expect in terms of two important aspects of adaptive behavior: stopping versus orienting.

II-42. 3-factor Hebbian learning rules in deep networks: an information bottleneck approach

Roman Pogodin

Peter Latham

Gatsby Computational Neuroscience Unit, University College London

RMN.POGODIN@GMAIL.COM

PEL@GATSBY.UCL.AC.UK

How can a baby learn to recognize objects? Such learning would involve plasticity in the visual system, but the rules guiding it are yet unknown. The deep learning community's solution for image recognition, backpropagation, is very efficient but biologically implausible: it requires neurons to know the weights in subsequent layers (the weight transport problem); it switches between the forward pass (computation) and the backward pass (learning); and it relies on precise labels (for classification). While feedback alignment, an approximation to backpropagation, solves the weight transport problem [Akrouf et al., 2019], it still needs precise labels and a separate backward pass implemented in a parallel network. As an alternative, 3-factor Hebbian rules are less efficient than feedback alignment but require only pre- and post-synaptic firing rates (without a backward pass) and a global signal (e.g.,

a neuromodulator that indicates the error) [Gerstner et.al., 2018]. We show that our form of 3-factor Hebbian learning is an instance of the information bottleneck approach [Ma et.al., 2019], in which networks learn to balance compression of the inputs (e.g., images) with correct prediction of the outputs (e.g., their labels). We also generalize the Hebbian update to more complicated forms of pre- and post-synaptic interaction such that they gain more learning power, but remain biologically plausible and do not need precise labels. This view allows us to build synaptic plasticity rules based on the task the network should perform and the amount of information available to a neuron; it also predicts what shape the third, neuromodulatory, factor should take for efficient learning.

II-43. Temporal basis decomposition reveals a linear relationship between neural activity and kinematics

David Sabatini
Matthew Kaufman

DAVIDASABATINI@GMAIL.COM
MATTKAUFMAN@UCHICAGO.EDU

The University of Chicago

Efforts to identify low-dimensional structure in neural activity have largely focused on finding basis sets consisting of instantaneous patterns of activity across neurons. This choice fits with the idea that neural activity has an instantaneous relationship with variables of interest, such as the kinematics of ongoing movement. Here, we test an alternative: identifying a basis set of activity patterns over time. This decomposition follows from the idea that neural circuits may have conserved, time-varying modes of activity, such as the coordinated, slow oscillations in firing rates found by Churchland et al. (2012). We tested this procedure on electrophysiology data from ~200 simultaneously-recorded neurons from monkeys performing 108 different straight and curved reaches. After performing this decomposition, we can analyze the loadings for these temporal components, which summarize how strongly each neuron participates in each temporal mode for a given reach. Using ordinary regression, we built linear maps from predictor variables to our loading matrices, and asked how well we could reconstruct the neural data. With nine temporal modes and this simple linear mapping, we could reconstruct ~85% of the movement-epoch neural activity from the neural preparatory activity. This high-quality reconstruction is consistent with preparatory activity determining how strongly each of several time-varying network modes is recruited. We could also relate the loadings to kinematics. Again, with simple regression against the loadings, given a reach's trajectory or the corresponding muscle activity, we could reconstruct 70-80% of the trial-averaged movement-epoch neural activity. Finally, we could invert this map and obtain single-trial estimates of kinematics. This procedure could reconstruct 85-90% of the variance in hand position over time, even during curved reaches. Temporal basis decomposition thus reveals a hidden, linear relationship between neural activity during movement and preparatory neural activity, muscle activity, and kinematics.

II-44. Experience dependent context discrimination in the hippocampus

Mark Plitt
Lisa Giocomo
Stanford University

MARKPLITT@GMAIL.COM
GIOCOMO@GMAIL.COM

The hippocampus is a medial temporal lobe brain structure that contains neural representations and circuitry capable of supporting declarative memory. Hippocampal place cells fire in one or few restricted spatial locations in a given environment. Between environmental contexts, place cell firing fields remap (turning on/off or moving to a new spatial location), providing a unique population-wide neural code for context specificity. However, the manner by which features associated with a given context combine to drive place cell remapping remains a matter of debate. Here we show that remapping of neural representations in region CA1 of the hippocampus is strongly driven by prior beliefs about the frequency of certain contexts, and that remapping is equivalent to an optimal estimate of the identity of the current context under that prior. This prior-driven remapping is learned early in

training and remains robust to changes in the behavioral task-demands. Furthermore, a simple auto-associative learning model is sufficient to reproduce these results. Our findings demonstrate that place cell remapping is a generalization of representing an animal's location. Rather than simply representing location in physical space, the hippocampus represents an optimal estimate of location in a multi-dimensional stimulus space.

II-45. Parallel emergence of mixed tuning to task variables and cell assembly coordination in mPFC during learning

Aleksander Domanski¹
Daniel Durstewitz²
Matt Jones¹

ALEKS.DOMANSKI@BRISTOL.AC.UK
DANIEL.DURSTEWITZ@ZI-MANNHEIM.DE
MATT.JONES@BRISTOL.AC.UK

¹University of Bristol

²Central Institute for Mental Health, Mannheim, Germany

The medial prefrontal cortex (mPFC) is a critical contributor to execution of complex behaviours, demonstrating remarkable flexibility in its neural dynamics during cognition. For example, during context-guided decision making by well-trained animals, neurons coordinate their firing into transiently activating cell assemblies. Concurrently, mPFC neuron firing displays non-linear mixed selectivity across task variables, facilitating decision-making based on high-dimensional representations of working memory items. It has been suggested that the mixed information tuning of mPFC neurons emerges through Hebbian learning in randomly wired networks [5] and that cell assembly dynamics of trained networks are a permissive substrate for nonlinear computation. Given this shared reliance on Hebbian plasticity, we should expect to see interrelated changes in both selectivity and assemblies over learning. However, this hypothesis remains largely untested; mixed selectivity and cell assembly dynamics have not been conclusively investigated in the same experimental preparation or evaluated across extended learning.

Here, we analysed two mPFC multiple single unit datasets from rats learning over days to perform Delayed Non-Match To Sample (DNMTS) or Serial Spatial Alternation (SSA) spatial working memory tasks. Under both paradigms, while assembly membership remained similarly sparse, mixed selectivity for contextual and spatial task variables increased during learning, mirroring increased behavioural decoding capacity. In the DNMTS dataset, mixed selectivity was strongest in mPFC neurons forming distributed cell assemblies with similarly-tuned hippocampal neurons. In the SSA dataset, mixed-selective neurons emerged and evolved with training. Cell assembly coordination was initially sparsely distributed across millimeters of cortical tissue, but shifted from a 1/distance dependence to increasingly link neurons separated by <2mm.

Together, these results firmly implicate both emerging mixed selectivity and the coordination of distributed cell assemblies in mPFC and suggest a synergistic role in learning to perform flexible computation.

II-46. The interplay between randomness and structure during learning in RNNs

Friedrich Schuessler¹
Alexis Dubreuil²
Francesca Mastrogiuseppe³
Srdjan Ostojic²
Omri Barak¹

FRSCHU@POSTEO.DE
ALEXIS.DUBREUIL@GMAIL.COM
FRAN.MASTROGIUSEPPE@GMAIL.COM
SRDJAN.OSTOJIC@ENS.FR
OMRI.BARAK@GMAIL.COM

¹Technion

²Ecole Normale Supérieure

³Gatsby Computational Neuroscience Unit, University College London

Learning in the brain always takes place on the background of existing "random" connectivity. Which role such

initial connectivity plays and how newly formed structure relates to the task learned are largely unsolved questions. One candidate structure is low-rank connectivity, which is the basis for many designed networks, and a tool for restricted network training (in the field of reservoir computing).

Here, we show that low-rank structure in the connectivity emerges generically when training networks to perform simple tasks. We numerically analyze artificial recurrent neural networks in which the entire connectivity is trained with gradient descent to perform a number of systems neuroscience tasks, and find that the newly formed structure is of low rank. Using a simplified task, we describe analytically the origin of this phenomenon. Specifically, we show how low-rank structures emerge sequentially during different phases of training. Finally, we show that the initial random connectivity accelerates learning, and uncover the underlying mechanisms.

Altogether, our study opens a new viewpoint on the synergistic roles of randomness and structure in task-performing neural networks.

II-47. Meta-learning Hebbian plasticity for continual familiarity detection

Danil Tyulmankov
Guangyu Robert Yang
Larry Abbott
Columbia University

DT2586@COLUMBIA.EDU
GYANG.NEURO@GMAIL.COM
LFA2103@COLUMBIA.EDU

Memories are stored and recalled throughout the lifetime of an animal, but many models of memory, including previous models of familiarity detection, do not operate in a continuous manner. We consider a family of models that recognize previously experienced stimuli and, importantly, operate and learn continuously. Specifically, we investigate a learning paradigm in which stimuli are presented in a streaming fashion with repetitions at various intervals, and the subject/model must report whether the current stimulus has previously appeared in the stream. We propose a feedforward network architecture with ongoing plasticity in the synaptic weight matrix. Parameters governing plasticity and static network parameters are meta-learned using gradient descent to optimize the continual familiarity detection process. This architecture, unlike recurrent networks without ongoing plasticity, generalizes easily over a range of repeat intervals even if trained with a single interval. We show that an anti-Hebbian plasticity rule (co-activated neurons cause synaptic depression) enables repeat detection over much longer intervals than a Hebbian one, and this is the solution most readily found by meta-learning. This rule leads to experimentally observed features such as repeat suppression in the hidden layer neurons. In contrast to previous theoretical work, the capacity of these networks remains constant across their lifetimes, meaning that pairs of stimuli with a given temporal separation are stored and recognized as familiar independent of the network's input history. We also consider learning rules that use an external gating circuit to control plasticity. Collectively, these models demonstrate a range of different psychometric curves that we compare to human performance.

II-48. Uncovering the organization of neural circuits with generalized phase-locking analysis

Shervin Safavi¹
Fanis Panagiotaropoulos²
Vishal Kapoor¹
Juan F Ramirez-Villegas³
Nikos Logothetis¹
Michel Besserve¹

RESEARCH@SHERVINSAFAVI.ORG
FANISPA@GMAIL.COM
VISHAL.KAPOOR@TUEBINGEN.MPG.DE
JUANFELIPE.RAMIREZVILLEGAS@IST.AC.AT
NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE
MICHEL.BESSERVE@TUEBINGEN.MPG.DE

¹Max Planck Institute for Biological Cybernetics

²NeuroSpin

³Institute of Science and Technology

Many neurophysiological recordings provide concurrent signals of two different nature: On the one hand, the time-stamps of action potentials reflect the information sent by individual neurons, and on the other hand, Local field Potentials (LFP) reflect multiple sub-threshold and postsynaptic mechanisms related to the underlying network activity. The synchronization between spiking activity and the phase of particular LFP rhythms has been used as an important marker to reason about the underlying cooperative network mechanisms. In order to extract in a systematic way coupling information from the largely multivariate data available in current recording techniques, we study a multivariate extension of spike-field coupling analysis. After whitening band-passed LFP signals, we collect normalized pairwise complex-value spike-field coupling coefficients in a rectangular matrix and summarize its structure with the largest singular value and the corresponding singular vectors. Singular vectors represent the dominant LFP and spiking patterns and the singular value, called generalized Phase Locking Value (gPLV), characterizes the strength of the coupling between LFP and spike patterns. We further investigate the statistical properties of the gPLV and develop an empirical and theoretical statistical testing framework. We apply the method to various simulated and experimental datasets. first, GPLA's performance is superior to univariate measures in the presence of large amounts of noise. Next, application of GPLA on simulations of hippocampal sharp-wave-ripples (SWR) reveals various characteristics of hippocampal circuitry (e.g. communication flow from CA3 to CA1 during the SWR) with minimal prior knowledge. Furthermore, application of GPLA on Utah array recordings in anesthetized macaque suggests a non-trivial coupling between spiking activity and LFP traveling wave in the ventrolateral Prefrontal Cortex (vIPFC). In summary, with GPLA, we can quantify, characterize and statistically assess the interactions between population spiking activity and mesoscopic network dynamics.

II-49. Impact of single neurons in cortical circuits with dense and strong connectivity

Juan Luis Riquelme
Mike Hemberger
Gilles Laurent
Julijana Gjorgjieva

JLUIS.RIQUELME@GMAIL.COM
MIKE.HEMBERGER@BRAIN.MPG.DE
GILLES.LAURENT@BRAIN.MPG.DE
GJORGJIEVA@BRAIN.MPG.DE

Max Planck Institute for Brain Research

Experimental studies have shown that individual neurons can impact the activity of their local network. Recent evidence in the reptilian cortex shows that just one spike can trigger repeatable cascades of activity in the surrounding neurons. Similar sequences of neuronal activations have been reported in other systems, such as mammalian cortex and the birdsong system, but we still lack a good mechanistic understanding. To study how single spikes impact local activity, we built a computational model of the three-layered turtle cortex. We found that the model, fitted to experimental data and with minimal assumptions about the connectivity, was able to reproduce sequences from single spikes comparable to those observed *ex vivo*. The model included rare but powerful synapses from a long-tailed distribution as reported experimentally. By simulating alternative models where the distribution is truncated, we showed that these strong connections are the main conduit through which individual excitatory neurons most reliably influence their local network. Activity in the responding network evolves by spreading away from the activated neuron and is sensitive to the global level of activity. Contrary to excitatory cells, we found that inhibitory neurons quickly lose the temporal precision of their response when background network activity increased. Therefore, we predict the presence of repeatable sequences of excitatory but not inhibitory spikes *in vivo*. finally, by studying the connections most frequently traversed and quantifying the motifs that are present, we describe how these strong connections are supported by an abundance of weak ones. Our work zooms into the fine network interactions triggered by a single spike of a single neuron in cortex, and explains how repeatable and temporally precise sequences of activity can be generated, with ample implications for information transmission in cortical networks.

II-50. Biologically plausible supervised learning via similarity matching

Shanshan Qin
Cengiz Pehlevan
Harvard University

SSQIN@SEAS.HARVARD.EDU
CPEHLEVAN@SEAS.HARVARD.EDU

End-to-end training of neural networks by optimizing an objective function is an extremely powerful idea, prompting the question whether the brain is using the same strategy. A main challenge for this proposal has been finding a biologically-plausible solution to the so-called “credit assignment” problem. We propose a novel solution, being motivated by observations in the ventral visual pathway and trained deep networks. In both, representations of objects in the same category become progressively more similar, while objects belonging to different categories becomes less similar. We use this observation as a local-to-layer learning goal in a deep network: each layer aims to learn a representational similarity matrix that interpolates between previous and later layers. We formulate this idea using a supervised deep similarity matching cost function and derive from it deep neural networks with feedforward, lateral and feedback connections, and neurons that exhibit biologically-plausible Hebbian and anti-Hebbian plasticity. We illustrate the plausibility of our approach in linear regression tasks.

II-51. Disentangling the roles of dimensionality and cell classes in neural computations

Alexis Dubreuil¹
Adrian Valente¹
Francesca Mastrogiuseppe²
Srdjan Ostojic¹

ALEXIS.DUBREUIL@GMAIL.COM
ADRIAN.VALENTE@ENS.FR
FRAN.MASTROGIUSEPPE@GMAIL.COM
SRDJAN.OSTOJIC@ENS.FR

¹Ecole Normale Supérieure

²Gatsby Computational Neuroscience Unit, University College London

The description of neural computations currently relies on two competing views: (i) a classical single-cell view that relates the activity of individual neurons to sensory or behavioural variables, and focuses on how different cell classes map onto computations; (ii) a more recent population view that instead characterises computations in terms of collective neural trajectories, and focuses on the dimensionality of these trajectories as animals perform tasks. How the two key concepts of cell classes and low-dimensional trajectories interact to shape neural computations is however at present not understood. Here we address this question by combining machine-learning tools for training RNNs with reverse-engineering and theoretical analyses of network dynamics. We introduce a novel class of theoretically tractable recurrent networks: low-rank, mixture of Gaussian RNNs. In these networks, the rank of the connectivity controls the dimensionality of the dynamics, while the number of components in the Gaussian mixture corresponds to the number of cell classes. Using back-propagation, we determine the minimum rank and number of cell classes needed to implement tasks of increasing complexity. We then exploit mean-field theory to reverse-engineer the obtained solutions and identify the respective roles of dimensionality and cell classes. We show that the rank determines the phase-space available for dynamics that implement input-output mappings, while having multiple cell classes allows networks to flexibly switch between different types of dynamics in the available phase-space. This overall picture leads to detailed predictions for the structure of neural selectivity in systems neuroscience experiments.

II-52. Neuromechanical modeling of Hydra behaviors

Hengji Wang
 Adrienne Fairhall
 University of Washington

HENGJW@UW.EDU
 FAIRHALL@UW.EDU

How behaviors arise from neural activity is one of the most fundamental questions in neuroscience. Small model systems pose the opportunity to develop complete models for the transformations from neural activity to behavior, taking into account the biomechanics of the body, including the properties of active muscle tissue. Transformation from neural firing to behavior requires appropriately modeling body and muscle dynamics, which can play important roles in transforming neural activation to mediate the behavioral outputs. We have constructed a model for the cnidarian Hydra which couples a biomechanical body model to a biophysical model for force generation by the epitheliomuscular network of Hydra, which is driven by a set of non-overlapping functional nerve nets. Our model successfully simulates Hydra's hydrostatic skeleton and produces some of the animal's major behaviors including contraction bursts (CB), active elongation, and bending. A novel aspect of our model combines the mechanisms of calcium influx and calcium release from internal stores in the epitheliomuscular cells to explain the multiple time scales of calcium dynamics that mediate different behaviors of Hydra. Our model serves as a testbed for exploring the correlations between neural activity and behavior and to reverse engineer the neuronal activity underlying distinct Hydra behaviors.

II-53. Serotonin neuron firing rates track expected uncertainty during dynamic decision making

Cooper Grossman
 Bilal Bari
 Jeremiah Cohen
 Johns Hopkins University

CGROSSM4@JHMI.EDU
 BBARI1@JHMI.EDU
 JEREMIAH.COHEN@JHMI.EDU

Reinforcement learning (RL) models propose that behavioral policies are learned through interactions between the nervous system and the environment. Within this framework, animals learn from discrepancies between expected and received outcomes of actions (reward prediction errors, RPEs). But the animals must “learn how to learn”; that is, the rate of learning from RPEs varies in accordance with statistics of the environment. For example, a single outcome of an action carries different information depending on the volatility of a probabilistic environment. Frequent environmental change results in recent rewards being more relevant and thus, drives faster learning. In a stable environment, slow learning is more optimal so that behavior is not driven by misleading short-term fluctuations in outcomes. This tuning of decision making is known as “meta-learning” and some have theorized that it is regulated by neuromodulatory systems.

One such neuromodulator is serotonin and prior research demonstrates that serotonin neurons are responsible for how quickly an animal adapts to a change in action-outcome contingencies. We therefore hypothesized that serotonin neurons may be involved in meta-learning. To test this hypothesis, we designed a dynamic foraging task for head-fixed mice and recorded electrophysiologically from serotonin neurons. We modified an RL model to include meta-learning. The model learns an estimate of the expected uncertainty in the environment by calculating a moving average of unsigned RPEs. How surprising an outcome is relative to this expected uncertainty determines how quickly the brain learns from that outcome. Intuitively, surprising outcomes carry more information because they may signal a change in the environment. Adding meta-learning to the model captured a unique feature of observed behavior. We found that firing rates of approximately 40% of serotonin neurons correlated with expected uncertainty over long timescales (tens of seconds). Thus, we demonstrate a quantitative link between serotonin neuron activity and behavior.

II-54. SMA-ACC network guides flexible strategy selection in rats

Mikhail Proskurin¹
 Maxim Manakov¹
 Maria Rysakova²
 Elena Kuleshova¹
 Andy Lustig¹
 Reza Behnam¹
 Dougal Tervo¹
 Alla Karpova¹

M.PROSKURIN@GMAIL.COM
 MANAKOV.MAKSIM@GMAIL.COM
 RYMARIK@GMAIL.COM
 KULESHOVA.ELENA.P@GMAIL.COM
 LUSTIGA@JANELIA.HHMI.ORG
 BEHNAMR@JANELIA.HHMI.ORG
 GOWANTERVO@GMAIL.COM
 KARPOVA@JANELIA.HHMI.ORG

¹HHMI Janelia Research Campus

²Institute of Higher Nervous Activity

The ability to efficiently acquire and organize knowledge is one of core aspect of intelligence. It often requires animals to make inferences and test specific hypotheses about the latent structured relationships in the environment. To study how brain may infer causal regularities in a complex environment, we designed an experimental framework that requires animals to do self-guided exploration and discover a latent sequence of actions that results in reward. To verify that animals indeed structure their choices into action sequences, we performed a number of behavioral tests and observed that animals can effectively discover and adapt to un-signaled changes in the rule (rewarded sequence) within a session. We found that rats perform a particular sequence when it is dominant (rewarded) but also as a part of resurgent exploratory bout. We used pharmacological inactivation to demonstrate functional involvement in this task of the rodent homologues of Anterior Cingulate Cortex (ACC) and Supplementary Motor Area (SMA). We revealed that SMA sends direct projections to ACC and demonstrated its functional relevance for the task. Targeted optogenetic suppression of SMA-ACC projections increased probability to repeat the target sequence but only in cases when animals encountered an unrewarded instance of that sequence. finally, we performed targeted electrophysiological recordings in both areas. We found that ACC ensemble representation associated with specific behavioral strategy scales with its relative prevalence. The sequence representation in ACC is modulated depending on the context: whether it is a dominant target sequence or it is part of an exploratory bout. In contrast, SMA ensemble activity associated with strategy representation remains largely stable across the contexts. Our findings suggest that SMA-ACC network may be arranged in a way where SMA provides a stable basis of sequence representation that ACC utilizes at the moments of uncertainty and the internally-guided need for further exploration.

II-55. Hierarchical Markov modeling of social behavior from 3D mouse pose

Caleb Weinreb¹
 Sandeep Robert Datta¹
 Alia Newman-Bouille¹
 Scott W Linderman²

CALEBSW@GMAIL.COM
 SRDATTA@HMS.HARVARD.EDU
 NEWMAN-BOULLE.A@HUSKY.NEU.EDU
 SCOTT.LINDERMAN@GMAIL.COM

¹Harvard University

²Stanford University

Computational approaches for unsupervised annotation of animal behavior have created new opportunities to investigate the structure of behavior and its relationship to neural activity. Yet these methods face a common challenge in behavior's hierarchical structure and its dependence on constantly changing sensory input. Commonly used methods that assume Markovian transitions between actions have a limited ability to capture these complex dynamics. During social behavior, for example, moment-to-moment kinematics form longer time-scale sequences such as aggression or investigation with transitions that depend on the actions of self and other. Social behavior presents the additional challenge that interacting partners can touch or occlude each other, making it difficult to accurately record their pose. Here, we use a novel algorithm to segment interacting mice, and develop a hierarchical Bayesian model to capture their behavioral dynamics across multiple timescales from 3D pose

measurements. The segmentation approach uses neural networks with a novel maximum likelihood approach to output a pixel-level labeling of each mouse, and performs with an accuracy comparable to the difference between two human annotators. To model joint behavior, we fit a family of models parameterized with varying degrees of mouse-mouse dependence or long time-scale dependence. Each additional dependency improved the likelihood of the model for held-out data, suggesting that the latent states of the model capture consistent behavioral patterns that are not detected by a purely Markovian approach. The most well-fit model defined three higher-order interaction states. We found a consistent pattern of state probabilities defined by the timing of fights between male mice, with state transitions in the subordinate mouse predicting an oncoming fight up to one minute before it occurs. Our model class can be generalized to other dynamic interactions such as between predator and prey, and suggests a framework for building on the initial success of Markovian methods for behavioral analysis.

II-56. Predictive coding explains the link between complex cells and the topographic organization of orientation maps in V1

Angelo Franciosini
Victor Boutin
Frederic Chavane
Laurent Perrinet

Aix Marseille Universite

ANGELO.FRANCIOSINI@UNIV-AMU.FR
VICTOR.BOUTIN@UNIV-AMU.FR
FREDERIC.CHAVANE@UNIV-AMU.FR
LAURENT.PERRINET@UNIV-AMU.FR

Cells in the primary visual cortex of mammals have historically been divided into two classes: simple and complex. Simple cells respond linearly to oriented visual stimuli while complex cells show various degrees of invariance with respect to the stimulus' phase (position). The existence of the two populations can be explained by hierarchical models where simple cells feed information into complex cells through a non-linear spatial pooling. Nevertheless, how the brain develops this representation remains an open question. One of the most successful theories to model hierarchical processing in the brain is Predictive Coding (PC): a framework introduced by Rao & Ballard that exploits feedback and feedforward connectivity to solve a Bayesian inference problem. We generalized PC in a convolutional network using a model called Sparse Deep Predictive Coding (SDPC). We demonstrate that this framework can easily replicate complex-like neurons when a non-linear pooling is included between the layers. In particular, we show that a large population of complex-like neurons, showing various degrees of phase invariance, emerge in the 2nd layer of the model when the pooling function is extended to include not only neighboring spatial locations but also neighboring neurons with different tuning properties. A striking emergent property of our model is that this non-linearity induces a topographical structure on the neurons of the network. This organization shows qualitatively strong similarities with columnar structures found in V1. We were able to reproduce similar effects with two different types of pooling (l2-pooling, max-pooling) and different network sizes, obtaining different degrees of topographical organization and different forms of complex-like cells populations. The novelty of this model lies in its ability to highlight the link between structure and function in a neural network. This study addresses a long-debated question on the function and role of the diversity of topographical structure in the visual cortex of mammals across species.

II-57. Optogenetically-evoked modulation of noise correlations in the visual cortex

Alan Akil¹
 Ariana Andrei²
 Robert Rosenbaum³
 Sorin Pojoga²
 Roger Janz²
 Valentin Dragoi²
 Kresimir Josic¹

AAKIL@MATH.UH.EDU
 ARIANA.R.ANDREI@UTH.TMC.EDU
 ROBERT.ROSENBAUM@ND.EDU
 SORIN.POJOGA@UTH.TMC.EDU
 ROGER.JANZ@UTH.TMC.EDU
 VALENTIN.DRAGOI@UTH.TMC.EDU
 JOSIC@MATH.UH.EDU

¹University of Houston

²University of Texas Health Science Center at Houston

³University of Notre Dame

Correlated variability is a ubiquitous feature of cortical networks, but its sources and impact on behavior are still unknown. Noise correlations between pairs of simultaneously recorded neurons have been studied, as they are an accessible way to quantify the collective activity of cells in a population. Changes in noise correlations have been linked to an assortment of behavioral conditions, including attention, arousal, as well as experience and learning.

Such correlation analyses have revealed crucial information about the functional architecture of cortical networks, providing information about network dynamics not discernible from single neuron recordings. Furthermore, correlated variability has been theorized to be both beneficial and detrimental for stimulus encoding, depending on the task and response properties of cells in which correlations are measured. However, there presently exists no method to causally manipulate correlations between neurons independent of other task parameters.

Here, we combine electrophysiology and computational modeling to better understand the mechanisms of correlated variability. Using optogenetic stimulation of a population of excitatory neurons in V1, we demonstrate that noise correlations can be modulated at both short and long-term timescales. First, we show that immediately following laser stimulation (hundreds of milliseconds), noise correlations between targeted neurons are dramatically reduced. Second, we show that correlations during laser stimulation drop over the course of many stimulation trials (minutes to hours), suggestive of concomitant plastic changes in the network. Both of these effects can be replicated in a model dynamically balanced network operating in a correlated state and subject to inhibitory plasticity. A theoretical analysis of this networks provides testable hypotheses about the mechanisms that underlie the experimentally observed changes in connectivity. We thus describe a novel method by which we can alter the functional connectivity in a target population to better understand the origins and impact of correlated variability in cortical networks, and an associated theoretical framework that allows us to understand the mechanisms driving these changes.

II-58. Continual weight updates and convolutional architectures for equilibrium propagation

Maxence Ernout¹
 Yoshua Bengio²
 Julie Grollier³
 Damien Querlioz¹
 Benjamin Scellier⁴

MAXENCE.ERNOULT@U-PSUD.FR
 YOSHUA.BENGIO@MILA.QUEBEC
 JULIE.GROLIER@CNRS-THALES.FR
 DAMIEN.QUERLIOZ@U-PSUD.FR
 BENJAMIN.SCELLIER@GMAIL.COM

¹Universite Paris Sud

²Montreal Institute for Learning Algorithms

³Unite Mixte de Recherche CNRS-Thales

⁴Universite de Montreal

Equilibrium Propagation (EP) is a biologically inspired alternative algorithm to backpropagation (BP) for training neural networks. It applies to RNNs fed by a static input x that settle to a steady state, such as Hopfield networks. EP is similar to BP in that in the second phase of training, an error signal propagates backwards in the layers of the network, but contrary to BP, the learning rule of EP is spatially local. Nonetheless, EP suffers from two major limitations. On the one hand, due to its formulation in terms of real-time dynamics, EP entails long simulation times, which limits its applicability to practical tasks. On the other hand, the biological plausibility of EP is limited by the fact that its learning rule is not local in time: the synapse update is performed after the dynamics of the second phase have converged and requires information of the first phase that is no longer available physically. Our work addresses these two issues and aims at widening the spectrum of EP from standard machine learning models to more bio-realistic neural networks. first, we propose a discrete-time formulation of EP which enables to simplify equations, speed up training and extend EP to CNNs. Our CNN model achieves the best performance ever reported on MNIST with EP. Using the same discrete-time formulation, we introduce *Continual Equilibrium Propagation* (C-EP): the weights of the network are adjusted continually in the second phase of training using local information in space and time. We show that in the limit of slow changes of synaptic strengths and small nudging, C-EP is equivalent to BPTT (Theorem~1). We numerically demonstrate Theorem~1 and C-EP training on MNIST and generalize it to the bio-realistic situation of a neural network with asymmetric connections between neurons.

II-59. Long-range horizontal connectivity in mammalian visual cortex for balancing performance and cost

Youngjin Park
Seungdae Baek
Se-Bum Paik
KAIST

YODAMASTER@KAIST.AC.KR
SEUNG7649@KAIST.AC.KR
SBPAIK@KAIST.AC.KR

Both our brain and deep neural networks(DNNs) successfully perform object recognition, but the visual pathway in the brain consists of fewer hierarchical layers than a DNN, presumably due to the restricted volume of the brain. Here, we suggest that the cortical long-range connections(LRCs) observed in various mammalian species can organize balanced local and global connections in the network, and they enable object recognition even in shallow neural networks. Using simulations of a model hierarchical network with convergent feedforward connections and LRCs from tree shrew data, we investigated the functional roles of LRCs for object recognition. first, we observed that the addition of LRCs to a shallow feedforward network significantly enhances the CIFAR-10 image classification performance so that it is comparable to that of a much deeper network. Second, through a network pruning with gradient-based optimization, we confirmed that LRCs could spontaneously emerge by minimizing the total connection length while maximizing classification performance. Although most of the initial connections were pruned during training, a certain portion of the very long connections survived until the end of training. Ablation of the surviving LRCs led to significant reduction of classification performance, which implies that LRCs are crucial for image classification. Lastly, we found that a combination of sparse LRCs and dense local connections organizes small-world type connectivity in a model network, and this dramatically increases performance per wiring cost. Furthermore, we found that performance enhancement by LRCs is strongly correlated with the small-worldness of a network, and this can explain species-specific existence of LRCs in the visual cortex. Taken together, we propose that long-range horizontal connectivity might be a key architecture of the visual cortex to implement parsimonious object recognition under the physical constraints of the brain.

II-60. Efficient and sufficient: generalized optimal codes

Luke Rast
Jan Drugowitsch
Harvard University

LUKERAST@G.HARVARD.EDU
JAN_DRUGOWITSCH@HMS.HARVARD.EDU

The efficient coding hypothesis proposes that neural codes are optimized to maximize task performance subject to various (e.g., metabolic) constraints, and has seen success explaining observed signatures of neural responses in a normative framework. Two dominant approaches to study optimized codes either: restrict analysis to simplistic models of how neural activity encodes the stimulus (e.g., specifically-tuned independently firing Poisson neurons), or: omit neural activity altogether, instead optimizing how information is distributed across the stimuli. Here, we unite the two approaches with a general analytic result that applies to a wide range of encoding models. It assumes only that the optimizable (monotonic, bounded) model parameters are sufficient statistics for the stimulus in the activity, and is otherwise agnostic to specifics of the encoding model and activity constraints. This allows for investigation of the effects of different features of the encoding, such as various activity costs, noise models, and population correlation structures, which could not be investigated fully in either specific encoding models or activity-agnostic approaches. In fact, our solution is sufficiently general to capture many previous solutions following from either approach, while remaining intuitive thanks to its relation to exponential family distributions. In particular, it clarifies the effects underlying the continuous shift from histogram equilibration (i.e., distributing neural resources according to stimulus probability), which is a hallmark of activity-agnostic solutions, to enforcing activity constraints, once those become dominant. This shift can be understood as a general property of efficient coding solutions, irrespective of the specifics of the encoding model, that, formally, results from the trade-off between a 'flat' base measure and constraint-enforcing energy terms. We highlight applications to neural populations with and without correlated noise.

II-61. Spatio-temporal structure of awake hippocampal replay

Emma Krause
Jan Drugowitsch
Harvard University

EMMAKRAUSE@G.HARVARD.EDU
JAN_DRUGOWITSCH@HMS.HARVARD.EDU

Hippocampal place cells in rodents represent the animal's current location in spatial navigation tasks. Strikingly, within sharp-wave ripple events during moments of immobility, this neural representation shifts to generating "simulated" spatio-temporal trajectories through the environment, which are believed to support planning in navigation—a phenomenon called awake hippocampal replay. Characterizing the structure of replay trajectories is essential for understanding their computational role in planning. For their detection and classification, however, the existing literature relies on heuristic methods, which discard the majority of ripple events, and most likely bias our computational interpretation. Here, we develop a novel, principled approach for revealing the spatio-temporal structure of replay events. We use Bayesian model comparison of state space models to characterize replay events by the presence/absence of structural features, such as trajectory-continuity, spatial diffusion, or momentum-dependence. Applied to hippocampal data collected during a foraging task in an open-field environment, we find that the vast majority of ripple events—even those discarded in previous studies—contain rich spatio-temporal structure. Most trajectories are characterized as combining spatio-temporal diffusion with momentum, leading to smooth, directed paths which mirror the physics of real movements. This stands in contrast to replay during sleep, which has recently been reported to resemble Brownian diffusion without momentum. Our findings argue against awake replay events being generated by local Markovian dynamics based on location only, and instead suggest them to emerge from network dynamics that incorporate a representation of velocity. Overall, the use of state space models provides a principled method for analyzing the spatio-temporal structure contained within sharp wave ripples that future work on their computational function can build upon.

II-62. Emergent evolutionary dynamics over recurrent neural computations

Czegel Daniel¹
Hamza Giaffar^{2,3}
Marton Csillag¹
Eors Szathmary¹

DANIELCZEGEL@GMAIL.COM
HGAFFAR@CSHL.EDU
MARTON.CSILLAG.11@GMAIL.COM
SZATHMARY.EORS@GMAIL.COM

¹Eotvos University

²Cold Spring Harbor Laboratory

³Watson School of Biological Sciences

Life, with its "endless forms most beautiful", is a result of a Darwinian evolutionary process, built on the principles of replication, selection and hereditary variation, combined with the enormous representational capacity of chemistry. The simplicity of these principles suggests that it is not impossible that they effectively operate at other levels of organizations, feeding on not-necessarily-chemical computational substrates. Neural substrates are certainly another example that produces endless forms most beautiful, generating a great diversity of action patterns and mental states. The endeavour of exploring the possible ways Darwinian evolutionary search might emerge as an effective high-level algorithmic mechanism from plasticity and activity dynamics of neural populations is called Darwinian neurodynamics (DN). In this work, we describe a framework for DN based on the emergent replication of dynamic patterns of neural activity between reservoir computing units. We show that an ensemble of such recurrent networks, capable of training one another based on previous evaluation of the signals they produced, can give rise to an effective evolutionary search that i) maintains information about already found solutions and ii) improves on them occasionally through the appearance and spread of beneficial mutations. We demonstrate the capability of such neural-evolutionary dynamics in solving combinatorial search problems. We illustrate the process and analysis methods on a prototypical example, the Traveling Salesman Problem (TSP), and then perform a more exhaustive analysis on tunably rugged NK landscapes. The proposed analysis methods, which combine ideas from computational neuroscience and phylogenetics, are generally applicable to any proposed DN architecture.

II-63. Fast and deep neuromorphic learning with first-spike coding

Julian Goltz¹
Andreas Baumbach¹
Sebastian Billaudelle¹
Oliver Breitwieser¹
Laura Kriener²
Akos Ferenc Kungl¹
Karlheinz Meier¹
Johannes Schemmel¹
Mihai A Petrovici²

JULIAN.GOELTZ@KIP.UNI-HEIDELBERG.DE
ANDREAS.BAUMBACH@KIP.UNI-HEIDELBERG.DE
SEBASTIAN.BILLAUELLE@KIP.UNI-HEIDELBERG.DE
OLIVER.BREITWIESER@KIP.UNI-HEIDELBERG.DE
KRIENER@PYL.UNIBE.CH
FKUNGL@KIP.UNI-HEIDELBERG.DE
MEIERK@KIP.UNI-HEIDELBERG.DE
SCHEMMEL@KIP.UNI-HEIDELBERG.DE
PETROVICI@PYL.UNIBE.CH

¹Heidelberg University

²University of Bern

For a biological agent operating under environmental pressure, energy consumption and reaction times are of critical importance. Similarly, engineered systems also strive for short time-to-solution and low energy-to-solution characteristics. At the level of neuronal implementation, this implies achieving the desired results with as few and as early spikes as possible. In the time-to-first-spike coding framework, both of these goals are inherently emerging features of learning. Here, we describe a rigorous derivation of error-backpropagation-based learning for hierarchical networks of leaky integrate-and-fire neurons. This narrows the gap between previous existing models of first-spike-time learning and biological neuronal dynamics, thereby also enabling fast and energy-efficient inference on analog neuromorphic devices that inherit these dynamics from their biological archetypes.

II-64. Goal-directed state space models of mouse behavior

Blue Sheffer¹
Sandeep Robert Datta²
Scott Linderman¹

BSHEFFER@STANFORD.EDU
SRDATTA@HMS.HARVARD.EDU
SCOTT.LINDERMAN@STANFORD.EDU

¹Stanford University

²Harvard University

A key challenge in neuroscience is to understand how neural circuits organize and select actions in order to achieve some goal. Unfortunately, in many experimental settings both the extent of an animal's behavioral repertoire and the goals that guide action selection are unknown. Here, we develop a probabilistic model and inference algorithm to simultaneously learn these action spaces and goals directly from behavioral data. We build on recent methods from the systems neuroscience community for segmenting behavioral data, and combine these methods with ideas from inverse reinforcement learning for discovering simple goals that can explain complex behavioral data. We demonstrate these goal-directed state space models of animal behavior in both simulation and in recordings of freely behaving mice. These techniques are a step toward extracting meaningful representations of behavior that could shed light on neural activity in the brain circuits that guide action selection and reinforcement learning.

II-65. Differential covariance of fMRI predicts structural connectivity and behavior

Yusi Chen¹
Qasim Bukhari²
Tiger Wuton Lin³
Terrence Sejnowski¹

CYUSI@UCSD.EDU
QBUKHARI@MIT.EDU
TIGER.W.LIN@GMAIL.COM
TERRY@SALK.EDU

¹Salk Institute for Biological Studies

²Massachusetts Institute of Technology

³Brain Corporation

Functional connectivity (FC) was first defined as the covariance of neural activities between brain regions extracted from resting state fMRI recordings (rs-fMRI). Previous research has suggested that FC reflects the neuronal baseline of the brain and disruptions in FC are indicative of neuropsychiatric disorders. There is now a rich repertoire of statistical algorithms used to define FC but few approaches have been made to systematically validate FC extracted from real recordings due to the lack of 'ground truth' connectivity patterns. Here, exploiting the link between FC and structural connectivity (SC), we attempt to systematically interpret and validate FC derived from differential covariance (dCov). Built upon the relationship between signal derivatives and the original signal, dCov have been shown to effectively reduce false positive connections in synthetic datasets and analytically proved to serve the purpose under proper assumptions. When applied to rs-fMRI recordings, dCov-FC is more similar to SC than covariance matrix and precision matrix. The result was replicated across subjects in both mouse and human datasets. In addition, in the Human Connectome Project (HCP) rs-fMRI dataset, subjects with a more integrated dCov-FC network tended to have shorter reaction time to multiple psychological tests. FC derived from dCov therefore could potentially provide a better description of brain anatomy and behavior.

II-66. Evolving the olfactory systems

Guangyu Robert Yang
 Peter Wang
 Richard Axel
 Larry Abbott
 Yi Sun
 Ashok Litwin-Kumar
 Columbia University

GYIANG.NEURO@GMAIL.COM
 PETERWANG724@GMAIL.COM
 RA27@COLUMBIA.EDU
 LFA2103@COLUMBIA.EDU
 YISUN@MATH.COLUMBIA.EDU
 AK3625@COLUMBIA.EDU

flies and mice are separated by 600 million years of evolution, yet they have evolved olfactory systems that share many similarities in their anatomic and functional organization. What functions do these shared features serve and why did they evolve in parallel? In this study, we address the evolutionary design of olfactory circuits by studying the properties of artificial neural networks trained to sense odors. We found that after learning to perform odor classification, layered neural networks quantitatively recapitulate structures inherent in the olfactory system, including the formation of glomeruli in a compression layer and sparse unstructured connectivity onto an expansion layer. We provide theoretical explanations for these results. Our work offers a framework for explaining the evolutionary convergence of olfactory circuits and gives insight into the anatomic and functional structure of the olfactory system.

II-67. Time continuity of early visual experience drives the development of ventral stream representations

Eis Annavini
 Mattia D'Andola
 Alessio Ansuini
 Davide Zoccolan

EIS.ANNAVINI@SISSA.IT
 MATTIA.DANDOLA@GMAIL.COM
 ALESSIOANSUINI@GMAIL.COM
 ZOCOLAN@SISSA.IT

SISSA– Scuola Internazionale Superiore di Studi Avanzati –

A long-standing hypothesis in the field of computational neuroscience is that sensory neurons are adapted to the statistical regularities of the signals they need to encode. Specifically, in the visual domain, unsupervised experience with temporal continuity of the visual input has been proposed as the key mechanism underlying the ability of visual neurons to tolerate transformations (e.g., translation and scaling) of visual objects. The causal role of unsupervised temporal learning (UTL) in determining the tuning of visual cortical neurons during postnatal development has been recently demonstrated in rat primary visual cortex (V1), but UTL is expected to govern the development of the whole cortical pathway underlying object recognition – i.e., the ventral visual stream. In our study, we causally investigated the role played by UTL in the developing ventral stream by rearing newborn rats either with natural movies (control group) or with their frame-scrambled versions (experimental group) for the whole duration of the critical period (60 days). We then performed acute neuronal recordings from the two extremes of the rat ventral stream (i.e., V1 and laterolateral area LL) in both groups of rats, as well as in naive animals (reared in normal cages), during presentation of natural objects with different shape, size, position, and luminosity. Consistently with previous studies, we found that, in naive rats, LL neurons displayed much lower sensitivity than V1 neurons to the luminosity and position of visual objects, as quantified by the performance of SVM classifiers in inferring these properties from neuronal population activity. Critically, such sensitivity difference was preserved in the control rats, but disappeared for the experimental group. This suggests that destroying the temporal continuity of early visual experience prevents ventral stream areas from developing their tolerance to variation of low-level properties, thus making object representations in these areas hardly distinguishable from those in V1.

II-68. A central source for fixational eye drifts

Nadav Ben-Shushan¹
 Nimrod Shaham²
 Mati Joshua¹
 Yoram Burak¹

NADAV.BEN-SHUSHAN@MAIL.HUJI.AC.IL
 NIMROD.SHAHAM@MAIL.HUJI.AC.IL
 MATI.JOSHUA@MAIL.HUJI.AC.IL
 YORAM.BURAK@ELSC.HUJI.AC.IL

¹The Hebrew University of Jerusalem

²Harvard University

During fixation and between saccades, our eyes undergo diffusive random motion called fixational drift. In the foveal region of the visual field, where receptive fields are dense, fixational drift causes stimuli to move across many receptive fields. The role of this motion in visual coding and inference has been debated in the past few decades. On one hand, it has been proposed that fixational motion may assist high acuity vision. On the other hand, it has been argued that fixational drift poses a challenge for high acuity visual inference, as it precludes the possibility to overcome spiking noise by simple temporal averaging. While the consequences of fixational drift have been studied from a normative perspective, the mechanisms that underlie this motion remained unknown. In particular, it has been unclear whether fixational drift arises from peripheral sources, or from central sources upstream of the muscles. To address this question, we analyzed simultaneous recording of neural activity of ocular motoneurons in the Rhesus monkey, alongside precise measurements of eye position obtained using implanted eye coils. Single motoneurons are too noisy to identify a correspondence between spiking activity and eye position during fixation over a single trial, but by analyzing data from 58 cells over hundreds of eye trajectories we show that the two are correlated, and the variability in activity of a single motoneuron accounts for much of the variability in eye motion during fixation. This is the first demonstration that fixational drift is correlated with central neural activity. Furthermore, several lines of reasoning indicate that much of the motion must arise from a central origin, upstream of the motoneurons. Using a line attractor model of the oculomotor integrator, we show that noise driven diffusion along the attractor accounts, with plausible parameters, for the amplitude and statistics of fixational drift.

II-69. Inhibitory output specificity and the speed-accuracy tradeoff in perceptual decision-making

James Roach¹
 Anne Churchland²
 Tatiana Engel¹

ROACH@CSHL.EDU
 CHURCHLAND@CSHL.EDU
 ENGEL@CSHL.EDU

¹Cold Spring Harbor Laboratory

²Cold Spring Harbor

During perceptual decision-making cortical neurons modulate their firing-rates to reflect animal choices. It has been recently shown that inhibitory neurons are as equally tuned for choice as excitatory neurons. However, existing circuit models of decision-making treat inhibitory neurons as an untuned, nonspecific pool that facilitates competition between excitatory neurons. We analyzed the impact of inhibitory choice tuning on decision-making performance by extending circuit models to account for the specificity of inputs to and outputs from inhibitory neurons. In this framework, inhibitory choice tuning arises from structured connections from excitatory to inhibitory sub-populations. For this tuning to have an effect on the circuit dynamics, inhibitory outputs also need to be specific. Inhibitory output specificity falls into two general classes: ipsiselective, where inhibitory neurons preferentially feedback to the excitatory neurons of the same tuning, and contraselective, where inhibitory neurons preferentially output to oppositely tuned excitatory neurons. These two classes of inhibitory output specificity have contrasting effects on the circuit dynamics underlying decision-making. Ipsiselective inhibition stabilizes recurrent activity of the excitatory sub-populations, slowing down decision-making and allowing the circuit to better estimate the category of the stimulus. Contraselective inhibition, on the other hand, maximizes competition between the excitatory sub-populations, which speeds up decision-making and frequently leads the circuit to misclassify the stimulus in a noise induced error. Thus, the specificity of inhibitory outputs places the circuit on one or another

side of the speed-accuracy tradeoff. We show how this tradeoff results from changes in the attractor dynamics that underlie decision-making in the circuit, specifically, due to changes in the time-constant of the unstable manifold of a saddle-point that separates the decision attractors. The model predicts divergent patterns of noise correlations between choice-selective populations for ipsiselective versus contraselective circuits. These model predictions provide a measure to identify which inhibitory motifs are present in cortical decision-making circuits.

II-70. Identifying and mitigating statistical biases in neural models of tuning and functional coupling

Pratik Sachdeva¹
 Jesse Livezey²
 Maximilian Dougherty²
 Sharmodeep Bhattacharyya³
 Kristofer Bouchard²

PRATIK.SACHDEVA@BERKELEY.EDU
 JLIVEZEY@LBL.GOV
 MAXDOUGHERTY@LBL.GOV
 BHATTASH@SCIENCE.OREGONSTATE.EDU
 KBOUCHARD@LBL.GOV

¹University of California, Berkeley

²Lawrence Berkeley National Laboratory

³Oregon State University

Neural activity is a function of factors in the external world (i.e., tuning) and other neurons, some of which may be simultaneously measured and analyzed (i.e., functional coupling), while others are not observed and contribute both shared and private variability. Statistical models (e.g., generalized linear models) have been used to describe neural responses with tuning and functional coupling, achieving high predictive accuracy while using parameters that provide insight into which internal and external factors are important and their relative importance. Past studies have demonstrated that the inclusion of coupling decreases the magnitudes of the tuning parameters. However, extracting conclusions about neural activity from model parameters requires that their estimates are unbiased. We assess the validity of parameters obtained through a tuning and coupling model by fitting it to data generated from the triangular model, a more complete model of neural activity that utilizes a latent state to model unobserved neurons. We demonstrate the existence of the simultaneous equations bias, which results in biased parameter estimates for the tuning and coupling model that potentially jeopardize past conclusions about neural computation. We propose two approaches to perform inference within the triangular model: first, an expectation-maximization approach that is constrained to account for non-identifiability within the triangular model, and second, a novel method called Iterated Two-Stage Factor Analysis (ITSFA). We validate these approaches on synthetic data and further demonstrate their consequential impact on tuning curves obtained from both single-unit recordings in monkey visual cortex and micro-electrocorticography recordings in rat auditory cortex. Specifically, we observe elevated tuning modulations relative to the tuning and coupling model, implying that the previously observed reduction in tuning may be an unmitigated parameter bias. Our analysis sheds light on the roles of the simultaneous equations bias and identifiability in neural models, and provides novel approaches to circumvent these issues.

II-71. Multi-dimensional perceptual decision making under dynamic belief states

Cheng Xue
 Lily Kramer
 Marlene Cohen
 University of Pittsburgh

XUECHENG19870106@GMAIL.COM
 LEK82@PITT.EDU
 COHENM@PITT.EDU

The world is rich with information in numerous feature dimensions. Yet we choose to inform our decisions using a very small proportion of external input based on an internal belief about its relevance. This selective mechanism is dynamic in itself, as beliefs emerge and evolve based on past experiences. To look into how beliefs are changed

to reflect reality, and how dynamically evolving beliefs about the relevance of sensory information affects how it is processed in the brain, we designed a two-feature discrimination task with stochastic task switching. The task requires macaque monkeys to make decisions based on the information they believe to be relevant, while reevaluating their belief in the behavioral rule on every trial. We recorded neuronal activity from visual cortical area V1 and parietal cortical area 7a while monkeys performed the task.

When faced with the same stimulus input, the monkeys were able to make decisions based on different feature dimensions depending on their evolving belief about the task context. Interestingly, there was an interaction between the monkeys' perceptual sensitivity and the strength their belief: when past experience challenges existing belief, the information encoded in visual cortex seems to be suboptimally used to guide decisions, worsening visual discrimination performance. We also found evidence for a 'sensitivity-flexibility trade-off', i.e. the better the subjects are at a particular task, the longer it takes them to adapt their belief to reality. We built a model to predict choices on individual trials based on the trial history. This framework will provide a novel platform for understanding the neuronal substrates of trade-offs between perception and flexibility during perceptual decision making under dynamic belief.

This work is supported by NIH grant R01EY022930, a grant from the Simons Collaboration for the Global Brain, and a McKnight Scholar Award.

II-72. Precise coupling of the thalamic head-direction system to hippocampal ripples

Guillaume Viejo^{1,2}
Adrien Peyrache^{1,2}

GUILLAUME.VIEJO@GMAIL.COM
ADRIEN.PEYRACHE@MCGILL.CA

¹McGill University
²Montreal Neurological Institute

The anterior thalamus is a key relay of neuronal signals within the limbic system. During sleep, the occurrence of hippocampal sharp wave-ripples (SWRs), believed to mediate consolidation of explicit memories, is modulated by thalamocortical network activity, yet how information is routed around the time of SWRs and how this communication depends on neuronal dynamics remain unclear. Here, by simultaneously recording ensembles of neurons in the anterior thalamus and local fields potentials in the CA1 area of the hippocampus, we show that the head-direction (HD) cells of the anterodorsal nucleus are set in stable states immediately before SWRs. This response contrasts with other thalamic cells that exhibit diverse couplings to the hippocampus, related to their intrinsic dynamics but independent of their anatomical location. Thus, our data suggest a specific and homogeneous contribution of the HD signal to hippocampal activity but a diverse and cell-specific coupling of non-HD neurons.

II-73. Kura-Net: Exploring systems of coupled oscillators with deep learning

Matthew Ricci
Aneri Soni
Thomas Serre
Yuwei Zhang
Minju Jung
Brown University

MATTHEW_RICCI_1@BROWN.EDU
ANERI_SONI@BROWN.EDU
THOMAS_SERRE@BROWN.EDU
ZYW@MAIL.NANKAI.EDU.CN
MINJU_JUNG@BROWN.EDU

The recent proliferation of powerful deep learning algorithms has initiated a fruitful dialogue between systems neuroscience and machine learning. However, one prevalent feature of neural systems remains largely unmodeled by contemporary statistical methods: the abundant and much speculated-upon presence of cortical oscillations.

To help fill this gap, we present a novel machine learning formalism for exploring the dynamical regimes of the celebrated Kuramoto model, a simple but hugely expressive model of oscillatory systems. Our formalism constructs the couplings and intrinsic frequencies of Kuramoto oscillators as the output of a deep neural network with learnable parameters and an input channel for external stimuli. We then formulate the Kuramoto dynamics as a recurrent neural network whose dynamical state can be evaluated by a loss function measuring a statistic of interest, such as global synchrony. The entire system, “Kura-Net”, is differentiable so that the parameters governing the Kuramoto dynamics can be learned through backpropagation. To expedite learning, we propose a new loss function which uses the rate of convergence in the Kuramoto dynamics to overcome discontinuities caused by Kuramoto phase transitions. We evaluate our method using two tasks purported to require oscillatory mechanisms in cortex: visual relation detection and perceptual grouping. Our system is able to learn both tasks on a training data set and in one case generalize to a testing set. Importantly, our model uses task performance to explain earlier results indicating that perceptual grouping using oscillators relies on a mix of short-range excitatory and long-range inhibitory couplings. We believe that this and similar methods could be used not only to study real-world oscillatory data but also to augment other machine learning formalisms.

II-74. Synergistic population coding through correlated variability in V1

Daniel Denman¹
Clay Reid²

DANIEL.DENMAN@CUANSCHUTZ.EDU
CLAYR@ALLENINSTITUTE.ORG

¹University of Colorado

²Allen Institute for Brain Science

Sensory stimuli are represented by the joint activity of large populations of neurons across the mammalian cortex. Information in such responses is limited by trial-to-trial variability. Because that variability is not independent between neurons, it has the potential to improve or degrade the amount of sensory information in the population response. How visual information scales with population size has been extensively modeled, yet remains an open empirical question, especially beyond the retina. Here, we use Neuropixels to simultaneously record tens to hundreds of single neurons in primary visual cortex (V1) and lateral geniculate nucleus (dLGN) of mice. These recordings allowed us to estimate information within and across V1 layers and dLGN. We found a mix of synergistic and redundant coding: synergy predominated in small populations (2-12 cells) before giving way to redundancy. The shared variability of this coding regime included global shared spike count variability at longer timescales, layer specific shared spike count variability at finer timescales, and shared variability in spike timing (jitter). Jitter linked ensembles span layers, and such ensembles carry more information than random or uncorrelated ensembles. Our results suggest fine time scale stimulus encoding may be distributed across physically overlapping but distinct ensembles in V1.

II-75. Variational Autoencoder account of the early visual hierarchy

Gergo Orban
Mihaly Banyai
David Nagy
Wigner RCP

ORGERGO@GMAIL.COM
BANYAI.MIHALY@WIGNER.MTA.HU
NAGY.G.DAVID@WIGNER.MTA.HU

Visual perception is the process of making inferences about latent variables in the environment based on observations. There is strong neural and behavioral support that the cortical implementation of this process is hierarchical and also represents uncertainty about the latents. Thus, given the success of supervised feed-forward neural networks in predicting mean neural responses, nonlinear probabilistic models of complex images, featuring hierarchically organized latents, are well-motivated to predict not only the mean but also higher-order statistics of neural responses. We fit hierarchical variational autoencoders, which inherently utilize both feed-forward and

recurrent connections, to image data in an unsupervised manner to make predictions about the representation of textures in early visual cortical areas. We demonstrate that consecutive latent layers learn increasingly compressed representations of stimuli, while also producing a disentangled representation of contrast. The proposed computational framework accounts for two distinct experimental observations in macaques. First, the linear decodability of the identity of texture stimuli has been found to be higher in V1 than in V2, while the reverse is true if the texture family of the stimulus is to be decoded. We show that our model has the same property, indicating that training discards local feature information from the second layer of latents present in the first, while retaining information about global stimulus statistics. Second, noise correlations in V1, in addition to being stimulus-specific, have been measured to vary according to the higher-order statistical content of the stimulus. We demonstrate that in our model, noise correlations in the first layer of latents show greater variance between stimuli from different texture families due to the effect of the top-down contextual prior from the second layer. These results show that hierarchical Bayesian models naturally extend feed-forward models to the probabilistic, unsupervised domain and account for a range of anatomical and electrophysiological observations.

II-76. For bursts with plateau potentials, GCaMP6-signal reflects plateau duration rather than burst size

Peter Ledochowitsch
Lawrence Huang
Gabriel Ocker
Michael Oliver
Jack Waters
Hongkui Zeng
Saskia de Vries
Michael Buice

LEDOCHOWITSCH@GOOGLEMAIL.COM
LAWRENCEH@ALLENINSTITUTE.ORG
GABEO@ALLENINSTITUTE.ORG
MICHAELO@ALLENINSTITUTE.ORG
JACKW@ALLENINSTITUTE.ORG
HONGKUIZ@ALLENINSTITUTE.ORG
SASKIAD@ALLENINSTITUTE.ORG
MICHAELBU@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Two-photon calcium imaging with genetically expressed calcium indicators (GECI), e.g. of the GCaMP6 family, is commonly used to monitor the spiking of large populations of neurons. Recovering action potentials (spikes) from fluorescence necessitates calibration experiments, often with simultaneous imaging and cell-attached recording. In this work, we present cell-attached recordings paired with simultaneous *in vivo* two photon calcium imaging from 35 neurons in 3 transgenic mouse lines. Our calibration indicated that 3 or more spikes were required to produce changes in fluorescence for which the chance of detection exceeds 50%. Moreover, neither a simple linear model nor a more complex biophysical model accurately predicted fluorescence for small numbers of spikes (1-3). Current attempts to infer the timing of action potentials from calcium-associated fluorescence typically model action potentials as the sole source of calcium influx in the neuron. In some neurons, we observed increases in fluorescence that were better explained by the duration of present plateau potentials than by the burst size (number of spikes contained in the burst preceding the transient), particularly in *Emx1-IRES-Cre* mouse line crosses. Our results indicate that deriving spike times from fluorescence measurements using models that assume no other source of calcium influx but action potentials may be an intractable problem in some mouse lines. Finally, we developed a lumped-element SPICE model of a patched neuron to investigate in simulation under which conditions plateau potentials may be observable in cell-attached recordings. Our simulations use biophysically realistic parameters that were either measured or obtained from literature. We successfully simulated realistic spike amplitudes for whole-cell, extracellular, and cell-attached recordings, respectively, and concluded from the simulations that our observed plateau potentials were most consistent with cell-attached recordings in which the cell membrane was intact, and the seal resistance was high (>100 MΩ).

II-77. Reorganization of phase space structure during network learning

Rainer Engelken
Larry Abbott
Columbia University

RE2365@COLUMBIA.EDU
LFA2103@COLUMBIA.EDU

The coordinated activity of biological neural circuits gives rise to complex input-output functions that can be mimicked mathematically in recurrent network models. Learning often reshapes network dynamics to a low-dimensional manifold. The mechanism behind this reorganization of the phase space is not well understood. How do subtle changes in the network structure cause drastic changes in the dynamics? Which mathematical tools are suitable for understanding the emergence of low-dimensional neural manifolds during learning? How does learning reshape the dimensionality of network activity?

To answer these long-standing questions in recurrent firing-rate network models, we present a spatiotemporally resolved analysis of dynamic stability during learning. Our approach applies to arbitrary network structures and firing-rate dynamics. Using concepts from dynamical systems theory, we calculate the attractor dimensionality and dynamic entropy rate of the network dynamics during learning based on the full Lyapunov spectrum. We calculate covariant Lyapunov vectors to map out stable and unstable manifolds.

We show that the dimensionality of the network activity after learning reflects the dimensionality of the task (fixed point, limit cycle, chaotic attractor). We analyze the learning dynamics and dynamic stability of different learning algorithms (FORCE, full-FORCE, backpropagation through time (BPTT)) and study when and how learning fails for excessively difficult tasks. A time-resolved analysis of instability reveals that failure of learning is often accompanied by an inability to stabilize the trajectory locally. We show that perturbations along the least-stable direction during training help to stabilize the trajectory and increase the basins of attraction.

Our study opens a novel avenue for characterizing the complex dynamics of rate networks and their reorganization during learning. We discuss how this not only gives a deeper understanding of the dynamics but can also help harness its computational capacities, e.g., for learning stable trajectories.

II-78. Meta-learning biologically plausible learning algorithms with feedback and local plasticity

Jack Lindsey¹
Ashok Litwin-Kumar^{1,2}

JACKWLINDSEY@GMAIL.COM
AK3625@COLUMBIA.EDU

¹Columbia University
²Neuroscience

In recent years, there has been renewed interest in developing biologically plausible learning rules that are effective in distributed and multi-layered networks. However, to date, local synaptic learning rules like those employed by the brain have failed to approach the performance of the backpropagation algorithm commonly employed in deep artificial neural networks. In this work, we employ meta-learning — optimization of the learning algorithm itself — to discover networks that learn in a biologically plausible manner, using feedback connections and local, Hebbian rules. Importantly, the networks we consider are free from several biologically implausible features of standard deep learning methods, as the feedback connections are not tied to the feedforward weights, and no explicit derivative computations are required. Our approach allows us to efficiently explore the space of architectures and learning rules that perform well under these constraints. Experiments show that meta-trained networks effectively use feedback to perform credit assignment in multi-layer architectures. We find that (1) feedback modulation of neural activity at multiple layers is crucial for task performance. (2) The resulting learning algorithm produces synaptic weight updates that deviate significantly from traditional gradient descent. (3) Simultaneously performing local unsupervised learning and feedback-modulated supervised learning benefits task performance. (4) Neurons can learn while multiplexing feedforward and error information. When compared with a state-of-the-art gradient descent-based machine learning algorithm, this model is competitive and even attains superior performance in

online continual and few-shot learning settings. Thus, this approach represents progress toward biologically plausible learning mechanisms that not only match modern machine learning approaches, but also overcome their limitations. Moreover, it enables a research program of using task-optimized neural networks — in which the task is learning itself — to elucidate the synergistic roles that feedback connections and plasticity rules play in neural circuits.

II-79. Biophysics-based frequency domain spatiotemporal pattern separation of the multichannel lfp signals

Weiwei Chen

Anton Sirota

Ludwig-Maximilians University Munich

CHEN@BIOLOGIE.UNI-MUENCHEN.DE

SIROTA@BIOLOGIE.UNI-MUENCHEN.DE

Multichannel high-density extracellular recording enables higher resolution of the spatial and temporal profile of mean-field neural dynamics. Identification of distinct sources of dynamics relies on the disentangling of the multichannel LFP signal based on the anatomy and biophysics of their synaptic inputs, as well as their geometry. However, accurate and robust source decomposition remains challenging due to the violation of the independence assumption as multiple inputs emerge from generally correlated dynamics, nonlinear and noninstantaneous mixing of distinct inputs via the dendrites, and, typically, undersampling in the signal space.

We propose a source separation approach in the frequency domain that overcomes some of these limitations. Meanwhile, frequency-domain ICA suffers from finding biophysically meaningful components and grouping the components from different frequencies. Here we resolve this problem by employing a biophysically-grounded model of LFP, linking the EPSCs to the spatiotemporal LFP patterns via the cable equation. It captures the main aspects of the dendritic mixing and linear spatial mixing process due to volume conduction. Based on this, we extend the Green's function-based inference and use it to guide the frequency domain ICA.

We test the performance on NEURON simulated LFP data generated by spatially overlapping and temporally correlated inputs. We first show that the ICA components match the input response patterns and the reconstructed source dynamics match those in the single pathway simulations. Afterward, we simulate mixtures arising under a high input coherence regime and find that our model still closely recovered true inputs, while conventional ICA results in incomplete unmixing. Furthermore, inverting the component with Green's function bases accurately recovers the input current distributions.

We conclude that biophysical modeling of the LFP signal in the frequency domain enhances source separation by guiding the algorithm to find a biological proper pattern and improves the interpretability of such decomposition analysis.

II-80. Combining MRI connectomics and reservoir computing to quantify the coding properties of large-scale brain networks

Estefany Suarez¹

Guillaume Lajoie²

Bratislav Misic¹

¹McGill University

²Universite de Montreal, Mila

LAURA.SUAREZ@MAIL.MCGILL.CA

LAJOIE@DMS.UMONTREAL.CA

BRATISLAV.MISIC@MCGILL.CA

Current theoretical accounts of brain function emphasize the idea that perception, cognition and complex behavior are the result of the interaction of multiple brain modules endowed with specialized information-processing capabilities. Brain connectomics offers an opportunity to investigate the link between the organizational features of brain networks and the spectrum of cortical functions. Here, the goal is to leverage reservoir computing models

to quantify the encoding and decoding capacity for distinct brain areas. We used brain connectivity measures obtained from MRI data to constrain the connections within the reservoir. Using this setup, we devised metrics to measure the encoding and decoding capacity of brain's functional modules with different underlying connectivity. We show that some brain modules are better encoders, while other are better decoders. These roles corroborate the computational properties of these networks, and recapitulate the information-processing hierarchy of the brain that spans from unimodal to transmodal areas. Specifically, we show that primary sensory and unimodal association areas located at the lower and intermediate levels of the brain's functional hierarchy, whose main role is to encode the features of the sensory experience, perform better as what we have denominated encoders. In contrast, brain regions at the highest levels of the functional hierarchy located in transmodal areas, whose major role is to act as multifocal integration points, perform better as what we have denominated decoders. Our findings highlight the role of the brain's underlying network architecture on the emergence of the spectrum of the computational properties of the brain's functional modules.

II-81. Experimentally observed features of sequential activity with temporally asymmetric Hebbian learning

Maxwell Gillett

Nicolas Brunel

Duke University

MAX.GILLETT@GMAIL.COM

NICOLAS.BRUNEL@DUKE.EDU

Sequential activity is a prominent feature of many neural systems, where it plays a key role in temporal information encoding. Numerous models have explored how neural activity can shape network connectivity to generate sequential activity through various synaptic learning rules, but few models can reproduce the diverse range of experimentally observed temporal statistics during sequence retrieval and over the course of multiple recording sessions. We find that a temporally asymmetric Hebbian learning rule provides a unifying framework for these observations, accounting for a wide range of statistical features through simple variation in learning rule parameters or addition of noise. In this rule, a sequence of input patterns are presented to the network, which result in a sequence of synaptic weight updates. Following learning, the sequence can be retrieved by an appropriate transient input. In such sequences, neurons can be characterized by their 'time fields', an interval of time in which they are active during sequence retrieval. We find that the distribution of time field centers is non-uniform, with an overrepresentation of fields at the beginning of the sequence, consistent with observations in multiple areas. We also find that depending on the details of the learning rule, the width of time fields can either increase with time (as in time cells in cortex and hippocampus), or remain constant in time (as in area HVC of adult zebra finch). Recent longitudinal recordings made in posterior parietal cortex and hippocampus have revealed sequential activity that at the single unit level is unstable over the course of many days, but in which task variables can still be stably read out. We find that ongoing synaptic dynamics due to noise or additional pattern storage can produce such unstable dynamics on long timescales, while still preserving decoding of stored sequential pattern activations at the population level.

II-82. Single neurons in rat orbitofrontal cortex predict confidence-guided time investment and choice updating

Torben Ott¹

Paul Masset²

Armin Lak³

Junya Hirokawa⁴

Adam Kepecs¹

¹Cold Spring Harbor Laboratory

²Harvard University

OTT.TORBEN@GMAIL.COM

PAUL_MASSET@FAS.HARVARD.EDU

ARMINLAK@GMAIL.COM

RATCORTEX@GMAIL.COM

KEPECS@CSHL.EDU

³Oxford University⁴Doshisha University

Every decision we make is accompanied by a sense of confidence about its likely outcome. Confidence judgments have been studied as a central aspect of metacognition in humans, but estimating the degree of confidence confers an adaptive advantage for any organism from the sophisticated to the mundane: deciding between fight and flight, choosing which drink to order or managing a stock portfolio. Which properties are required for a neural representation of confidence? first, neural activity should reflect a confidence computation, i.e. reflect the statistical likelihood that a proposition is correct, $P(\text{choice}=\text{correct} \mid \text{evidence, choice})$. Second, neural activity should be related to decisions based on confidence, i.e. predict confidence-guided behavior. A prerequisite for metacognition is an abstract and centralized confidence representation with two additional criteria: Abstract confidence representations should generalize across the source of information used for probability computations, and generalize across different types of confidence guided behaviors. Neurons in several brain areas meet the first requirement; however, it is unknown if there are neural representations of confidence that fulfill these criteria and thus establish abstract confidence representations. Here we tested these criteria by training rats to make perceptual choices based on either an ambiguous smell or sound and invited them to invest variable time for a potential reward after each decision. We developed normative decision models to show that decision confidence guided rats' trial-by-trial time investment behavior as well as rats' trial-by-trial updates in choice bias. Single neurons in orbitofrontal cortex (OFC) neurons represented statistical confidence and predicted both confidence-guided behaviors. Neurons predicted time investment and choice updating in the next trial irrespective of the sensory modality, olfactory or auditory, used to make a choice. Orbitofrontal cortex thus provides a generalized encoding of confidence consistent with a metacognitive process that is useful for mediating confidence-guided economic decisions.

II-83. Modeling behaviorally relevant neural dynamics with a novel preferential subspace identification (PSID)

Omid Sani¹Bijan Pesaran²Maryam Shanechi¹

GHASEMSA@USC.EDU

BIJAN@NYU.EDU

SHANECHI@USC.EDU

¹University of Southern California²New York University

In addition to any specific behavior of interest, neural activity often relates to other brain functions as well as internal states with brain-wide representations such as thirst. Dissociating these behaviorally irrelevant neural dynamics from the behaviorally relevant ones is key in understanding how neural dynamics explain behavior. We develop a novel algorithm termed PSID for dissociating and modeling behaviorally relevant neural dynamics. PSID dissociates and models the temporal structure in behaviorally relevant neural dynamics in terms of a latent brain state. first, with extensive numerical simulations, we show that PSID learns behaviorally relevant neural dynamics much more accurately, using markedly lower-dimensional latent states and orders of magnitude less training data. Next, we use PSID to study motor cortical dynamics in two non-human primates performing naturalistic 3D reach, grasp and return movements while 27 joint angles in their shoulder, elbow, wrist and fingers are being tracked. PSID reveals that the behaviorally relevant neural dynamics are markedly lower dimensional than implied by standard methods. In addition, PSID uncovers distinct rotational dynamics in neural activity during reach and return epochs that are missed by standard methods and are significantly more predictive of behavior. finally, PSID learns more accurate behaviorally relevant dynamics from almost all recording channels and for almost all joints. Taken together, PSID provides a general new tool for studying and modeling behaviorally relevant neural dynamics that can otherwise go unnoticed.

II-84. Normalized reinforcement learning predicts asymmetric and adaptive choice behavior

Kenway Louie

KL837@NYU.EDU

New York University

Novel, uncertain, or dynamic environments require organisms to learn appropriate behavior based on environmental feedback. This learning is widely modeled in psychology, neuroscience, and computer science by prediction error-guided reinforcement learning (RL) algorithms. While standard RL models assume a linear reward function, growing empirical evidence suggests that neural activity in diverse brain areas is a saturating, nonlinear function of reward; however, the computational and behavioral implications of nonlinear RL models are unknown. Here, we characterize the behavior of an RL algorithm with a specific nonlinear value function implemented via divisive normalization. We show that normalized reinforcement learning (nRL) models exhibit comparative advantages over standard RL models under biologically-relevant conditions, including adaptive noise filtering for signal-dependent noise and improved choice under skewed reward distributions. Furthermore, nRL models reproduce empirical differences in choice behavior driven by outcomes better or worse than expected, a finding typically attributed to asymmetric learning rates; under nRL with a single learning rate, both the degree and direction of asymmetry is controlled by overall reward expectation. These results suggest that nonlinear RL reflects an adaptive response to ecological constraints, and argue for an incorporation of biologically valid value functions in computational models of learning and decision-making.

II-85. Modeling the population-level statistics and structure of hippocampal place cell responses

Man Yi Yim¹

Jae Sung Lee²

Albert Lee²

Thibaud Taillefumier¹

Ila fiete³

MANYI.YIM@UTEXAS.EDU

LEEJ13@JANELIA.HHMI.ORG

LEE@JANELIA.HHMI.ORG

TTAILLEF@AUSTIN.UTEXAS.EDU

FIETE@MIT.EDU

¹University of Texas at Austin

²HHMI Janelia Research Campus

³Massachusetts Institute of Technology

In large environments, individual hippocampal place cells can have multiple place fields with highly diverse statistics. The distribution of the number of fields per cell is long-tailed and cells with higher field propensity in one environment tend to fire more across environments. It is unclear what factors drive these statistics, and whether they arise because of dynamical constraints or an optimization of some coding property. The finding seems information-theoretically counterintuitive: efficient codes are typically whitened so that each coding element has an equal contribution probability. Thus we explore through computational modeling the hypotheses that populations of place cells compete to respond to their drive from grid cells and sparse external spatial cues, and that the grid-cell input influences the field statistics of place cells. We show that the observed population statistics of place fields can be robustly replicated with associatively learned or random grid-cell-to-place-cell connections in single environments. Further, the model predicts that relative field propensities across cells will remain unchanged across environments because the average response of grid cells is invariant across environments (once averaged across a length-scale of the period). Finally, we make new predictions for the structure of place-field distributions: A) The distribution of inter-field intervals (Ifis) should be structured, with peaks at specific values that can be derived analytically. B) The discrete spatial frequencies associated with grid periods should be overrepresented in the power spectral density (PSD) of place cells, even in the presence of non-grid spatial input and noise. These serve as robust predicted signatures of information flow from grid to place cells.

II-86. Predictive coding via optimal sampling regulates processing of odor dynamics

Morgan Brown
Yashar Ahmadian
Roma Shusterman
University of Oregon

MORGANALLEN@GMAIL.COM
YASHAR@UOREGON.EDU
ROMAS@UOREGON.EDU

During active sensing, we control sensory information by adjusting our sampling behavior. Different sampling routines are deployed for different tasks. In olfaction, odor identification or localization tasks evoke faster sniffing. However, if modulating sniff rate benefits neural processing remains unknown. Here we present an investigation into how variations in sampling shape the processing of odor gradients in the olfactory bulb. This computation is critical to mice because efficient sensing of concentration changes enables rapid localization of odor sources or early detection of predators. Here we present the first data indicating that fast sampling enables the olfactory system to encode odor gradients efficiently, utilizing a coding model which filters out redundant information. By delivering odor stimuli to awake mice with sniff locked changes in concentration we investigated the stability of the olfactory bulb response (MT cells) as a function of variations in sniff rate. We found that sniff rate strongly affects the processing of time-varying input. Exclusively at high sniff rate, response to identical odor concentration was suppressed. This suppression increases efficiency and enhances the signal associated with subsequent changes. The observed response exhibited a sharp transition with sniff rate. We further characterized the observed transition in odor coding via the machinery of GLM modeling, building sniff and neural response kernels as a function of odor. This enabled us to quantify the observed transformation the olfactory coding and provided a benchmark to compare between cells and sniff frequencies. We then replaced odor with direct optogenetic stimulation and established kernels to optogenetic activation of MT cells, bypassing the complex processing of the nasal cavity and olfactory sensory neurons. This approach enabled further classification of the mechanism driving these changes, specifically investigating if it is feed forward or feedback. This work represents a step toward understanding active sampling in odor information processing.

II-87. Minimax dynamics of optimal excitatory-inhibitory balanced networks

Qianyi Li
Cengiz Pehlevan
Harvard University

QIANYI_LI@G.HARVARD.EDU
CPEHLEVAN@SEAS.HARVARD.EDU

Excitation-inhibition (E-I) balance is ubiquitously observed in the cortex. Recent studies suggest an intriguing link between balance on fast timescales, tight balance, and efficient coding of information with spikes. By deriving a spiking neural network from a greedy-minimization of a reconstruction error cost function, Boerlin et al. (2013) showed that tight balance arises from correcting for coding errors on a fast timescale. While a powerful method for designing optimally coding balanced networks, the resulting networks typically disobey Dale's law. In this work, we take a principled approach to optimal balanced networks of E-I networks by considering a greedy spike-based optimization of constrained minimax objectives. In these networks, tight balance arises from correcting for deviations from the minimax optimum. We predict specific neuron firing rates in the network by solving the minimax problem, going beyond statistical theories of balanced networks. finally, we design minimax objectives for predictive coding (reconstruction) of an input signal and associative memory, and derive from them E-I networks that perform the computation.

II-88. Compressed interpretable models of neurons in cortical area V4

Reza Abbasi-Asl¹
Bin Yu²

REZA.ABBASIASL@UCSF.EDU
BINYU@BERKELEY.EDU

¹University of California, San Francisco

²University of California, Berkeley

Functions of neurons in cortical area V4 is less understood compared to other areas in ventral visual pathway such as V1, V2 and IT. Recently, models based on convolutional neural networks (CNNs) have shown promise in predicting the activity of V4 neurons and interpreting these models has opened new doors to understand their functions. However, these interpretations are based on models with hundreds of convolutional filters. Therefore, it is challenging to present a sparse set of filter basis for neurons. To address this, we propose two algorithms to remove redundant filters in CNN-based models of V4 neurons. CAR compression prunes filters in a CNN based on their contributions to the image classification accuracy. Similarly, RAR prunes filters based on their contribution to the neural response prediction accuracy. Both CAR and RAR give rise to a new set of simpler accurate models for V4 neurons. These models achieve almost similar (for CAR) or higher (for RAR) accuracy compared to the original model. Using the compressed models, we are able to characterize each V4 neuron with a sparse set of filters. We show that V4 neurons are modeled via curvature or texture filters, as well as other more complicated filters. Our results present one of the first efforts to bridge between large-scale convolutional models of neurons and interpretable sparse networks.

II-89. Balanced spiking networks with conditionally Poisson neurons

Camille Rullan¹
Jonathan Pillow²

CER454@NYU.EDU
PILLOW@PRINCETON.EDU

¹New York University

²Princeton University

Balanced spiking networks (BSNs) provide a powerful framework for implementing arbitrary linear dynamical systems in networks of integrate-and-fire neurons [Boerlin et al 2013]. However, the classic BSN model requires near-instantaneous transmission of spikes between neurons, which is biologically implausible. Introducing realistic synaptic delays leads to a pathological regime known as “ping-ponging”, in which different populations spike maximally in alternating time bins, causing network output to overshoot the target solution. Here we document this phenomenon and provide a novel solution: we show that a network can have realistic synaptic delays while maintaining accuracy and stability if neurons exploit conditionally Poisson firing. Formally, we propose two alternate formulations of Poisson balanced spiking networks: (1) a “local” framework, which replaces the hard integrate-and-fire spiking rule within each neuron by a “soft” threshold function, such that firing probability grows as a smooth nonlinear function of membrane potential; and (2) a “population” framework, which reformulates the BSN objective function in terms of expected spike counts over the entire population. We show that both approaches offer improved robustness, allowing for accurate implementation of network dynamics with realistic synaptic delays between neurons. Moreover, both models produce positive correlations between similarly tuned neurons, a feature of real neural populations that is not found in the original BSN. This work unifies balanced spiking networks with Poisson generalized linear models and suggests several promising avenues for future research

II-90. Motor signals promote rapid basis set formation in model cerebellar granule cell layer

Jesse Gilmer¹
Ioannis Delis^{2,3}
Abigail Person¹

JESSE.I.GILMER@GMAIL.COM

ABIGAIL.PERSON@UCDENVER.EDU

¹University of Colorado

²I.Delis@leeds.ac.uk

³University of Leeds

The cerebellum supports predictive encoding that improves movement, sensory processing, and cognition. Purkinje cells, the principal neurons of the cerebellar cortex, receive inputs from parallel fibers originating from granule cells, and learn to adjust synaptic weights through instructive feedback from the inferior olive. This associative learning rule is hypothesized to be supported by information preprocessing by the cerebellar granule cell layer (GCL) that sparsens associative information into a temporally diverse 'basis set'. However, rates of learning differ widely across paradigms, ranging from 10s of trials in saccade adaptation to 100s of trials in associative eyeblink conditioning. This discrepancy raises the question of whether learning rules or preprocessing mechanisms differ between these tasks. A recent study reported that classical conditioning coupled with locomotion sped learning, suggesting that movement-related signaling may be beneficial to learning. We hypothesized that time-variant physiological signals associated with ongoing movements would support robust pattern separation in the GCL, enhancing learning, while time invariant associative signals used in typical classical conditioning paradigms result in an impoverished 'basis set', making learning more difficult. To test this prediction, we used a dataset of complex EMG signals measured during directional reaching movements in humans as inputs to a model GCL. We find that physiological-like inputs produce a temporally varying basis set with parsimonious physiological thresholding and synchrony assumptions, which greatly improves the speed and accuracy of prediction in the model. By contrast, temporal feature extraction from time-invariant predictors requires additional mechanisms that have not been physiologically substantiated. We conclude that physiologically relevant inputs to the GCL engage robust temporal sparsening mechanisms that enhance associative learning that may explain diverse learning speeds across tasks. These findings point to information preprocessing by the GCL as a critical component of the learning algorithm implemented by the cerebellum.

II-91. Functional relevance of sensory- and task-related representations in posterior parietal cortex

Matthijs oude Lohuis
Pietro Marchesi
Umberto Olcese
Cyriel Pennartz

M.N.OUDELOHUIS@UVA.NL

P.MARCHESI@UVA.NL

U.OLCESE@UVA.NL

C.M.A.PENNARTZ@UVA.NL

University of Amsterdam

Animals need to construct representations of multimodal sensory inputs, and flexibly map these onto appropriate motor actions. Previous work has identified the posterior parietal cortex (PPC) as a key locus that integrates sensory inputs from different modalities together with contextual demands to support this flexible decision making. However, several recent studies disagree about two topics of major interest: (1) whether PPC is causally involved in decision-making depending on different sensory modalities (audition vs. vision), and (2) whether its involvement varies as a function of task structure (relying on memory, accumulation of evidence, conflict, etc.). Here we investigated whether PPC plays a causal role in the detection of auditory and/or visual stimuli when no evidence accumulation nor a short-term memory component are explicitly needed. To this aim, we trained head-fixed mice in a multisensory change detection task and recorded neural activity in PPC using multi-electrode laminar probes. We found that neurons in PPC respond to both auditory and visual features and show strong detection-related activity. Tensor decomposition and population decoding of the heterogeneous activity patterns revealed

that distinct neuronal populations in PPC encode the task-relevant auditory and visual features, and also the upcoming behavioral response. Surprisingly, however, optogenetic silencing of PPC had no effect on visual or auditory hit rates, whereas silencing of V1 impaired visual - but not auditory - detection. Although PPC thus showed robust encoding of sensory- and task-related variables, we found no immediate functional relevance to behavior. These results highlight the importance of cautiously interpreting the relevance of task-related neural activity. Conceivably, the PPC participates in a larger cortical network constructing contextual representations of multimodal inputs, but is not causally required to perform the change detection task. PPC activity might become behaviorally-relevant when trial-by-trial utilization of this contextual information is necessary.

II-92. Controlling processing speed via cell-type specific perturbations in spiking neural circuits

David Wyrick
Luca Mazzucato
University of Oregon

DWYRICK@UOREGON.EDU
LUCA.MAZZUCATO@GMAIL.COM

To thrive in dynamic environments, animals can generate flexible behavior and rapidly adapt responses to a changing context, goals to attain, and the animal's internal state. Behavioral flexibility may manifest as faster and more accurate stimulus responses to expected stimuli and during task engagement; or slower and less accurate responses in the presence of distractors. Experimental evidence suggests contextual modulations occur early in the cortical hierarchy and may be implemented via afferent projections to sensory cortex, induced by neuromodulatory pathways or top-down feedback from higher areas. However, the circuit mechanisms driving such modulations remain elusive. Here, we elucidate the effects of afferent perturbations on the speed and accuracy of sensory processing, using biologically plausible models of cortical circuit, based on recurrent spiking networks, where excitatory and inhibitory neurons are arranged in clusters. Network activity shows rich temporal dynamics, unfolding through sequences of metastable attractors even in the absence of external stimuli, capturing empirical distributions of firing rates and autocorrelation times. How do stimuli and perturbations interact with the attractor dynamics to shape responses? We considered four types of cell-type specific perturbations, targeting either excitatory/inhibitory interneurons, or GABAergic/glutamatergic synapses. We uncovered new paradoxical effects, such as improved performance via increased variability, originating from the complex interaction between perturbations and metastable attractor dynamics. Crucially, perturbations effects were explained in terms of how they affected the network's temporal dynamics; but they could not be predicted by traditional measures such as changes in stimulus selectivity or neural variability. This surprising result stands in stark contrast to E/I balanced circuits, where contextual modulations are accounted for by changes in selectivity and variability. Our results establish a theoretical framework for explaining contextual modulations of sensory processing and suggest new ways to dissect neural circuit architectures and dynamics via cell-type specific manipulation experiments.

II-93. A novel computational framework for the role of nucleus accumbens dopamine in information processing

Munir Kutlu
Jennifer Zachry
Patrick Melugin
Erin Calipari
Vanderbilt University School of Medicine

GUNES.KUTLU@VANDERBILT.EDU
JENNIFER.E.ZACHRY@VANDERBILT.EDU
PATRICK.R.MELUGIN@VANDERBILT.EDU
ERIN.CALIPARI@VANDERBILT.EDU

The reward prediction error (RPE) theory has been extremely successful in explaining how dopamine controls the formation of associations between stimuli and rewards. However, the RPE theory may prove insufficient as recent findings indicate that the role of dopamine during learning extends beyond the prediction of a reward.

Thus, the field needs a new framework incorporating existing theories on RPE with the newly emerging findings. Here, using an in-vivo fluorescent dopamine sensor (dLight1.1) combined with computational modeling and machine learning, we characterize the role of the dopamine input to the nucleus accumbens (NAc) within multiple information processing dimensions (e.g., value, novelty, outcome encoding). Our results demonstrate that current error prediction-based models are not capable of explaining the dopamine signal to cues and outcomes with negative values. Specifically, NAc dopamine peaks match the intensity of shock, a predicted negative outcome, and decrease with repeated presentations of shocks of the same intensity. This indicates that the dopamine signal follows the perceived saliency or novelty of outcomes. Importantly, we demonstrate that the dopamine signal is present when an expected shock is omitted. However, this peak is smaller than when the shock was present suggesting it signals an internal representation of the predicted outcome rather than unexpectedness of the event. We thus propose a new real-time neural network model, which combines error-prediction terms with novelty-based updating mechanisms and a bimodal outcome value term. Testing the prediction of our model, we find that NAc dopamine primarily encodes the “internal representation” of present as well as predicted but absent stimuli in the environment, which is independent of stimulus value, subsequent action, and prediction of future behavior. Together, our analysis offers a complete picture of the multifaceted action of NAc dopamine during learning that supersedes dopamine as solely encoding prediction error.

II-94. Task-selective connectivity improves rapid task switching in an interpretable model of SC

Sean Bittner¹
Carlos D Brody²
John Cunningham¹
Alex Piet³
Chunyu Duan⁴

SEANRBITTNER@GMAIL.COM
BRODY@PRINCETON.EDU
JPC2181@COLUMBIA.EDU
ALEX.PIET@GMAIL.COM
ANNDUAN2@GMAIL.COM

¹Columbia University

²Princeton University

³Allen Institute for Brain Science

⁴Chinese Academy of Sciences

Prefrontal cortex has been traditionally thought of as the brain region most central to executive processing. However, optogenetic inactivation of the superior colliculus (SC) in rats rapidly switching between tasks based on a contextual cue suggests that SC is also a principal contributor to flexible routing. Recordings in this experiment also motivated the definition of functional neuron types, resulting in a dynamical model that was used to reason about task responses. Recently, machine learning research has produced a methodology that can identify distributions of biologically interpretable model parameters consistent with high-level features describing neuronal computation, like rapid task switching accuracy. We used this methodology to gain an intuitive understanding of how this model of SC executes this task, and produced experimentally testable predictions about changes in effective connectivity in SC throughout learning of the behavioral paradigm.

II-95. An anatomically and functionally constrained model of direction selectivity in *Drosophila*

Jacob Portes
Jessica Kohn
Larry Abbott
Rudy Behnia
Columbia University

JPP2139@COLUMBIA.EDU
JRK2177@COLUMBIA.EDU
LFA2103@COLUMBIA.EDU
RB3161@COLUMBIA.EDU

A flurry of recent experimental studies have upended the canonical “correlation” model for invertebrate motion detection. While the models proposed by these studies can robustly predict direction selectivity, they rely on inputs that have not yet been demonstrated to be part of circuit anatomy. Furthermore, these models do not incorporate direct measurements of cellular inputs to the system. Combining newly acquired electrophysiological data with connectomic constraints, we construct a novel, parsimonious model of direction selectivity in the *Drosophila* OFF visual pathway. We performed *in vivo* whole-cell patch-clamp recordings of cells presynaptic to the direction selective T5 cell in *Drosophila*. In addition to extracting recordings with high temporal accuracy, we note a diversity of responses, including (1) previously unreported ON-OFF depolarization in a T5 input, and (2) rapid temporal adaptation in the presence of the neuromodulator octopamine (OA). We combined these observations with connectomic data to build a state-dependent, linear-nonlinear-type model of T5 based only on its presynaptic inputs. This biologically constrained model reconciles experimental data from Wienecke et al. (2018) and Gruntman et al. (2019), and suggests a biologically grounded framework for motion detection in invertebrates - a critical step in understanding computational principles of motion processing across visual systems.

II-96. LOOPER: A tool for the semi-supervised extraction of behaviorally relevant dynamics from observations of neural data

Connor Brennan
Alex Proekt
University of Pennsylvania

SHARSNIK@GMAIL.COM
PROEKT@GMAIL.COM

Animal and human behavior is encoded in neuronal activation patterns. Yet, neuronal activity varies dramatically from trial to trial. Furthermore, because of nonlinearities in neuronal dynamics, simply averaging activity across trials does not always yield representative results. This variability arises because there are many variables that are not directly observable, but nonetheless influence neuronal activity. These hidden variables span the latent space in which the neural dynamics unfold. While techniques exist to approximate the latent space, this in and of itself is not conceptually revealing, as the latent space is typically high-dimensional. Nevertheless, trajectories traced by the nervous system in the high-dimensional space can be low dimensional objects well approximated by simple and conceptually revealing models. Here we present LOOPER, a GUI software for the modeling and visualization of low dimensional dynamics extracted from high-dimensional noisy neural recordings and other multivariate stochastic nonlinear dynamical systems. We treat the nervous system as a stochastic dynamical system and estimate a discrete approximation of the time evolution operator of this system using diffusion maps. The chief innovation here is the introduction of asymmetric distances between points in phase space. This allows us to model the temporal evolution of the system. Diffusion maps are then clustered into loops. The manifold is thus parameterized by only two variables: loop identity and phase along each loop. Thus, dynamics are succinctly approximated by a system of loops with transitions between them. This description allows for both quantitative modeling and qualitative understanding of the underlying dynamics. This approach has been verified by modeling whole brain activity in *C. elegans*. We validate LOOPER on several dynamical systems including a trained recurrent neural network to demonstrate how a qualitative understanding of the dynamics informs the behavioral strategy used by the system to solve a task.

II-97. Normative evidence-accumulation and decision-making in foraging

Zachary Kilpatrick¹
Jacob Davidson²
Ahmed El Hady^{3,4}

ZPKILPAT@COLORADO.EDU
JDAVIDSON@AB.MPG.DE
ATAREK1@GMAIL.COM

¹University of Colorado, Boulder

²Max Planck Institute of Animal Behavior

³Princeton University

⁴HHMI

Foraging tasks provide a way to study how ecological and evolutionary forces have shaped neural circuits underlying decision-making (Mobbs et al 2018). It is unknown under what conditions animals can learn and efficiently exploit environmental statistics to optimize foraging decisions, and how the brain processes information to make such decisions.

We consider a canonical and realistic foraging task - the patch-leaving problem - and formulate a probabilistic inference approach to foraging decisions, where an ideal observer's belief is associated with a posterior over patch yield rates. We compare these decision rules with a recent evidence-accumulation model of patch-foraging decisions (Davidson & El Hady 2019) where different heuristic 'strategies' were suggested to be efficient based on the structure of the environment. We show that evidence-accumulation can be approximated by a drift-diffusion process, generating a normative account of the decision rules proposed in Davidson & El Hady (2019).

Our computational model is a baseline theory for designing experiments and interpreting data from animals freely foraging in patchy environments, and provides a framework to test neural and cognitive mechanisms underlying patch leaving decisions. The model results suggest that environmental parameters can be used to control observables like the animal's reaction (patch-leaving) times or valuation of food encounters as evidence, which could be observed by recording or imaging from decision-making, motor-planning, or reward-coding regions of the brain.

II-98. Learning sequences of correlated patterns in recurrent networks

Subhadra Mokashe
Nicolas Brunel
Duke University

SM713@DUKE.EDU
NICOLAS.BRUNEL@DUKE.EDU

What determines the format of memory representations in cortical networks is a subject of active research. During memory tasks, the retrieval of stored memories is characterized either by the persistent elevation in the firing rate of a set of neurons ('persistent activity') [Funahashi et al., 1989] or by ordered transient activation of different sets of neurons ('sequential activity') [Harvey et al., 2012]. Multiple theoretical studies have shown that temporally symmetric Hebbian learning rules give rise to fixed point attractor representation of memory (e.g., [Pereira and Brunel, 2018] and references therein), while temporally asymmetric learning rules lead to a dynamic sequential representation of memories (e.g., [Gillett et al., 2019] and references therein). These studies assume that inputs to the network during learning have no temporal correlations.

The sensory information received by brain networks is likely to be temporally correlated. We study temporally asymmetric Hebbian learning rules in a recurrent network of rate-based neurons in the presence of temporal correlations in the inputs and characterize how the inputs shape the network dynamics and memory representation using both numerical simulations and mean-field analysis. We show that the network dynamics depend on the temporal correlations in the input stream the network receives. For weakly correlated inputs, the network exhibits sequential activity, while for highly correlated inputs, the network settles into a fixed point attractor. We find that correlations increase the sequential memory capacity of the network. Non-linear learning rules increase the range of timescale of correlation for which the networks represent the memories as sequential activity in the network. We also show that the network maintains a sequential representation, both in the case of sequences of discrete patterns and in the continuum limit. Our work thus suggests that the correlation time scales of inputs at the time of learning have a strong influence on the nature of network dynamics during retrieval.

II-99. Dynamic long-range coordination patterns in macaque motor cortex

David Dahmen¹
 Moritz Layer¹
 Lukas Deutz²
 Paulina Anna Dąbrowska¹
 Nicole Voges³
 Michael von Papen¹
 Sonja Grun¹
 Markus Diesmann¹
 Moritz Helias¹

D.DAHMEN@FZ-JUELICH.DE
 M.LAYER@FZ-JUELICH.DE
 SCLD@LEEDS.AC.UK
 P.DABROWSKA@FZ-JUELICH.DE
 NICOLE.VOGES@GMX.COM
 MVONPAPEN@GMAIL.COM
 S.GRUEN@FZ-JUELICH.DE
 M.DIESMANN@FZ-JUELICH.DE
 M.HELIAS@FZ-JUELICH.DE

¹Research Centre Julich

²University of Leeds

³Aix-Marseille University

The cortex is a network of networks that is organized on various spatial scales. Short information pathways are realized by potential direct connections on the smallest scale, and by structured white matter connectivity on the largest scale. What about the intermediate, mesoscopic scale? On the scale of a single cortical area, for example the primary motor cortex in macaque monkey, the majority of connections is governed by connection probabilities that fall off with distance on characteristic length scales of a few hundred microns. Yet, in massively parallel recordings of motor cortex spiking activity in awake and resting macaques we find strongly correlated neurons almost across the whole Utah array, which covers an area of $4 \times 4 \text{ mm}^2$. Positive and negative correlations form salt-and-pepper patterns in space that are seemingly unrelated to the underlying short-range connectivity profiles. Whilst additional complex connection and input structures could potentially give rise to such patterns, we here show that the latter emerge naturally in a balanced network state close to instability: in such a state, interactions are mediated via a multitude of parallel paths through the network, giving rise to long-range correlations between individual neurons despite short-range connections. Using methods from statistical physics and disordered systems, we derive that the width of the covariance distribution decays exponentially at large distances with space constants much larger than the local connectivity profile. These predictions are directly confirmed by our experimental data. The found mechanism for long-range cooperation is not imprinted in specific connectivity structures, but rather dynamically results from the network state. Thereby, neuronal coordination patterns are not static, but dynamically changing over time in a state-dependent manner, as we demonstrate by comparing two different behavioral epochs of a reach-to-grasp experiment.

II-100. Higher visual areas contribute to visually guided detection behavior in the mouse

Hannah Goldbach¹
 Bradley Akitake¹
 Caitlin Leedy¹
 Mark H Histed²

HCGOLDBACH@GMAIL.COM
 BRADLEY.AKITAKE@NIH.GOV
 CAITLIN.LEEDY24@GMAIL.COM
 MARK.HISTED@NIH.GOV

¹National Inst. of Mental Health

²NIMH/NIH

Primary visual cortex (V1) in the mouse projects to a variety of brain areas, including several secondary/higher visual cortical areas (HVAs), frontal cortex, and basal ganglia [1,2]. While silencing V1 strongly impairs visually guided behavior [3,4], it is not known which, if any, higher visual cortical areas are involved in simple visual behaviors. We studied whether higher visual areas are involved in a contrast increment detection task (fig. 1). In such a simple behavior, animals might require V1 activity but not activity in higher visual areas. Instead, we find that there are several regions surrounding V1 that are involved in task performance, a lateral location (in visual areas RL, AL and LM), and a medial location near what has been termed parietal cortex. Behavioral

effects are specific to the retinotopic locations of activity evoked by the small Gabor stimulus animals detect. Inactivation of area PM does not affect performance. However, GCaMP imaging (fig. 2) shows that area PM does respond to the flashed Gabor stimulus, but with longer latency than responses in lateral areas (LM, AL, RL), perhaps explaining why PM inhibition produces little change in behavior. finally, we find the principal effect of HVA inhibition is to suppress the illuminated area, with little effect on V1. GCaMP imaging in V1 with and without HVA stimulation shows that our optogenetic inhibition has little effect on the visually evoked response of V1 ($\Delta dF/F0 = 0.09 \pm 0.08\%$ at 1mW, mean \pm SEM, N=3). Together, these results show that a network of cortical areas in mouse posterior cortex, beyond just V1, have activity that is used by the animal during performance of even simple contrast change visual behaviors.

II-101. Neural dynamics across brains of socially interacting animals

Lyle Kingsbury
Shan Huang
Jun Wang
Ken Gu
Peyman Golshani
Emily Wu
Weizhe Hong

LYLEKINGSBURY@GMAIL.COM
SHANHUANGYM@GMAIL.COM
WANGJ9669@GMAIL.COM
KENQGU@GMAIL.COM
PGOLSHANI@MEDNET.UCLA.EDU
YE.WU@UCLA.EDU
WHONG@MEDNET.UCLA.EDU

University of California, Los Angeles

Social interactions involve some of the most complex decision-making tasks that animals must navigate to secure their survival and success, as they require integration of internal state variables and contextual information with real-time decisions of others. To date, social neuroscience has mostly focused on behaviors in individual animals in order to interrogate the neural circuit computations that underlie social interactions. But a full understanding of the social brain, and how it is engaged during complex, dyadic interactions, requires a broader picture that reflects the dynamic nature of real-time social engagement. A major open question is whether and how emergent properties of neural systems, such as interbrain synchronization, may arise across brains of individuals to shape their interaction.

Here, we used dual microendoscopic calcium imaging to investigate neural dynamics in the prefrontal cortex of two mice engaged in free social interactions. We find that interacting animals exhibit synchronized brain activity in multiple social contexts that emerges from single cells in each that encode their own behavior and that of their social partner. Interbrain synchrony also predicts future social interactions as well as dominance relationships during social competition, suggesting a functional link between multi-animal neural dynamics and the real-time evolution of dyadic encounters.

II-102. Transient chaotic dimensionality expansion by recurrent networks

Christian Keup
Tobias Kuhn
David Dahmen
Moritz Helias
Research Centre Julich

C.KEUP@FZ-JUELICH.DE
T.KUEHN@FZ-JUELICH.DE
D.DAHMEN@FZ-JUELICH.DE
M.HELIAS@FZ-JUELICH.DE

Cortical neurons communicate with spikes, which are discrete events in time. Functional network models often employ rate units that are continuously coupled by analog signals. Is there a benefit of discrete signaling? By a unified mean-field theory for large random networks of rate and binary units, we show that both models have identical second order statistics. Their stimulus processing properties, however, are radically different: We discover a chaotic submanifold in binary networks that does not exist in rate models. Its dimensionality increases

with time after stimulus onset and reaches a fixed point that depends on the synaptic coupling strength. Low dimensional stimuli are transiently expanded into higher-dimensional representations that live within the manifold. We find that classification performance peaks when stimulus dimensionality matches the submanifold dimension; typically within a single neuronal time constant. Classification shows a high resilience to noise that exceeds rate models by orders of magnitude. Our theory mechanistically explains all these observations.

These findings have several implications. 1) Optimal performance is reached with weaker synapses in discrete state networks compared to rate models; implying lower energetic costs for synaptic transmission. 2) The classification mechanism is robust to noise, compatible with fluctuations in biophysical systems. 3) Optimal performance is reached when each neuron in the network has been activated only once; this demonstrates efficient event-based computation with short latencies. 4) The presence of a chaotic sub-manifold has implications for the variability of neuronal activity; the theory predicts a transient increase of variability after stimulus onset. Our theory thus provides a new link between recurrent and chaotic dynamics of functional networks, neuronal variability, and dimensionality of neuronal responses.

II-103. Foveal action for the control of extra-foveal vision

Ziad Hafed
Xiaoguang Tian
Tong Zhang

ZIAD.M.HAFED@CIN.UNI-TUEBINGEN.DE
TXGXP88@GMAIL.COM
ZHANGTONGDORA@HOTMAIL.COM

University of Tuebingen

Visual processing is frequently interspersed with saccades, which are associated with strong peri-movement changes in neural and perceptual sensitivity. With fixed gaze, for highly controlled experiments on visual and cognitive processing like covert attention (Krauzlis et al., 2013), tiny microsaccades still occur. During the past two decades, it became clear that microsaccades are not random (Hafed and Clark, 2002; Engbert and Kliegl, 2003) but instead exhibit highly predictable correlations with enhanced or decreased peripheral visual sensitivity (in multiple brain areas; Chen et al., 2015) and perception (Hafed, 2013; Tian et al., 2016). These time-locked changes appear at eccentricities 1-2 orders of magnitude larger than the eye movement endpoints themselves. However, it remains unclear whether it is microsaccades, perhaps through their motor preparatory activity, that causally influence visual sensitivity, or whether visual sensitivity is itself modulated independently of microsaccades, maybe through oscillatory brain-state fluctuations; in this case, it is such modulations that “leak” into the motor system and trigger microsaccades. Here, motivated by the idea that each microsaccade is purposeful (Willeke et al., 2019), we tested the former hypothesis. We used real-time retinal image stabilization to introduce tiny foveal motor errors at fixation (0.03-0.18 deg) and causally drive microsaccades in an experimentally-controlled direction. We then presented peripheral (>5 deg) visual stimuli congruent or incongruent with microsaccade direction and recorded monkey superior colliculus (SC) activity. In both monkeys and humans, we also probed behavioral performance (reaction time in monkeys; perceptual contrast sensitivity in humans). We found enhanced visual sensitivity and perceptual performance for microsaccade-congruent stimuli as a result of causally-generated microsaccades. Thus, foveal motor activity is sufficient to influence peripheral visual sensitivity. Theoretically, this implies that lateral connections in foveally-magnified visual-motor structures are dictated by anatomical coordinates, rendering foveal and extra-foveal eccentricities functionally much closer to each other than in visual coordinates.

II-104. Using noise to probe network structure and prune synapses

Rishidev Chaudhuri
University of California, Davis

RCHAUDHURI@UCDAVIS.EDU

Many networks in the brain are sparsely connected, and the brain prunes synapses during development, learning and, perhaps, sleep. Determining how the brain finds and maintains sparse network structure is important both

to understanding the brain's remarkable energy efficiency (and replicating it in artificial neural networks) and to understanding changes in connection density in disease and across the lifespan. How could the brain decide which synapses to prune? In a recurrent network, determining the importance of a synapse between two neurons is a hard computational problem, depending on the role that both neurons play and on all possible pathways of information flow between them.

Noise is ubiquitous in neural systems, and often considered an irritant to be overcome. Here we suggest that noise could play a functional role in synaptic pruning, allowing the brain to probe network structure and determine which synapses are redundant, by comparing the pattern of noise-driven correlations to the strength of a synapse. We construct a simple local anti-Hebbian plasticity rule that prunes synapses using only synaptic weight and the noise-driven covariance of the neighboring neurons. The plasticity rule is task-agnostic—it seeks only to preserve certain properties of network dynamics. Thus, it could act alongside learning or during separate pruning epochs (e.g., sleep), and does not restrict the learning rule in any way.

We prove that for a class of linear networks the pruning rule preserves multiple useful properties of the original network (including resting-state covariance and the spectrum), even when the fraction of removed synapses approaches 1. The proof leverages powerful recently-developed concentration of measure techniques that are finding increasing use in computer science and quantum information theory. These tools are useful to analyze large networks and should be of independent interest to neuroscience.

II-105. Automatic fitting of spiking network models to neuronal activity reveals limits of model flexibility

Shenghao Wu¹
Chengcheng Huang²
Adam Snyder³
Matthew Smith¹
Brent Doiron²
Byron Yu¹

SHENGHAW@ANDREW.CMU.EDU
HUANGC@PITT.EDU
ADAM.SNYDER@ROCHESTER.EDU
MATTSMITH@CMU.EDU
BDOIRON@PITT.EDU
BYRONYU@CMU.EDU

¹Carnegie Mellon University
²University of Pittsburgh
³University of Rochester

A major goal of theoretical neuroscience is to develop network models that can reproduce key aspects of neuronal activity recorded in the brain. A given model typically reproduces some, but not all, of the desired features of neuronal recordings. We term the ability of a network model to match different statistics of neuronal data the model flexibility. To characterize model flexibility and improve existing network models, we need 1) incisive measures to compare neuronal recordings with the activity produced by network models, and 2) automatic methods to fit network model parameters to data to understand where existing models fall short. To address these needs, we propose to use population activity statistics based on dimensionality reduction that goes beyond the commonly-used single-neuron and pairwise statistics. Furthermore, we develop a Bayesian optimization algorithm to efficiently fit model parameters to data. We then used our algorithm to study which aspects of neuronal activity recorded in macaque V4 can be reproduced by classical balanced networks (CBN) and spatial balanced networks (SBN). We found that CBNs are less flexible compared to SBNs in fitting population activity from visual area V4 and discovered a strong trade-off between population statistics and single-neuron Fano factors, thereby revealing limits on the model flexibility. These insights can be used to guide the development of future network models whose activity resembles neuronal recordings even more closely.

II-106. Development of negative prediction-error neurons in the presence of three inhibitory interneuron types

Loreen Hertaeg
Sprekeler Henning
TU Berlin

LOREEN.HERTAEG@TU-BERLIN.DE
H.SPREKELER@TU-BERLIN.DE

Internal predictions strongly influence how we perceive the world. To make these predictions accurate, they need to be compared to sensory inputs to detect mismatches, which are then used to refine our inner model of the world. Recent advances have given new insights into the neuronal circuitry that underlies the computation of those mismatches [1-4]. It has been suggested that in V1, L2/3 pyramidal cells (PCs) can serve as negative prediction-error (nPE) neurons by simultaneously processing excitatory feedback and inhibitory feedforward input at their distal dendrites. The required dendritic inhibition is thought to be mediated by somatostatin-expressing (SOM) interneurons [4]. How mismatch responses are influenced by other interneuron types and how experience-dependent plasticity could shape inhibitory circuits to support mismatch responses, however, is not resolved.

To investigate constraints on cortical interneuron circuits that allow PCs to act as nPE neurons, we analyzed a network model comprising PCs, parvalbumin-expressing (PV), SOM and vasoactive intestinal peptide-expressing (VIP) interneurons. We show that nPE neurons can only emerge if PV interneurons are inhibited sufficiently by both SOM and VIP interneurons. We furthermore show that the circuit can learn to compare internal predictions with sensory input by means of plasticity of synapses onto PV neurons and from SOM neurons onto the apical dendrites of PCs. While a simple homeostatic plasticity rule suffices when PCs do not receive feedforward input at the soma, a non-local credit assignment problem emerges when PCs also receive sensory inputs. Finally, we demonstrate that the proportion of nPE neurons in a heterogeneous network can be controlled by the inputs onto SOM and VIP neurons.

In summary, we show how nPE neurons can emerge in a network comprising three types of interneurons, an important step toward a biologically-founded circuit theory of predictive coding.

II-107. Hexadirectional coding of decision trajectories through abstract and discrete spaces

Seongmin Park
Douglas Miller
Erie Boorman
University of California, Davis

IMRAVICK@GMAIL.COM
DSMILLER@UCDAVIS.EDU
EDBOORMAN@UCDAVIS.EDU

Recent findings suggest the hippocampal-entorhinal (HC-EC) system may serve a general mechanism for representing and navigating cognitive maps of non-spatial tasks. These map-like representations can be used to guide flexible goal-directed decisions. However, it is unclear whether this system uses the same principles to map the dimensions of abstract and discrete decisions in the absence of continuous sensory feedback, such as whom to collaborate with. We developed a novel task to address this question. During training, human participants learned the status of people in independent 2-dimensional (2-D) social hierarchies piecemeal from binary comparisons. The overall structure of the social network has never shown to participants. In fMRI, participants were asked to choose the better collaborator for an individual in the social network between two by comparing their 'growth potential (GP),' which corresponded to the area drawn by the relative rank of two individuals in the 2-D ability space. First, the dissimilarity in the pattern of HC and EC activity increase with the pairwise Euclidean distance between individuals in the 2-D social network, suggesting that social dimensions learned piecemeal are integrated into a 2-D cognitive map. Second, the EC, medial prefrontal cortex (mPFC), temporoparietal junction (TPJ), and posteromedial cortex all display hexa-directional signals for the direction of inferred trajectories over the cognitive map that guide decisions. The grid-like signal was specific for six-fold periodicity and consistent across sessions more than a week apart. Finally, the HC, EC and ventromedial prefrontal cortex (vmPFC) encode

the decision-value, suggesting their roles in comparing values of decision options from the multidimensional cognitive map. Our findings show that a grid-like code in the human brain is extended to encode inferred trajectories of decision-making over an abstract and discrete cognitive space, which suggests a general mechanism for how the brain implements model-based decision-making.

II-108. Diverse behavioral strategies and neural coding during a visual change detection task

Alex Piet
Nick Ponvert
Marina Garrett
Anton Arkhipov
Yazan Billeh

ALEXPJET@GMAIL.COM
NICK.PONVERT@ALLENINSTITUTE.ORG
MARINAG@ALLENINSTITUTE.ORG
ANTONA@ALLENINSTITUTE.ORG
YAZANB@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

How are representations in sensory cortex influenced by immediate stimuli, behavioral context, and task relevant goals? To investigate neural coding during behavior we developed a systematic pipeline to train mice to perform a visual change detection task during two-photon calcium imaging ($n = 32$ mice, 192 sessions). Head fixed mice are shown serially flashed images. When the image identity changes, the mice must lick to receive a water reward. Using a time-varying regression model, we quantify the influence of several strategies on each mouse's behavior on an image by image basis. A diverse range of strategies were required to fit mouse behavior on this task. Individual mice vary their strategies and performance within each session. A single dimension of task-centric to timing-centric behaviors accounts for 80% of the variance across mice. Our calcium imaging dataset spans cell-type specific driver lines (3 cre-lines) across several visual areas during both active behavior and passive viewing. V1 cells are modulated by task demands, responding strongly to image changes during active, but not passive, behavior. We have quantified mouse behavior during a visual change detection task, while simultaneously recording a large standardized neural dataset.

II-109. An inference perspective on urgency in decision-making: A drunkard's walk case study

Maximilian Puelma Touzel¹
Raymond Chua²
Yoshua Bengio³
Paul Cisek¹
Guillaume Lajoie¹

MAX.PUELMATOUZEL@GMAIL.COM
RAYMOND.CHUA@MAIL.MCGILL.CA
YOSHUA.BENGIO@MILA.QUEBEC
PAUL.CISEK@GMAIL.COM
LAJOIE@DMS.UMONTREAL.CA

¹Universite de Montreal

²McGill University

³Montreal Institute for Learning Algorithms

Agents are often tasked with deciding early to maximize reward rate [1]. Evidence accumulation-to-bound [2] and urgency-gating models [3] each fail to wholly reproduce experimental results of single tasks, but each bring useful ingredients. Here, we provide an intuitive theory of time-constrained decision-making, combining these ingredients in the context of the well-known, yet under-studied 'tokens task', which challenges agents to use prediction to capitalize on early decisions. We support the theory with the development and analysis of a solution by a neurally plausible reinforcement learning (RL) algorithm, by an interpretable optimal solution, and with a qualitative match to measured neural recordings for urgency and commitment time from non-human primates [4] and human behavioral data [3]. Our approach offers 3 novel attributes. First, the agent employs a compressed representation of future trajectories, inspired by the recently proposed successor state representation in hippocampus [5], as a pow-

erful and learnable balance between model-based and model-free RL approaches. Second, the agent exploits its evidence accumulation model to compute a real-time posterior estimate of a trial's value that we find decays with time, introducing a bias towards high value trials. We show that hyperbolic discounting, the prevalent form found in primates, emerges naturally if the agent's memory is limited to storing mean values and uses a least-biased prior. Third, the resulting urgency signal in our model is an estimate of trial difficulty and has a particularly simple form that combines intrinsically generated subjective confidence, experimentally imposed time pressure, and time linearly. These dependencies are also exhibited by the urgency signal measured in basal ganglia, providing neural evidence for our formulation. Finally, we use a variant of Q-learning, a neurally plausible learning algorithm, on the task and achieve primate-level performance. Our formulation of time-constrained decision-making provides an experimentally-grounded inference perspective on reward-based learning.

II-110. Excitatory and inhibitory tone shapes decision regimes in hierarchical neural networks to maximize reward across environments

Warren Pettine¹
 Kenway Louie¹
 John D Murray²
 Xiao-Jing Wang¹

WARREN.PETTINE@GMAIL.COM
 KL837@NYU.EDU
 JOHN.MURRAY@YALE.EDU
 XJWANG@NYU.EDU

¹New York University
²Yale University

People are constantly faced with decisions between alternatives defined by multiple attributes, where the true value of each attribute is at times clear, and other times heavily obscured. How such attribute information is processed and integrated by the brain is controversial. Some contend that a single pathway performs these operations[1], while others hold that is distributed among multiple processing streams[2]. We investigated how neural network structure defines behavior in a multi-attribute, multi-alternative decision task, with varying levels of environmental noise. By tuning the levels of excitation and inhibition (E/I) in the intermediate processing layer of the hierarchical network, we found that the network would enter distinct decision-making regimes. One regime functioned in line with the single-pathway hypothesis, and another regime performed according to the distributed-decision hypothesis. Specific decision-regimes were found to be optimal depending on the level of noise in the environment. We then used these principles to investigate a leading etiological theory of autism (E/I imbalance), and predict that the decision patterns and overall performance will differ from neurotypicals as environmental noise levels increase.

[1]Padoa-Schioppa, C. & Conen, K. E. Orbitofrontal Cortex: A Neural Circuit for Economic Decisions. *Neuron* 96, 736–754 (2017). [2]Cisek, P. Making decisions through a distributed consensus. *Current Opinion in Neurobiology* 22, 927–936 (2012).

II-111. A cognitive map in the primate entorhinal cortex for mental navigation

Sujaya Neupane
 Mehrdad Jazayeri

SNEUPANE@MIT.EDU
 MJAZ@MIT.EDU

Massachusetts Institute of Technology

A 'cognitive map' is thought to organize the underlying structure of a contingent task in a unitary space. Neural correlates of the allocentric representation of space have been discovered in rodent hippocampus pointing towards a spatial map based on Euclidean geometry. However, it is not known how the nervous system represents cognitive maps and utilizes them to make flexible inferences. To address this question, we designed a novel task for monkeys in which they had to use a joystick to travel mentally (without sensory feedback) between two arbitrary points on a one-dimensional map. On each trial, animals were presented with a start image and a target image

chosen randomly from a previously learned ordered sequence of images. Animals had to use their memory of the sequence to infer the position of the target relative to the start image, and move the joystick in the correct direction and for the appropriate length of time to move from the start image to the target. Importantly, the entire process of traversing the image sequence had to be performed mentally as the images were not presented during the movement. Single neurons in the entorhinal cortex (EC) were strongly modulated during the mental travel and exhibited two broad firing rate patterns. Some neurons had temporally sparse firing rates with multiple peaks similar to grid-cell activity observed during spatial navigation. Others had robust ramping activity modulated by distance. At the population level, the neural states representing the images were ordered similarly to the sequence. Furthermore, the Euclidean distances between image pairs in the neural space exhibited systematic biases that were present in the animal's behavioral reports. These findings are consistent with EC establishing a low-dimensional embedding of task structure that can be used as a cognitive map for mental navigation.

II-112. Economic choices with simultaneous or sequential offers rely on the same neural circuit mechanism

Weikang Shi¹
Sebastien Ballesta²
Camillo Padoa-Schioppa¹

SHI.WEIKANG@WUSTL.EDU
BALLESTA@UNISTRA.FR
CAMILLO@WUSTL.EDU

¹Washington University in St Louis
²University of Strasbourg

Economic choices entail assigning values to the available options; a decision is then made by comparing values. Studies in which monkeys chose between two juices identified in the orbitofrontal cortex (OFC) different groups of cells encoding the value of individual options (offer value), the binary choice outcome (chosen juice) and the chosen value. These variables capture both the input and the output of the choice process, suggesting that these groups of neurons constitute the building blocks of a decision circuit. Experimental and theoretical results support this hypothesis. However, current notions emerged from studies where offers were presented simultaneously; conversely, it was suggested that choices between goods offered (or attended) sequentially rely on different mechanisms. Ballesta et al. recently examined OFC activity during economic decisions under sequential offers. They found different groups of cells seemingly analogous to those identified under simultaneous offers: group1 cells encoded the value of one particular juice whenever that juice was offered; group2 cells encoded the value of the juice currently offered and, subsequently, the chosen value; group3 cells encoded in a binary way the juice type present on the monitor and, subsequently, the chosen juice. Here we directly examined the relation between these cell groups and those previously identified under simultaneous offers. Monkeys performed economic decisions alternating between the two task modalities in a trial-by-trial way. Our results indicate that decisions under sequential offers and under simultaneous offers rely on the same neural circuit.

II-113. Suboptimal evidence integration reveals adaptive, nonlinear mechanisms in spatial vision

Takahiro Doi
Johannes Burge
University of Pennsylvania

DOI.TAKAH@GMAIL.COM
JBURGE@SAS.UPENN.EDU

How humans integrate ambiguous and conflicting evidence has been a focus of perception and decision-making research for decades. Identifying the source of suboptimality can reveal the information integration strategy adopted by the nervous system to solve ecologically relevant tasks in natural environments. Linear receptive fields (i.e., integration windows) are optimal for certain fundamental tasks, including target detection and spatial averaging. Behavioral suboptimality in these tasks has often been ascribed to improperly shaped integration

windows or to internal noise. Here we report that both factors together are insufficient to account for observed evidence integration performance in simple laboratory visual averaging tasks. Instead, the observed suboptimality is consistent with adaptive, stimulus-by-stimulus nonlinear integration mechanisms. Human observers judged the average luminance or average stereoscopic depth of spatially variable test stimuli relative to a reference surface; multiple levels of spatial variability were examined. The experimental design allowed the ideal observer to perform the task without error at all levels of spatial variability. The optimal strategy is to simply average the stimulus across space. Human accuracy, however, deteriorated with increased spatial variability, even though the human integration window estimated by reverse correlation was near optimal. Together, these findings violate the predictions of standard models, ruling out linear mechanisms that are fixed across stimuli. Next, we examined whether a stimulus-specific, nonlinear integration strategy could account for the data. Observers made consistent choices across repeated presentations of identical stimuli that could only be accounted for with an adaptive, stimulus-specific mechanism that is sensitive to the spatial layout of the relevant information: choice consistency decreased substantially when the spatial locations of the stimulus values were randomized. These nonlinear mechanisms, while suboptimal in our laboratory task, may reflect an integration strategy that improves performance in more natural contexts cluttered with object boundaries and illumination variation.

II-114. Time-dependent decision urgency in dynamic reward and evidence-quality conditions

Nicholas Barendregt¹

Kresimir Josic²

Zachary Kilpatrick¹

NICHOLAS.BARENDREGT@COLORADO.EDU

JOSIC@MATH.UH.EDU

ZPKILPAT@COLORADO.EDU

¹University of Colorado, Boulder

²University of Houston

In a constantly-changing world, making accurate and rewarding choices requires adaptive evidence accumulation and decision strategies that respect the dynamic nature of the environment. A foraging animal must decide whether to continue exploiting their current territory or risk exploring new areas for potentially better alternatives. To make these decisions effectively, they must account for the possibility of unexpected environmental changes (e.g., in food availability). There is growing interest in understanding how subjects make decisions in such dynamic contexts, with most work focusing on descriptive models of tasks where the strength of observed evidence changes between trials. For instance, the "urgency-gating model" (UGM) is a neurally plausible instantiation of a collapsing decision bound model derived under a sequential sampling assumption. To better frame theories of time-dependent decision criteria in human subjects, we extend normative Bayesian strategies to "dynamic-context" environments, where the reward or evidence quality for a correct choice changes within a trial according to a discrete-state Markov process. Using a Bayesian framework and dynamic programming principles we find optimal decision thresholds for ideal observers that maximize their average reward per unit time. This results in rich threshold dynamics that depend strongly on task parameters: thresholds can collapse as in the UGM, exhibit non-monotonic behavior when context change times are known in advance, or jump discretely when context changes in a Markovian fashion. Using these thresholds, we leverage recently developed, efficient solution methods for dynamic drift-diffusion models [3] to compute response time statistics of the ideal observer. Our results can aid in designing more informative experiments, and provide a touchstone for interpreting behavioral data from subjects making decisions in dynamically-changing environments.

II-115. From single neurons to behavior in the jellyfish *Aurelia aurita*

Fabian Pallasdies
 Sven Goedeke
 Wilhelm Braun
 Raoul-Martin Memmesheimer
 University of Bonn

FABIANPALLASDIES@GMAIL.COM
 SGOEDEKE@UNI-BONN.DE
 WBRAUN@UNI-BONN.DE
 RM.MEMMESHEIMER@UNI-BONN.DE

Understanding how neuronal activity leads to behavior in animals is a central goal in neuroscience. Since jellyfish are anatomically relatively simple animals with a limited behavioral repertoire, modeling their nervous system opens up the possibility to achieve this goal. Further, such modeling provides insight into the origins of nervous systems, as both their taxonomic position and their evolutionary age imply that jellyfish resemble some of the earliest neuron-bearing, actively-swimming animals. Here we develop the first neuronal network model for the nerve nets of jellyfish. The present study focuses on the neuro-muscular control of the swimming motion in a moon jelly *Aurelia aurita* in its medusa stage of development. Incorporating the available experimental observations and measurements, we develop a bottom-up multi-scale computational model of single neurons, nerve nets, steered muscles and the bell. The proposed single neuron model disentangles the contributions of different currents to a spike. The network model identifies factors ensuring non-pathological activity and suggests an optimization for the transmission of signals. Using fluid-structure hydrodynamics simulations, we then explore how the nervous system generates and shapes different swimming motions. Activating the two known nerve nets of the bell, we find that different delays between their activations can lead to well-controlled, differently directed movements. Our model bridges the scales from single neurons to behavior, allowing for a comprehensive understanding of jellyfish motor control.

II-116. Mixture of poisson models for neural correlations and coding

Sacha Sokoloski¹
 Amir Aschner¹
 Ruben Coen-Cagli²

SACHA.SOKOLOSKI@EINSTEIN.YU.EDU
 AASCHNER@MAIL.EINSTEIN.YU.EDU
 RUBEN.COEN-CAGLI@EINSTEIN.YU.EDU

¹Albert Einstein College of Medicine
²Albert Einstein College Of Medicine

Theories of neural coding describe how neural responses encode information about stimuli. In particular, correlations in neural population responses have been shown to significantly impact information processing in theoretical neural circuits. Although modern parallel recording techniques can measure neural correlations in the brain, it has proven challenging to validate theories of their effects on biological neural circuitry due to a lack of models that both are compatible with these theories and can be efficiently fit to data.

To address this challenge we developed a novel model based on conditional finite mixtures of independent Poisson distributions. The model captures neural correlations by mixing component distributions, and the number of components can be cross-validated to match the complexity of target correlations. The model also affords closed-form expressions for both its density function and fisher information, which are critical for evaluating its coding properties. Moreover, we derived an expectation-maximization algorithm to fit the model to data from large neural populations, and have validated it extensively on synthetic data.

We demonstrate that our model successfully captures noise correlations in the responses of macaque V1 neurons to oriented gratings. We then show that estimates of the fisher information of recorded V1 neurons based on our model compare favourably with established methods. finally we show that a multilayer neural network cannot distinguish real V1 data from data synthesized by a conditional mixture model with a sufficient number of components. In conclusion, our framework should allow researchers to quantitatively assess the optimality of information processing in various biological neural circuits.

II-117. Understanding encoding and redundancy in grid cells using partial information decomposition

Praveen Venkatesh
Pulkit Grover
Carnegie Mellon University

VPRAVEEN@CMU.EDU
PULKIT@CMU.EDU

The brain is capable of performing reliable computation using neural circuitry that has high internal variability; this can only be possible if the brain controls for this variability through redundancy. Understanding such mechanisms will likely require quantifying the degree of uniqueness, redundancy and synergy that exists in the neural code. Recently, the literature in information theory has developed a structured framework called Partial Information Decomposition (PID) [Bertschinger et al. 2014; Lizier et al., 2018], which defines the unique, redundant and synergistic interactions between two or more variables in conveying information about a particular message. While the idea of decomposing mutual information into unique, redundant and synergistic components has roots in neuroscience [Schneidman et al., 2003], it has not received significant attention therein since these recent developments. In this abstract, we show two instances in which the PID provides insights into neural encoding, using computational models of grid cells based on prior work by Sreenivasan and fiete [2011]. (1) We identify redundant information in the grid cell code, which we associate with error correction capability, and identify redundant and unique information when error correction is in effect. (2) We find that determining the animal's location at a coarse spatial resolution requires the aggregation of synergistic information from just as many grid cell networks, as required to determine location at a much finer resolution: this might have implications for downstream networks that utilize grid cell outputs [fiete et al., 2008]. More generally, our results suggest that computing the PID in different settings might provide fine-grained insights into neural coding and function that are harder to attain with simpler correlation-based analyses. For instance, recent work by Venkatesh et al. [2019] explicitly connects the PID with a formal definition for information flow in the brain.

II-118. Isolating attention components in superior colliculus: perceptual sensitivity, criterion and motor bias

Supriya Ghosh¹
John Maunsell²
¹University of Chicago
²The University of Chicago

SGHOSH5@UCHICAGO.EDU
MAUNSELL@UCHICAGO.EDU

Like those in visual cerebral cortex, neurons in the superior colliculus (SC) of monkeys are strongly modulated by attention. It has been a point of contention whether attention-related enhancement of neuronal activity in the SC is associated with selective increase in behavioral d' in neurons' receptive fields or is instead associated with changes in the animal's decision threshold, which can be closely linked with motor planning. We trained rhesus monkeys in a novel visual orientation-change detection task that precisely dissociates perceptual sensitivity (d'), decision criterion and motor planning using delayed choice targets appearing in random locations far from the attended locations. Monkeys' attentional performance (d') and decision criterion at the neuron's receptive-field location were independently controlled, and sequentially switched between high and low levels within short blocks of trials by adjusting relative reward probabilities associated with stimuli in the left and right visual hemifields. Multiple SC neurons were recorded simultaneously while the animal did the task. Results from two animals showed that SC neurons increased their spiking rates when behavioral sensitivity (d') increased at their receptive field location. This attentional modulation was stronger for visuomotor neurons compared to purely visual neurons and motor neurons. Unlike the behavioral d' , change in decision criterion had no effect on spiking of these same SC neurons. These results confirmed that SC neurons contribute selectively to specific attentional states, and increased attentional performance can enhance spiking of SC neurons independent of decision criterion or motor planning towards the receptive field. This behavioral task has special significance for isolating distinct behavioral components associated with attention states in sensory-motor brain areas.

II-119. Formalizing a perceptual-mnemonic theory of the medial temporal lobe

Tyler Bonnen
 Dan Yamins
 Anthony Wagner
 Stanford University

TYLER.RAY.BONNEN@GMAIL.COM
 YAMINS@STANFORD.EDU
 AWAGNER@STANFORD.EDU

Within the mammalian brain, the Medial Temporal Lobe (MTL) is known to enable flexible (e.g. one-shot) memory-guided behaviors. Does the MTL also play a role in perception? This question has led to an enduring debate, where competing theories have been forced to rely on informal, descriptive accounts of stimulus properties (e.g. stimulus “complexity”). In order to formalize the perceptual demands on the MTL, here we combine meta-analytic, computational, and behavioral approaches, focusing on the role of Perirhinal cortex (PRC) in concurrent visual discrimination (“odddity”) tasks. We expect that PRC-lesioned behavior should reflect the pattern of behaviors supported by a linear readout of the primate Ventral Visual System (VVS). If a stimulus set is perfectly separable within the VVS (i.e. ceiling performance) then this experiment is not “diagnostic” of PRC involvement in perception. We formalize this null hypothesis for Perirhinal function by leveraging a stimulus-computable proxy for the primate VVS: task-optimized convolutional neural networks (CNNs). We first collect stimuli and behavior from published experiments identified by our meta-analysis. We observe that multiple experiments are not diagnostic of PRC’s involvement in perception; these studies should not be used to evaluate the involvement of PRC in perception, no matter their results. With the remaining studies, we observe a striking correspondence between model and PRC-lesioned behaviors. Moreover, PRC-intact behaviors demonstrate a significant improvement over what could be expected from a linear readout of the VVS. This relationship between human and model performance holds across multiple (feedforward and recurrent) architectures, as well as across models optimized to perform both object recognition and identity discrimination. This computational and meta-analytic approach offers converging evidence that Medial Temporal Lobe structures may be required for a certain class of perceptual behaviors.

II-120. The effect of boundaries on grid cell patterns

Mauro Miguel Monsalve Mercado¹
 Christian Leibold²

MM5461@COLUMBIA.EDU
 LEIBOLD@BIO.LMU.DE

¹Columbia University

²University of Munich

Grid cells form a most efficient code for space. However, geometrical distortions, defects, and environmental influences present in patterns of the local network question the universality of this code (Krupic et al., 2018, Science, 359). We argue that the existence of these observations is a natural result of the mechanisms leading to triangular patterns in grid cells. In particular, we show that feedforward models are a natural framework to study their development, since they implement a Turing pattern formation process directly over physical two-dimensional space. We found that in this class of models, during learning, the activation of spatially unstructured feedback inhibition in the grid cell network removes the inherent bias in the weight growth process induced by the presence of physical boundaries, which is indicative for Turing type pattern formation without inhibitory feedback. The model leads to the eventual formation of almost perfect triangular patterns in enclosures of arbitrary shape without any special treatment for the boundaries, such as periodic or fading boundary conditions. Moreover, the model produces a coherent population of grid cells sharing similar orientation, spacing, and field size. In addition, it predicts that grid cells asymptotically self-organize into three phase groups to cover the environment in an approximately uniform way. finally, the model gives a general account of how interaction among grid cells gives rise to the observed geometrical distortions of the patterns depending on the shape of the recording enclosure and premature termination of learning, and reproduce some of them via unprimed simulations.

II-121. Optimal dendritic processing of spatiotemporal codes

Brendan Bicknell
Michael Hausser
University College London

B.BICKNELL@UCL.AC.UK
M.HAUSSER@UCL.AC.UK

The transformation of patterns of synaptic input to action potential output is central to information processing in the brain. For every neuron, this operation is shaped by the morphology and biophysics of the dendritic tree. Synaptic potentials are subject to saturation and attenuation due to passive cable properties, yet active conductances and spatial interactions may compensate, or even enhance processing of particular input features. Dendritic integration, as the filter through which cells communicate, thus presents both fundamental constraints and additional opportunities for computation. How biological or artificial networks may capitalise on this layer of processing remains an open question. A key challenge is to quantify the interplay between synaptic input statistics and dendritic biophysics, and how these interactions can be shaped by learning. However, most previous work has been limited to restricted input regimes, or reduced models that lack the constraints of dendritic physiology. Here, we show that dendritic processing is maximally effective when information is embedded in both the spatial and temporal structure of physiological input patterns. We augment a model of a layer 2/3 pyramidal neuron with a set of variational equations that report how the somatic voltage depends on the history of input to the dendritic tree. We use this to construct an algorithm with which we train the model to perform a nonlinear feature-binding task. In the optimal regime, sparse bursts of synaptic input provide a substrate for computation far exceeding the capacity of purely rate or temporally coded input. When stimulus features are spatially segregated, performance depends critically on voltage-dependent NMDA receptors, predicting a mechanism for multimodal selectivity through summation of dendritic spikes.

II-122. Dimensionality control in the critical regime of balanced networks

Eric Shea-Brown¹
Stefano Recanatesi¹
David Dahmen²
Moritz Helias²

ETSB@UW.EDU
STEFANO.RECANATESI@GMAIL.COM
D.DAHMEN@FZ-JUELICH.DE
M.HELIAS@FZ-JUELICH.DE

¹University of Washington

²Research Centre Jülich

The dimensionality of a network's collective activity indicates the number of modes in which activity is organized. A small number of modes suggests a compressed, low dimensional neural code and possibly high robustness. In an orthogonal view, neural networks can be considered dynamical systems which undergo transitions between different regimes; it has been found that computational performance peaks right at those critical points. Here we use both theory and new, high-density neuropixel recordings to show that there is in fact a deep and universal relation that dictates a tight link between these two qualitatively distinct concepts of dimensionality and criticality. We first demonstrate this link theoretically: computations beyond the mean-field approximation allow us to express the participation ratio, a measure of the network's dimensionality, in terms of a single number, the spectral bound of the network's connectivity – the largest real part of its eigenvalues. The latter controls transitions between regular and chaotic dynamics. The theory explains why dimensionality can be controlled most effectively close to this transition point. Surprisingly, the theory also predicts that both measures of global collective information processing, the spectral bound and the participation ratio, sensitively depend on local features of the connectivity – predominantly on small connection motifs involving pairs and triplets of cells. Analysis of massively simultaneous *in vivo* recordings at single cell resolution via neuropixel probes confirms that various cortical and subcortical areas in mouse indeed operate close to the point of optimal dimensionality control. This suggests that dimensionality control could be accomplished via highly local synaptic learning rules in many brain regions.

II-123. Encoding object transformations across distinct areas of rodent visual cortex

Juliana Rhee
Cesar Echavarría
Javier Masis
David Cox
Harvard University

RHEE@G.HARVARD.EDU
CECHAVARRIA@G.HARVARD.EDU
JMASIS@G.HARVARD.EDU
DAVIDCOX@FAS.HARVARD.EDU

The ability to recognize objects despite the enormous variation with which they can appear (e.g., due to viewing angle, scale, or lighting) is a computational feat widely believed to be unique to the primate visual system. Recently, rats have been shown to perform invariant object recognition tasks, and a subset of areas in visual cortex (V1, LM, LI, LL) has been proposed as functionally analogous to the hierarchical pathway underlying object recognition in primate ventral visual cortex. Previous studies have used electrophysiology in anesthetized rats and resulting population-level analyses have largely relied on aggregating few simultaneously recorded cells across many sessions. The present study takes advantage of optical methods, which allow simultaneous access to multiple brain regions with single-cell resolution, and presents their first application to an animal model of invariant object recognition.

Awake, head-fixed rats expressing GCaMP in areas V1, LM, and LI were imaged using large-field two-photon microscopy while shown a battery of visual stimuli. Low-level response properties were tested with receptive field measurements and drifting gratings. High-level properties were tested with objects that varied through identity-preserving or identity-changing transformations. At the single-neuron level, relative differences in low-level properties between visual areas were in agreement with electrophysiology studies in primates and in anesthetized rats. However, population-level analyses of simultaneously recorded neurons revealed surprising differences from previous studies of how these visual areas may support object recognition. Object representations were most separable in V1 and least separable in LI, and view-specific information was present across all visual areas tested. In addition, when trained to classify objects at one view, V1 populations showed successful generalization to novel, untrained views. These results provide new insights into how rats may provide both an alternate and complementary model for invariant recognition.

II-124. The sensorimotor strategies and neuronal representations of whisker-based shape recognition in mice

Chris Rodgers
Ramon Nogueira
Esther Greeman
Stefano Fusi
Randy Bruno
Columbia University

CCR2137@COLUMBIA.EDU
RN2446@COLUMBIA.EDU
EG3026@COLUMBIA.EDU
SF2237@COLUMBIA.EDU
RB2604@COLUMBIA.EDU

Humans and other animals can identify objects by active touch, requiring the coordination of exploratory motion and tactile sensation. For instance, mice recognize objects by scanning them with their whiskers, just as we do with our fingertips. The brain must integrate these exploratory motor actions with the resulting tactile signals in order to form a holistic representation of object identity. To elucidate the behavioral strategies and neural mechanisms underlying this ability, we developed a shape discrimination task for head-fixed mice that challenged them to discriminate concave from convex shapes. As a control, we trained a separate group of mice on a simpler shape detection task that simply required them to detect the presence of either shape regardless of its curvature. To identify the behavioral strategies that mice used to perform these tasks, we measured a comprehensive suite of sensorimotor variables about whisker motion and touch, and trained classifiers to use these variables to predict the presented stimulus and the mouse's choice. This analysis revealed that mice compared the number of con-

tacts across whiskers during shape discrimination, whereas during detection they summed them across whiskers. We also recorded populations of neurons in the barrel cortex, which processes whisker input, to identify how it encoded the sensorimotor variables that mice use to recognize shape. We observed whisker-specific coding for motion (especially in deep inhibitory neurons) and touch (especially in superficial neurons). This whisker-specific pattern in the neural responses was similar to the weights that the behavioral classifier used to decode shape identity. We suggest a predictive coding model in which inhibitory whisker motion signals reformat contact responses so that shape identity may be more easily read out by downstream areas. More generally, a similar computation of comparing across multiple sensors may underlie object recognition in other brain areas and species.

III-1. Motor cortical representation and decoding of attempted handwriting in a person with tetraplegia

Frank Willett¹
Donald Avansino¹
Leigh Hochberg²
Jaimie Henderson¹
Krishna Shenoy¹

WILLETT2@GMAIL.COM
DAVANSIN@STANFORD.EDU
LEIGH_HOCHBERG@BROWN.EDU
HENDERJ@STANFORD.EDU
SHENOY@STANFORD.EDU

¹Stanford University

²Brown University

Handwriting is a fine motor skill in which straight and curved pen strokes are strung together in rapid succession. Because handwriting demands fast, richly varying trajectories, it could be a useful tool for studying how motor cortex generates complex movements. Additionally, attempted handwriting movements could be decoded by a brain-computer interface (BCI) and translated to text in real time, restoring the ability to communicate to people with severe paralysis. Here, we investigated the neural representation and decodability of attempted handwriting movements in a person (“T5”) paralyzed from the neck down (C4, ASIA-A spinal cord injury). We recorded from two microelectrode arrays in hand knob of precentral gyrus (a premotor area analogous to macaque PMd).

first, we tested whether we could neurally decode handwritten sentences. We recorded a dataset where T5 imagined writing 102 sentences in a self-paced manner. We then trained a recurrent neural network (RNN) to transform the neural data into character probabilities at each timestep. The RNN correctly labeled 95.6% of all characters on held-out data with no language model, generating understandable text at speeds exceeding prior records for intracortical BCIs (66 characters per minute vs. 39 correct characters per minute).

Next, to understand how motor cortex generates handwriting movements, we recorded neural activity while T5 imagined writing individual letters in an instructed delay paradigm. We analyzed the preparatory activity observed before each letter was written and found that it represented upcoming letter features only within a short time-horizon. That is, letters that differed only near the end (e.g. m and n) had seemingly identical preparatory representations. We also found that preparatory activity continuously re-engaged throughout the movement period. These results suggest that brain areas other than motor cortex orchestrate the generation of curved trajectories, and may feed input into motor cortex via preparatory dimensions throughout the movement.

III-2. Reinforcement regulates context-dependent timing variability in thalamus

Jing Wang¹
 Eghbal Hosseini¹
 Nicolas Meirhaeghe^{1,2}
 Adam Akkad¹
 Mehrdad Jazayeri¹

JINGWANG.PHYSICS@GMAIL.COM
 EHOSEINI@MIT.EDU
 NMRGHE@MIT.EDU
 AKKAD@MIT.EDU
 MJAZ@MIT.EDU

¹Massachusetts Institute of Technology

²Harvard-MIT Division of Health Sciences and Technology

While motor learning reduces variability, motor variability could facilitate learning. This paradoxical relationship has made it challenging to tease apart behavioral and neural signatures of variability that degrade performance from those that improve it. We tackled this question using a time interval production task involving multiple behavioral contexts; i.e., different motor effectors and different timing intervals. In both humans and monkeys, behavioral variability was governed by two context-specific processes: a slow process due to memory drifts across trials that degraded performance and a fast reinforcement learning process that modulated variability depending on the outcome of the preceding trial. We developed a generative Reward Sensitive Gaussian Process (RSGP) model that validated the presence of these two components of variability. The model also revealed that the fast reinforcement-dependent regulation of variability was an adaptive process whose function was to calibrate behavioral responses in the face of memory drifts. Recently, it has been shown that monkeys exert flexible control over the motor timing through adjustment of a tonic signal in higher-order thalamus that sets the speed at which activity in cortico-basal ganglia circuits evolve toward an action-triggering state. Accordingly, we hypothesized that the signals in the same region of the thalamus might be subject to adjustment of variability through reinforcement learning. Simultaneous recording from thalamic neurons validated this hypothesis: thalamic signals underwent context-specific slow fluctuations in register with the behavior, and their trial-by-trial variability was modulated by reward along the task-relevant dimension. These findings indicate that the nervous system counters memory drifts by calibrating thalamic input through reinforcement. More generally, our work provides an example of reinforcement acting upon neural variability to regulate exploratory behavior needed to improve performance.

III-3. Dynamic recombination of temporally coincident presynaptic modules underlies responses in cortex

Arjun Bharioke¹
 Stuart Trenholm²
 Adrian Wertz³
 Daniel Hillier¹
 Emilie Mace⁴
 Keisuke Yonehara⁵
 Balazs Rozsa⁶
 Karl-Klaus Conzelmann⁷
 Rava Azeredo da Silveira^{1,8}
 Botond Roska¹

A.BHARIOKE@IOB.CH
 STUART.TRENHOLM@MCGILL.CA
 ADRIAN.WERTZ@WEB.DE
 DANIEL.HILLIER@IOB.CH
 EMILIE.MACE@IOB.CH
 KEISUKE.YONEHARA@DANDRITE.AU.DK
 ROZSABAL@KOKI.HU
 CONZELMANN@GENZENTRUM.LMU.DE
 RAVA@IOB.CH
 BOTOND.ROSKA@IOB.CH

¹Institute of Ophthalmology Basel

²Montreal Neurological Institute

³Friedrich Miescher Institute

⁴Max Planck Institute of Neurobiology

⁵Aarhus University

⁶Hungarian Academy of Science

⁷Ludwig-Maximilians-University Munich

⁸Ecole Normale Supérieure

As the seat of sensory perception, multi-sensory integration, decision-making and other complex computations, the neocortex has been extensively studied. However, a detailed understanding of the spatiotemporal topology by which individual cortical neurons collect into modules that then contribute to the activity of other neurons, is still unknown. To understand how specific modules of presynaptic neurons are organized, we used single-cell-initiated monosynaptically-restricted rabies to express gCaMP6s transynaptically, in a single cortical neuron and its population of presynaptic input neurons. We then used a rapid random-access, acousto-optic-deflection 2-photon microscope to record, in vivo, simultaneously from both the output cell and its input population – sampling the response dynamics of an isolated local cortical circuit, for a variety of inputs. Repeating this over multiple circuits, we characterized both the stimulus response dynamics of individual neurons and their spontaneous activity. We observed that neurons within each presynaptic network show significantly higher spontaneous pairwise correlations than randomly sampled cortical neurons. This effect was increased for correlations between the presynaptic neuron and their postsynaptic partners, and was consistent across pairwise spatial distances. This suggests that temporal coincidence of inputs is actively sampled for in cortical circuits, and is an important property in these circuits. Identifying modules of correlated neurons within the presynaptic population, we observe the neurons within each module show similar responses to part of the full stimulus space. Examining postsynaptic responses, we observe that each coincides with responses in different combinations of functionally distinct presynaptic modules, suggesting that the temporal alignment provides a substrate for recombining different inputs in the output neuron, at different time points. Hence, our results suggest that functional modules of presynaptic neurons, defined by temporal coincidence, are dynamically recombined to generate cortical neuronal responses, and provide the first measurements of the organization of functional connectivity within individual cortical circuits in vivo.

III-4. Linking structure to function in a model of early color processing

Matthias Christenson
Sarah Heath
Maia Weisenhaus
Larry Abbott
Rudy Behnia
Columbia University

MPC2171@COLUMBIA.EDU
SLH2171@CUMC.COLUMBIA.EDU
MW3228@COLUMBIA.EDU
LFA2103@COLUMBIA.EDU
RB3161@COLUMBIA.EDU

Models of sensory processing have historically abstracted underlying biological circuits, due to unknown connectivity and/or complexity. In contrast, the use of tractable and anatomically well-characterized model organisms such as the fruit fly *Drosophila melanogaster* allows us to utilize biological constraints in models of sensory processing to understand underlying circuit mechanisms and make more accurate predictions. This approach has been used to dissect motion vision circuits, but investigations into color vision - a salient visual feature for many animals - have been limited. Here, we investigate the circuit mechanisms of the early color circuit of the fruit fly and assess its information processing capabilities. Using in vivo two-photon calcium imaging and genetic manipulations, we measure the chromatic tuning properties of photoreceptor outputs and propose a dual circuit mechanism that compares inputs within one point in space and across space. We built an anatomically constrained model that simulates our observed responses, while also incorporating underlying connectivity. Our model is able to quantitatively reproduce our experimental observations without fitting synaptic weights. Instead, we used electron-microscopy-derived synaptic count, an anatomically defined measure, as a proxy for synaptic weight, thereby linking structure to function. We were able to use our model to ask how this early color circuit affects downstream processing. We find that while early chromatic processing allows for a robust representation of color in downstream circuits, the biological constraints in our model appear to limit the range of this representation. This limit may reflect the behaviorally relevant chromatic processing requirements of the fruit fly. In summary, the detailed anatomy of an early color circuit predicts the chromatic tuning properties of its outputs and an anatomically constrained model allows us to infer the effects of early color processing on downstream chromatic encoding.

III-5. Internal models for active sensing in freely swimming electric fish

Avner Wallach
Nate Sawtell
Columbia University

AW3057@COLUMBIA.EDU
NS2635@COLUMBIA.EDU

Animals move their sensory organs to perceive the world. The resulting sensory activity reflects complex interactions between self- and externally generated components. It is widely hypothesized that the nervous system needs to distinguish between these components to accurately perceive the external world. One system where this has been studied is the electrosensory lobe (ELL) of weakly electric fish, where copies of motor commands, termed corollary discharge, are used to predict self-generated sensory input. These predictions, termed “negative images,” are integrated with afferent input to improve the sensory representation of external stimuli. However, past studies were largely restricted to simple, highly-stereotyped electromotor behavior in immobilized preparations. To examine how the negative image mechanism operates in active exploration we developed methods for tracking behavior and recording neuronal activity at the input, intermediate, and output layers of ELL in freely swimming fish, over multiple days. We found that during free exploration, sensory modulations due to the fish’s movements (e.g. tail bending) were strongly influenced by the external context, such as the position of the tank walls relative to the fish. Hence, the cancelation of these modulations required an internal model that can generalize to different external conditions. Remarkably, putative intermediate neurons conveying the negative image signal exhibited context-specific activity and output layer activity was largely motor-invariant. Since pure corollary discharge signals convey no information about the external world, these results imply that negative images must incorporate an additional source of input. Using an ELL model and Bayesian estimation theory we show that such input could be provided by fast sensory feedback. Our findings provide the first demonstration of negative images in freely behaving animals and suggest an essential role for feedback in the generalization of internal models to changing external conditions.

III-6. Multi-animal pose tracking and predictive modeling of social interactions

Talmo Pereira
Mala Murthy
Nathaniel Tabris
Shruthi Ravindranath
Junyu Li
Joshua Shaevitz
Princeton University

TDP@PRINCETON.EDU
MMURTHY@PRINCETON.EDU
TABRIS@PRINCETON.EDU
SHRUTHI@PRINCETON.EDU
JUNYULI@PRINCETON.EDU
SHAEVITZ@PRINCETON.EDU

The desire to understand how the brain generates and patterns behavior has driven rapid methodological innovation, in part enabled by the success of deep learning for computer vision applications, to quantify and model natural animal behavior. A first step in many of these analyses are recently-developed methods for deep-learning-based markerless pose estimation. Here we present two contributions to this domain: (1) a novel deep learning framework for multi-animal pose tracking with modular components to solve the sub-problems of region proposal, detection, grouping and tracking, flexibly customizable in any combination to specialize to a wide array of common experimental paradigms, including for social behavior; and (2) an agent-based artificial neural network model design that learns representations grounded in sensorimotor transformations and applies them to a large dataset of freely interacting pairs of fruit flies (*D. melanogaster*) to discover structure in their social dynamics.

III-7. The functional architecture of glutamatergic release onto CA1 neurons

Jason Climer
 Michael Adoff
 Heydar Davoudi
 Daniel Dombeck
 Northwestern University

JASON.R.CLIMER@GMAIL.COM
 MICHAELADOFF2018@U.NORTHWESTERN.EDU
 HEYDAR.DAVOUDI@NORTHWESTERN.EDU
 D-DOMBECK@NORTHWESTERN.EDU

To understand the computation done by neurons, it is critical to understand how they transform their presynaptic input into somatic output in the form of action potentials. For example, area CA1 of the hippocampus is necessary for spatial memory and hosts place cells, pyramidal neurons that fire as mice visit distinct locations in an environment. They are surrounded by silent and non-place cells, which may be place cells when the mouse is in a different context. Place cells receive inputs from spatial and non-spatial neurons in other regions, notably cells in hippocampal area CA3. In contrast to other regions containing neurons with well-defined receptive fields, there is no anatomical organization to the arrangement of the tuning of place cells. Thus, many models of place firing rely on random presynaptic inputs being amplified post-synaptically to generate selective firing. Advances in functional indicators and in vivo subcellular two-photon imaging provide means to examine the presynaptic inputs. Using resonant scanning two photon microscopy and the fluorescent glutamate sensor iGluSnFR.A184S, we mapped glutamate release onto the dendrites of CA1 pyramidal neurons with micron scale resolution and single vesicle sensitivity as mice navigated a linear track in virtual reality. By combining this with calcium imaging using the red-shifted calcium indicator jRGECO1a, we determined how excitatory inputs contribute to the type and tuning of somatic output in these neurons. Place neurons receive more spatially tuned inputs than their non-place peers, and this input arrives preferentially in the receptive field of the neuron. Surprisingly, this tuned input is organized along the dendrites, arriving in functional clusters along the dendrite. These data imply that postsynaptic strength alone does not explain place cell firing; presynaptic inputs are tuned and organized to determine if and where the cell will fire.

III-8. Tongue kinematics reveal cortex-dependent corrections as the mouse tongue reaches for, and misses, targets

Teja Bollu
 Brendan Ito
 Sam Whitehead
 Brian Kardon
 Mei Liu
 James Redd
 Jesse Goldberg
 Cornell University

TB387@CORNELL.EDU
 BSI8@CORNELL.EDU
 SAMCWHITEHEAD@GMAIL.COM
 BMK27@CORNELL.EDU
 ML2294@CORNELL.EDU
 REDD.JAMES@GMAIL.COM
 JHG285@CORNELL.EDU

Precise control of the tongue is necessary for drinking, eating, and vocalizing. Yet because tongue movements are fast and difficult to resolve, neural control of lingual kinematics remains poorly understood. We combine kilohertz frame-rate imaging and a deep learning based artificial neural network to resolve 3D tongue kinematics in mice performing a cued lick task. Cue-evoked licks exhibit previously unobserved fine-scale movements which, like a hand searching for an unseen object, were produced after tongue protrusions and were directionally biased towards remembered locations. Photoinhibition of anterolateral motor cortex (ALM) abolished the fine-scale movements, resulting in well-aimed but hypometric licks that missed the spout. Electrophysiological recordings in ALM identified neural correlates of cue-evoked licks and the component fine-scale movements at both single neuron and at the population level. Our results show that cortical activity is required for online corrections during licking and reveal novel, limb-like dynamics of the mouse tongue as it reaches for, and misses, targets.

III-9. Precise optical probing of perceptual detection in olfactory circuits

Jonathan Gill
Gilad Lerman
Hetince Zhao
Benjamin Stetler
Shy Shoham
Dmitry Rinberg
New York University

JVG219@NYU.EDU
GILADLERMAN@GMAIL.COM
HETINCE.ZHAO@NYULANGONE.ORG
BENJAMIN.STETLER@NYULANGONE.ORG
SHOHAM@NYU.EDU
DMITRY.RINBERG@NYULANGONE.ORG

Establishing causal links between patterns of neuronal activity and perception is crucial for understanding the function of sensory systems. Across sensory systems, stimulus evoked activity can vary between stimuli along several feature dimensions, including changes in firing rate and the timing of action potentials. A key question is which features of stimulus evoked activity (e.g. rate, synchrony, or latency) are used to guide behavior? In the mouse olfactory bulb, inhalation of different odors leads to changes in the set of neurons activated, as well as when neurons are activated relative to each other (synchrony), and the onset of inhalation (latency). Here, we probed the sensitivity of mice to perturbations across each stimulus dimension using holographic two-photon (2P) optogenetic stimulation of olfactory bulb neurons, with cellular and single action potential resolution and millisecond precision. We found that mice can detect single action potentials evoked synchronously across <20 olfactory bulb neurons. Mice exhibited this sensitivity for artificial ensembles of mitral cells (excitatory projection neurons, mean threshold 10.3 ± 2.6 , $n=3$), as well as mixed ensembles of mitral and granule cells (excitatory and inhibitory neurons, mean threshold 13.9 ± 5.7 , $n=3$). Further, we discovered that detection depends strongly on the synchrony of activation across neurons, with detectability falling to near-chance levels with an imposed stimulus jitter ≥ 20 ms ($p < 0.05$, fisher's exact test, $n=3$), while detection performance was minimally perturbed by changes in the latency of activation relative to inhalation ($p=0.75$, Kruskal-Wallis test, $n=3$). These results reveal that mice are acutely attuned to single neurons and action potentials in olfactory circuits, and that synchrony across neurons may be a critical feature supporting the perceptibility of sparse ensemble activity signals.

III-10. Neural mechanisms of selection in visual working memory

Matthew Panichello
Timothy Buschman
Princeton University

MATT.PANICHELLO@GMAIL.COM
TBUSCHMA@PRINCETON.EDU

Working memory is critical to cognition and, yet, it has a severely limited capacity. We mitigate this capacity limitation by carefully controlling the contents of working memory. For example, when holding multiple items in working memory, we can 'select' one of them (to the exclusion of others) to guide behavior (Myers, Stokes, and Nobre, 2017). However, the neural mechanisms underlying this selection process are unknown. To investigate the neural basis of selection, we trained monkeys to perform a visual working memory task that required them to hold two items in working memory and then, after a cue, select one of them for a behavioral response. Using large-scale neural recordings in prefrontal, parietal, and visual cortex, we found that prefrontal cortex plays a leading role in directing selection. Furthermore, these control signals matched the signals used for directing attention, suggesting a common neural code for cognitive control of both internal and external representations. Finally, we found that both selection and attention bias competitive dynamics between stimulus representations in prefrontal and visual cortex. Together, our results provide novel insight into the neural mechanisms that control the contents of working memory.

III-11. Multiple overlapping hypothalamus-brainstem circuits drive rapid threat avoidance

Matthew Lovett-Barron¹
 Ritchie Chen¹
 Susanna Bradbury¹
 Aaron S Andalman¹
 Mahendra Wagle²
 Su Guo²
 Karl Deisseroth¹

MATTLB@STANFORD.EDU
 RITCHEN@STANFORD.EDU
 SBRADBUR@STANFORD.EDU
 ANDALMAN@STANFORD.EDU
 MAHENDRA.WAGLE@GMAIL.COM
 SU.GUO@UCSF.EDU
 DEISSERO@STANFORD.EDU

¹Stanford University

²University of California, San Francisco

Animals avoid environmental changes that threaten homeostasis in order to promote survival. The hypothalamus mediates such behaviors in vertebrates, but it is unknown whether different threats act through separate or overlapping hypothalamic populations. We address these issues in zebrafish, by examining multiple hypothalamic cell types during avoidance responses to homeostatic threats: increased heat, salinity, or acidity. Using cellular registration of large-scale neural activity recordings to multiplexed in situ gene expression, we find that distinct threats drive multiple peptidergic cell types across the hypothalamus with no one-to-one correspondence between activity patterns and neuropeptide gene expression. Cellular characterization and manipulations revealed that multiple peptidergic cell types (including cells releasing oxytocin, corticotropin-releasing hormone, and/or vasopressin) are predominantly glutamatergic, play shared roles in rapid avoidance behavior, and converge onto a common set of brainstem premotor neurons required for avoidance. These results demonstrate that the rapid avoidance of environmental conditions that threaten homeostasis – an essential survival behavior - is robustly driven by multiple hypothalamic cell types that converge in the brainstem.

III-12. Action selection in continuous time and space in the basal ganglia

Sean Escola
 James Murray
 Columbia University

GSE3@COLUMBIA.EDU
 JM4347@COLUMBIA.EDU

While action selection has traditionally been conceived as a discrete process, there are many situations in which a continuous array of possible actions might be chosen successively at arbitrary points in time. Motivated by new experiments involving learned forelimb movements in mice, we model action selection in continuous time and space using a biophysically realistic model arm driven by a neural network constructed to emulate the cortico-basal ganglia architecture. Similar to experimental data, the simulated behavior undergoes selection and refinement through training. Investigating the activity of simulated basal ganglia neurons from the trained network, we find that it exhibits coactivation among neural populations during movement, mixed selectivity for behavioral variables among neurons, encoding of both action selection and vigor, and dopamine-like reward prediction error signals. These results are congruous with several previously published and hitherto unexplained results from the basal ganglia literature, while also providing new predictions to be tested in future experiments.

III-13. The role of TRN for feedback-enhanced surround suppression in mouse dLGN

Gregory Born¹
 Sinem Erisken²
 Felix Schneider¹
 Agne Klein²
 Martin Spacek¹
 Milad Mobarhan³
 Davide Crombie¹
 Chu Lan Lao¹
 Gaute Einevoll^{4,3}
 Laura Busse¹

BORN@BIO.LMU.DE
 SINEM.ERISKEN@GMAIL.COM
 FELIX.SCHNEIDER@CAMPUS.LMU.DE
 AGNE.KLEIN@GMAIL.COM
 MSPACEK@MM.ST
 M.H.MOBARHAN@IBV.UIO.NO
 D.CROMBIE@CAMPUS.LMU.DE
 CHULAN.LAO@MED.UNI-MUENCHEN.DE
 GAUTE.EINEVOLL@NMBU.NO
 BUSSE@BIO.LMU.DE

¹Ludwig-Maximilians-University Munich

²University of Tubingen

³University of Oslo

⁴Norwegian University of Life Sciences

Neurons in the dorsolateral geniculate nucleus (dLGN) of the thalamus are suppressed by stimuli extending beyond the classical receptive field (RF) into the surround. The mechanisms underlying surround suppression in dLGN are difficult to disentangle since dLGN relay cells receive converging inputs from excitatory retinal ganglion cells, local inhibitory neurons, and corticothalamic (CT) feedback of both signs emerging from layer 6 (L6) of primary visual cortex (V1). To investigate the effect of CT feedback on surround suppression in dLGN, we combined extracellular single-unit recordings in head-fixed mice with network modeling.

In dLGN of PV-Cre mice, we probed responses to drifting gratings of different sizes while disrupting CT feedback on half of the trials by optogenetically activating V1 parvalbumin positive cells. We found that CT feedback promotes surround suppression in dLGN by enhancing responses to small stimuli and reducing responses to large stimuli.

Because L6 CT neurons are excitatory they can inhibit relay cells only indirectly via geniculate interneurons or inhibitory neurons in the thalamic reticular nucleus (TRN). To explore via which route feedback enhances surround suppression in dLGN, we simulated dLGN size tuning curves with the enhanced difference of Gaussians model (Mobarhan et al., 2018). Manipulating the spatial scale of the model's inhibitory feedback component revealed that it needed to be at least twice the size of the inhibitory feedforward component to capture experimental results.

We next measured RFs in the visual part of TRN and dLGN and found that RFs in TRN were significantly larger. Finally, probing size tuning in TRN of PV-Cre mice while optogenetically manipulating CT feedback showed that CT feedback substantially enhances responses in TRN, in particular for large stimulus sizes.

In summary, both model predictions and experimental results suggest that CT feedback sculpts surround suppression in dLGN likely via recruitment of inhibitory neurons in TRN.

III-14. A theory of learning with a constrained weight distribution

Ben Sorscher¹
 Weishun Zhong²
 Daniel D Lee³
 Haim Sompolinsky⁴

BSORSCH@STANFORD.EDU
 WSZHONG@MIT.EDU
 DDL46@CORNELL.EDU
 HAIM@FIZ.HUJI.AC.IL

¹Stanford University

²Massachusetts Institute of Technology

³Cornell University

⁴The Hebrew University of Jerusalem and Harvard University

The hypothesis that synaptic plasticity underlies learning and memory has spurred extensive experimental investigation into the structure and distribution of synaptic weights in the brain, and yielded powerful theoretical models of learning. Early theoretical studies used techniques from statistical physics to calculate quantities of fundamental interest, such as capacity and generalization, in simple learning models like the perceptron. However, these theories often predict the optimal distribution of synaptic weights is Gaussian, in apparent conflict with numerous anatomical studies which demonstrate that synaptic weights (1) are predominantly either excitatory or inhibitory, and don't change sign during learning, and (2) closely follow a lognormal distribution in cortical layers 2, 3, 5, hippocampus, and cerebellum. The distribution of synaptic weights thus imposes a fundamental biological constraint on learning, and calls for a theory that can be matched with these precise experimental measurements.

Recent work has investigated perceptron learning under specific weight distributions, such as binary or sign-constrained weights. Here, we derive a general theory of perceptron learning under a generic synaptic weight distribution, including the lognormal distribution. Our theory shows that constraining the weight distribution reduces the perceptron capacity by a simple factor with interesting information-theoretic and geometrical interpretations. Further, we propose a novel learning algorithm, based on the theory of optimal transport, allowing us to optimize linear classifiers with prescribed weight distributions. Finally, our theory predicts that if a student perceptron has a prior on the distribution of a teacher perceptron's weights, the student can substantially improve generalization performance by incorporating that prior as a constraint. Simulations with the optimal transport algorithm confirm this prediction. These novel, surprisingly elegant, interpretable results provide insight into the computational consequences imposed by biological constraints on the distribution of synaptic weights, and offer a theoretical framework for learning that can be matched with known experimental results.

III-15. Latent-state models reveal a state-dependent contribution of the striatum to decision-making

Iris Stone
 Scott Bolkan
 Ilana Witten
 Jonathan Pillow
 Princeton University

ISTONE@PRINCETON.EDU
 SBOLKAN@PRINCETON.EDU
 IWITTEN@PRINCETON.EDU
 PILLOW@PRINCETON.EDU

A central problem in systems neuroscience is to understand the relationship between sensory stimuli, neural activity, and sensory decision-making behavior. Traditional models assume that these relationships are fixed. Here we develop a latent-variable model to show that—to the contrary—time-varying internal states play a key role in modulating the effects of both sensory inputs and neural activity on decisions. Our approach consists of a Hidden Markov Model (HMM) with discrete latent states that correspond to different decision-making strategies. Decisions in each state are governed by a generalized linear model (GLM) with state-specific weights, which describe how task variables and neural activity are combined to generate a choice. We applied this model to experimental data from mice performing a two-choice decision-making task in virtual reality. In this new experiment, mice received optogenetic inhibition on a random 20% of trials targeting one of two principle output pathways of the dorsomedial striatum (DMS): the D1-receptor expressing direct pathway (D1R) or the D2-receptor expressing indirect pathway (D2R). In conducting our analysis, we first used cross-validation to compare models with different numbers of states and found that the a 3-state GLM-HMM exhibited an 18% improvement in log-likelihood compared to a null model. This meant that a typical session (~200 trials) within the test dataset of 25,705 trials was 16.6 times more likely under the GLM-HMM than the standard GLM. Furthermore, by inspecting the identified latent states, we discovered that not all decision-making strategies are equally dependent on striatal activity. Indeed, the effects of DMS inhibition were much larger in one of the three states, implying that the contribution of DMS to decision-making depends on internal state. These results reveal the complex, state-dependent relationship between striatal activity and decision-making in mice and have major implications for the study of decision-making behavior in other brain areas.

III-16. Remembrance of things practiced: A two-pathway circuit for sequential learning

James Murray
Sean Escola
Columbia University

JM4347@COLUMBIA.EDU
GSE3@COLUMBIA.EDU

The intertwined effects of practice and recency on memory and learned behavior have long been studied in psychology and neuroscience. We develop a bottom-up, mechanistic theory of learning and memory addressing the combined effects of recency and practice, beginning with a classical single-neuron model and applying insights from this model to the network level and to the implementation of learned motor behaviors in the brain. Combining error-based and associative learning, we mathematically derive the forgetting curve for a single neuron and its dependence on practiced repetition, showing how highly practiced memories or behaviors can become far more resistant to being overwritten by later learning. At the network level, error-based learning is responsible for initial gains in performance, while associative learning gradually transfers control of the downstream population from one input pathway to the other. We interpret the model neurobiologically by identifying the inputs as cortex and thalamus, and the downstream population as striatum. This model is consistent with recent experiments in rodents showing that execution of learned skills requires thalamic inputs to striatum but not cortical inputs. More broadly, it provides a framework for understanding the neural basis of habit formation and the automatization of behavior through practice.

III-17. Local and global organization of synaptic inputs on cortical dendrites

Jan Hendrik Kirchner
Julijana Gjorgjieva
Max Planck Institute for Brain Research

JAN.KIRCHNER@BRAIN.MPG.DE
GJORGJIEVA@BRAIN.MPG.DE

Synaptic inputs on cortical dendrites are organized with remarkable subcellular precision at the micron level. This organization emerges during early postnatal development through patterned spontaneous activity and manifests both locally where nearby synapses tend to share functional properties, and globally with distance to the soma. Recent experimental studies reveal species-specific differences in this organization between mouse, ferret and macaque visual cortex: While in the mouse, synapses are retinotopically organized along the proximal-distal axis of the dendrite, no such organization is present in the ferret or macaque. Instead, here synapses are organized into local clusters according to orientation preference. We propose a computational framework that combines activity derived from retinal waves with functional and structural plasticity to generate these different types of organization across species, as well as across scales by including attenuating backpropagating action potentials. Within this framework, a single anatomical factor – the size of the visual cortex and the resulting magnification of visual space – can explain the observed differences. This allows us to make predictions about the organization of synapses also in other species and indicates that the proximal-distal axis of a dendrite might be central in endowing a neuron with its powerful computational capacities.

III-18. Probing the neural substrates of movement across the rodent behavioral repertoire

Jesse Marshall¹
 Diego Aldarondo²
 William Wang¹
 Tim Dunn³
 Gordon Berman⁴
 Bence Olveczky¹

JESSE_D_MARSHALL@FAS.HARVARD.EDU
 DIEGOALDARONDO@G.HARVARD.EDU
 W_WANG@COLLEGE.HARVARD.EDU
 DUNN.TW@GMAIL.COM
 GORDON.BERMAN@EMORY.EDU
 EMBERBENCE@GMAIL.COM

¹Harvard University

²Harvard

³Duke University

⁴Emory University

Mammals are generalists, capable of flexibly re-using the same behaviors across a range of environmental contexts. However, existing experimental paradigms in mammals currently lack the behavioral diversity and kinematic resolution needed to examine which features of movement the brain encodes and how this encoding varies across contexts. To extend the range of behaviors and contexts that can be studied, we developed a new behavioral monitoring system, CAPTURE, that combines motion capture and deep learning to track the 3D movements of twenty points on a freely-behaving rat's trunk and appendages, continuously over weeks-long timescales in naturalistic environments. CAPTURE has sub-millimeter and millisecond-timescale spatial and temporal resolution, and exhibits substantial gains in precision over existing 2D convolutional network approaches to behavioral tracking. We developed a comprehensive behavioral analysis platform alongside CAPTURE that allowed us to detect known and novel behaviors, behavioral sequences, and states in our continuous kinematic recordings. This allowed us to comprehensively phenotype behavioral changes following administration of drugs, striatal lesions, and in animal models of disease. We then combined CAPTURE with continuous neural recordings in the dorsolateral striatum, a brain region with a known, if debated, role in controlling diverse aspects of movement and behavior. Within individual behavioral states, striatal neurons showed tuning to kinematic and behavioral variables. However across states, these tuning properties changed, in a manner that improved the within-state but not across-state behavioral decodability. This suggests that striatal representations undergo state-specific remodeling in order to facilitate motor coding, and reinforces the need to interpret the neural correlates of behavior in a context-specific manner.

III-19. flygenectors: large-scale dynamics of internal and behavioral states in a small animal

Evan Schaffer
 Neeli Mishra
 Wenze Li
 Matthew Whiteway
 Jason Freedman
 Kripa Patel
 Venkatakaushik Voleti
 Liam Paninski
 Larry Abbott
 Elizabeth Hillman
 Richard Axel
 Columbia University

ESS2129@COLUMBIA.EDU
 NM2786@COLUMBIA.EDU
 WL2351@COLUMBIA.EDU
 M.WHITWAY@COLUMBIA.EDU
 JWIZ9191@GMAIL.COM
 KBP2115@COLUMBIA.EDU
 VENKATA.VOLETI@GMAIL.COM
 LMP2107@COLUMBIA.EDU
 LFA2103@COLUMBIA.EDU
 EH2245@COLUMBIA.EDU
 RA27@COLUMBIA.EDU

Internal states, such as hunger and arousal, elevate the probability of a set of behaviors and persist on longer timescales than the behaviors that they predict. Recent evidence suggests that internal states are represented

throughout cortex in rodents and in many neuropil regions in *Drosophila*. Here, we ask at the single-neuron level how internal states are represented throughout the brain and on what timescales neural activity predicts behavior in *Drosophila*. We use SCAPE (Swept Confocally Aligned Planar Excitation) microscopy to perform calcium imaging throughout the entire brain with single-cell resolution at speeds exceeding 10 brain volumes/second. We explore two internal states, arousal, in flies freely running on a spherical treadmill, and hunger, in food-deprived flies consuming sugar. We define internal state as neural activity that predicts behavior on long timescales. To determine the timescale with which individual neurons best predict behavior, we define a regression model in which the activity of each cell is proportional to behavior filtered with unique time constant τ . In freely running flies, we observe that most neurons across the brain are highly correlated with locomotion. While the median timescale is short (8.8s), the distribution of timescales across all cells is broad, with some neurons correlated with locomotion on a much longer timescale ($\tau > 60$ s), putatively representing arousal. Similarly, in flies consuming sugar, we identify distributed ensembles of neurons possessing either a short timescale ($\tau < 10$ s), likely representing reward, or a long timescale ($\tau > 60$ s), putatively representing hunger or satiety. Finally, we predict the causal role of different neural ensembles in behavior by examining sequences in subspaces of the dynamics. Thus, we have identified dimensions of global neural dynamics (flygenvectors), including subspaces predicting behavior on long timescales that may support internal states of arousal and hunger.

III-20. Studying neural codes of odor perception using a novel behavioral paradigm in mice

Hirofumi Nakayama¹
Dmitry Rinberg²

NAKAYH01@NYUMC.ORG
RINBERG@NYU.EDU

¹New York University Neuroscience Institute

²Neuroscience Institute, New York University Langone Health

Similar sensory objects should evoke similar neural activity responsible for encoding these objects. Such a rule can be a powerful guiding principle for studying sensory processing. However, this approach requires understanding of the structure of perceptual spaces, which can be more easily achieved in humans, while neural codes are usually measured in animals. Here we develop a behavioral paradigm to measure perceptual distances in mice, propose an approach to study principles of neural coding by comparing perceptual and neural distances, and apply it to study olfactory coding. To measure perceptual distances across many odors we trained mice to report if two sequentially presented odors are the same or not. We assume that the perceptual distance between two odors is proportional to the probability of mice categorizing them as different ones. We test that such behavioral readout is consistent with predictions about perceptual distances between odor mixtures and with chemical similarities between odors. Using this behavioral paradigm we constructed a matrix of perceptual distances for many odors and applied it to study neural codes. We measure neural activity in response to the same set of odors at the level of glomeruli, and on the next level, mitral/tufted cells, using Ca imaging (wide-field and 2P). We constructed multiple pairwise distance matrices based on different features of the code and compared them with the perceptual distance matrix. We found that OB neuronal activity during earlier periods from odor onset is more relevant to odor perception. This result is consistent with the previous behavioral studies, but the first time demonstrated in a much more general behavioral paradigm, not limited to binary odor discrimination. Our approach can be extended to other brain areas and sensory systems and allows us to find both the structure of perceptual space in animals and perceptually relevant neural codes.

III-21. Developmental and evolutionary principles of olfactory circuit designs

Naoki Hiratani
Peter Latham

N.HIRATANI@UCL.AC.UK
PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, University College London

A fundamental goal of neuroscience is to understand the design principle of neural circuits. That is difficult to do in complete generality, so here we ask a slightly simpler question: can we predict the number of cells in various layers of a neural circuit by optimizing performance? We address this question in the context of the olfactory circuit, which shows clear scaling laws: in mammals, the number of layer 2 neurons in piriform cortex is proportional to the number of glomeruli to the $3/2$ power (Srinivasan & Stevens, 2019), while in insects, the number of Kenyon cells is roughly cubic in the number of glomeruli (literature survey). We model the olfactory system as a three-layered nonlinear neural network, and analytically derive scaling laws by estimating the network size that optimizes, over the lifetime of the animal, its ability to predict rewards associated with odors. Although having more neurons increases the capacity of the network, having too many makes developmental tuning of synaptic weights difficult due to overfitting. Applying this tradeoff, we find that the optimal population sizes follows the scaling law observed in mammals. This scaling is robust: it holds in full batch optimization and stochastic gradient learning, and under various choices of nonlinearity. We extend the framework to the case when a fraction of the olfactory circuit can be genetically specified, not developmentally learned, and numerically demonstrate that when there are a small number of glomeruli, this makes the scaling steeper, as is observed among insects. This suggests that it may be possible to understand neural circuit from an optimality point of view, with optimization that occurs over evolutionary timescales.

III-22. Natural gradient learning for spiking neurons

Elena Kreutzer
Mihai A Petrovici
Walter Martin Senn
University of Bern

ELENA-KREUTZER@GMX.DE
PETROVICI@PYL.UNIBE.CH
SENN@PYL.UNIBE.CH

Due to their simplicity and success in machine learning, gradient-based learning rules represent a popular choice for synaptic plasticity models. While they have been linked to biological observations, it is often ignored that their predictions generally depend on a specific representation of the synaptic strength. In a neuron, the impact of a synapse can be described using the state of many different observables such as neurotransmitter release rates or membrane potential changes. Which one of these is chosen when deriving a learning rule can drastically change the predictions of the model. This is doubly unsatisfactory, both with respect to optimality and from a conceptual point of view. By following the gradient on the manifold of the neuron's firing distributions instead of one that is relative to some arbitrary synaptic weight parametrization, natural gradient descent provides a solution to both these problems. While the computational advantages of natural gradient are well-studied in ANNs, its predictive power as a model for *in vivo* synaptic plasticity has not yet been assessed. By formulating natural gradient learning in the context of spiking interactions, we demonstrate how it can improve the convergence speed of spiking networks. Furthermore, our approach provides a unified, normative framework for both homo- and heterosynaptic plasticity in structured neurons and predicts a number of related biological phenomena.

III-23. Algorithmic basis of foraging decisions

Sergey Shuvaev¹
 Sarah Starosta¹
 Duda Kvitsiani²
 Adam Kepecs¹
 Alexei Koulakov¹

SSHUVAEV@CSHL.EDU
 SSTAROST@CSHL.EDU
 KVITSI@DANDRITE.AU.DK
 KEPECS@CSHL.EDU
 KOULAKOV@CSHL.EDU

¹Cold Spring Harbor Laboratory
²Aarhus University

In ethologically relevant conditions, we often face sequential stay-or-leave decisions about whether to engage with the current option, or to search for a better one. These decisions include employment, time investment, partner selection, etc. To model such decisions in mice, behavioral ethology uses foraging paradigms, in which animals decide when to leave depleting food resources. To explore the trial-by-trial choice strategy and its neural correlates, we designed a foraging task for mice and performed recordings of dopamine neuron activity in the Ventral Tegmental Area (VTA) – a brain area involved in reward-guided behaviors (Schultz et al, 2015). We compared the observed mouse behavior to a range of computational models and show that individual trial-by-trial decisions rely on many previous rewards. We identified a novel decision rule that accounts for animals' behavior in a range of conditions, including different inter-trial delays, reward depletion rates, and random restarts. Specifically, we found that the animals compare the next expected reward to the exponential average of the previous rewards – the decision rule we named the leaky MVT, for some similarity with the conclusions of the Marginal Value Theorem (MVT; Charnov, 1976; Constantino & Daw, 2015). This behavior strategy can be explained by the Bayes optimality of exponential filtering in dynamic environments discovered previously (Yu & Cohen, 2009; Piet et al, 2018). We further show that the leaky MTV decision rule may be learned via a reinforcement learning (RL) paradigm called R-learning (Schwartz, 1993), but is not consistent with classical V-, or Q-learning paradigms. finally, we show that dopaminergic signaling in the VTA best correlates with the Reward Prediction Error (RPE) of R-learning, pointing to the potential learning mechanism that optimizes stay-or-leave choices. Overall, our work links the MVT, RL and Bayesian approaches to offer a decision rule, learning mechanism and computational rationale for stay-or-leave decisions.

III-24. Efficient routing of information to high-dimensional neural representations

Samuel Muscinelli¹
 Ashok Litwin-Kumar^{1,2}
¹Columbia University
²Neuroscience

SPM2176@COLUMBIA.EDU
 AK3625@COLUMBIA.EDU

A dominant idea in neuroscience is that task-relevant variables are encoded by high-dimensional, distributed neural representations. Experiments have demonstrated the importance of such representations for task performance, and models argue that numerous brain areas, from neocortex to cerebellum, have connectivity capable of producing them. A standard simplifying model assumption is that these areas receive random, uncorrelated input, but a widespread connectivity motif consists of "bottlenecks" that process input with specific statistical structure and route it to regions that form expanded representations. In the olfactory bulb, for instance, axons of olfactory receptor neurons of a given type converge on a single glomerulus from which signals are routed to an expanded representation in piriform cortex. In the cortico-cerebellar system, outputs of layer-5 cells across cortex converge to the pontine nuclei in the brainstem. In turn, the pontine output innervates an expansion layer of cerebellar granule cells. We investigate the functional role of pre-expansion bottlenecks, and hypothesize that such structures have evolved to compress and decorrelate their input, thereby maximizing the efficiency of the subsequent expansion by removing redundancies and noise. Our theory predicts that connectivity onto the compression layer should reflect the statistical structure of the input, and we show analytically the existence of an optimal compress-

sion level that depends on the level of input noise. Moreover, using simulations we show that the presence of a compressive layer improves the performance of the cerebellum in learning a forward model. Interestingly, such improvement can be obtained using biologically plausible learning rules at cortico-pontine synapses. Our results show how markedly different instances of the same motif emerge from the same underlying principle, and support the idea that compressive layers, including the pontine nuclei, are more than mere relay stations.

III-25. Expressivity of expand-and-sparsify representations

Sanjoy Dasgupta

DASGUPTA@ENG.UCSD.EDU

University of California, San Diego

A simple sparse coding mechanism appears in the sensory systems of several organisms: a d -dimensional input x is mapped to much higher dimension $m \gg d$ by a random linear transformation, and is then sparsified by a winner-take-all process in which only the positions of the top k values are retained, yielding a k -sparse m -dimensional vector z . We study the benefits of this representation for subsequent learning. We first show a universal approximation property, that arbitrary continuous functions of x are well approximated by linear functions of z , provided m is large enough. This can be interpreted as saying that z unpacks the information in x and makes it more readily accessible. We give bounds on how large m needs to be for this approximation to hold, as a function of the input dimension d and the smoothness of the target function. We show, furthermore, that with a slight variant of the thresholding process, the representation is adaptive to manifold structure in the input space.

III-26. Functional organization of posterior cortex during flexible navigation decisions

Shih-Yi Tseng¹

SHIHYI_TSENG@G.HARVARD.EDU

Selmaan Chettih²

SC4551@COLUMBIA.EDU

Charlotte Arlt¹

CHARLOTTE_ARLT@HMS.HARVARD.EDU

Roberto Barroso-Luque³

BARROSOLUQUER@UCHICAGO.EDU

Christopher Harvey¹

CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

¹Harvard University

²Columbia University

³University of Chicago

flexible decision-making during spatial navigation underlies core behaviors such as foraging and predator avoidance, and is crucial to survival in dynamic, real-world environments. These behaviors involve multiple cognitive processes, including sensory processing, navigation planning, motor execution, and adjusting the rules by which these processes are integrated according to experience. Here we investigated the functional organization of the encoding of these processes across posterior cortex with single neuron resolution. We developed a flexible decision-making task for mice navigating a virtual Y-maze, where the rules determining associations between visual cues and reward locations changed unexpectedly within single sessions. Mice adapted to rule switches over tens of trials, exhibiting variability in the across-trial adaptation timescale and their within-trial navigation trajectories. Behavioral modeling of across-trial strategy identified periods of random guessing, biased choices, and high rule-following. Interestingly, this strategy was embodied in the dynamics of choice formation within single trials, as estimated from navigation trajectories. Photoinhibition of retrosplenial cortex or posterior parietal cortex decreased rule-following and delayed choice formation, leading to impaired task performance. We collected calcium imaging data during behavior from hundreds of thousands of neurons tiling posterior cortex. Using single-neuron generalized linear models and population decoding, we quantified neural representations of multiple task and movement-related variables. Neural activity was better described by task variables, including visual cues and the navigational trajectory related to the mouse's choice, than by movement, except in the anterior-lateral portion of

parietal cortex. Representations were distributed but their strength exhibited reliable, quantitative gradients across cortical space. Analysis of these gradients identified three functional modes preferentially encoding visual cues, navigational plans, and ongoing movements. Each mode was distributed but enriched in visual, retrosplenial, and parietal areas, respectively. Our results suggest distinct functions for these areas within a network that broadly distributes information useful for navigation-based decisions across posterior cortex.

III-27. Perceptual straightening of natural videos arises from a cascaded computation

Robbe Goris¹
Yoon Bai¹
Olivier Henaff²
Corey Ziemba¹

ROBBE.GORIS@UTEXAS.EDU
YOONBAI@UTEXAS.EDU
OLIVIER.HENAFF@GMAIL.COM
ZIEMBA@UTEXAS.EDU

¹University of Texas at Austin
²DeepMind, London

A fundamental goal of sensory processing is predicting future states of the environment. Visual prediction is difficult because the stream of images on the retina evolves according to irregular, curved trajectories. We posit that the visual system seeks to transform natural inputs such that neural representations follow straighter, more predictable trajectories. Previous work has revealed that the human visual system and macaque V1 populations indeed selectively straighten natural videos. Which visual computations underlie temporal straightening? Naturally occurring temporal transformations are diverse and complex, making it unlikely that a single transformation can provide a general solution. We therefore hypothesize that perceptual straightening is achieved through a series of hierarchical transformations. Here, we present psychophysical, physiological, and computational evidence for this hypothesis. We reasoned that if straightening arises from a hierarchical computation, then it should be diminished by visual crowding – a mid-level process that limits visual functions that rely on hierarchical computations such as object recognition. Crowding is strong in the visual periphery, but weak near the fovea. We performed a psychophysical experiment and found that perceptual trajectories for cluttered natural scenes are straightened less in the periphery than in the fovea. This suggests that straightening partly arises from mid-level processing. We then performed a physiological experiment and compared the straightness of neural population trajectories elicited by natural image sequences in V1 (early vision) and V2 (mid-level vision) of awake, fixating macaques. We confirmed that V1 selectively straightens natural videos. This effect increases significantly in V2. Which cortical computation underlies this property? Hierarchical convolutional networks optimized for object recognition do not straighten natural videos. But a biologically inspired model in which V2 units nonlinearly combine V1 inputs does account for our findings. Temporal straightening may thus be an objective that shapes the hierarchical transformations of the primate visual system.

III-28. Strong disinhibition establishes long neuronal timescales critical for working memory

Robert Kim¹
Terrence Sejnowski²

R8KIM@UCSD.EDU
TERRY@SALK.EDU

¹University of California, San Diego
²Salk Institute for Biological Studies

Cortical areas important for working memory (WM), such as prefrontal cortex, utilize neurons with long timescales and stable temporal receptive fields to support reliable representations of stimuli. Despite recent advances in experimental techniques, the mechanisms for the emergence of timescales long enough to support WM are unclear and challenging to investigate experimentally. In order to study the circuit mechanisms necessary for stable

temporal receptive fields, we employed spiking recurrent neural networks (RNNs) trained via a gradient-descent method to perform a WM task. By comparing the results from the network model with the results obtained from rhesus monkeys trained on similar WM tasks, we first establish that the spiking RNN model is a valid model for probing neuronal timescales. We found that both network model and macaque prefrontal cortex display heterogeneous timescales and utilize units with long timescales to maintain short-term memory. To our surprise, analyzing the emerging connectivity patterns of the RNNs revealed that dense inhibitory-to-excitatory connections combined with strong inhibitory-to-inhibitory connections were critical for generating units with long timescales, which in turn could be recruited for WM maintenance. Our findings are closely aligned with previous experimental findings underscoring the importance of disinhibition exerted by vasoactive intestinal peptide (VIP) interneurons and suggest that higher cortical areas could utilize a disinhibitory circuit motif embedded in a dense inhibitory network to support WM.

III-29. Network mechanisms of long-term spatial memory consolidation

Andres Grosmark¹

Fraser Sparks¹

Matt Davis²

Attila Losonczy¹

AG3633@COLUMBIA.EDU

FS2594@COLUMBIA.EDU

MATTJDAVIS@GMAIL.COM

AL2856@COLUMBIA.EDU

¹Columbia University

²University of Texas, Austin

The mechanisms by which sharp-wave/ripple (SWR)-related activity leads to stable or unstable spatial-coding outcomes remain little understood. Moreover, the absence of well-defined Ca²⁺-analogues of established biomarkers of offline activity such the LFP-detected SWRs, have restricted the accessibility of offline activity to Ca²⁺-imaging. We performed combined chronic CA1 LFP recordings with simultaneous fast (60 Hz) two-photon Ca²⁺ imaging to stably track the activity of genetically identified mouse CA1 principal neurons over weeks of online spatial behavior and offline resting. Our results demonstrate: 1) that joint LFP recording and Ca²⁺-imaging is an effective tool for studying fast neural dynamics, such as SWR-related responses, 2) stable and unstable cross-day spatial coding of CA1 place cells, 3) stimulus-specific and long-lasting (~24 h) reactivation of RUN-related ensembles, 4) the first reported Ca²⁺-imaging of SWR-associated place cell sequence replay of a recent non-local spatial experience, 5) that SWR-recruitment and contribution to RUN-related reactivation, predict long-term spatial stability outcomes, and 6) that PRE to POST changes in excitability and replay, as well as their stabilizing effect on long-term memory traces, are most prominent for locations farther from the reward. These results for the first time directly link CA1 place cell SWR-recruitment dynamics to their long-term (multi-day) memory consolidation. Furthermore, the results suggest that, consistent with CA3's role in both SWR generation and pattern completion and contrasting RUN-epoch offline activity, POST-learning SWR's preferentially consolidate the memory representations of the trajectories between rewarded locations. Therefore, POST-epoch SWR's may promote the stabilization of long-term representations whose statistics more closely reflect the relative uniformity of the underlying spatial stimulus as compared to the strong reward bias observed in the animals' spatial behavior.

III-30. Structural and functional differences in the head direction circuit of two insect species

Ioannis Pisokas¹

Stanley Heinze²

Barbara Webb³

I.PISOKAS@SMS.ED.AC.UK

STANLEY.HEINZE@BIOL.LU.SE

BWEBB@INF.ED.AC.UK

¹University of Edinburgh

²Lund University

³The University of Edinburgh

Ants, bees and other insects have the ability to return to their nest or hive using a navigation strategy known as path integration. Similarly, fruit flies employ path integration to return to a previously visited food source while moths use it to migrate. An important component of path integration is the ability of the animal to keep track of its heading relative to salient visual cues. A highly conserved brain region known as the central complex has been identified as being of key importance for the computations required for an insect to keep track of its heading. However, the details of the underlying heading tracking neural circuit as well as its operation are not well understood. We sought to address this shortcoming by deriving the effective underlying neural circuits of two species, the fruit fly and the locust. Our analysis revealed that regardless of the anatomical differences between the two species the essential circuit structure has not changed. Both effective neural circuits have the structural topology of a ring attractor with an eight-fold symmetry. However, despite the strong similarities between the ring attractors in the fruit fly and the locust, there remain differences. We found that two apparently small anatomical differences have significant functional effect on the ability of the two circuits to track fast rotational movements and to maintain a stable heading signal. Our results highlight that even seemingly small differences in the distribution of dendritic fibres can have a significant effect on the dynamics of the effective ring attractor circuit with consequences for the behavioural capabilities of the animal. These differences, emerging from morphologically distinct single neurons, call into question the validity of broad generalisations drawn from studies in any one species, such as *Drosophila*, and highlight the importance of a comparative approach to neuroscience.

III-31. Neural and computational basis for organizing episodic experience into event units

Chen Sun
Kelsey Allen
Susumu Tonegawa

Massachusetts Institute of Technology

CSUN@MIT.EDU
KRALLEN@MIT.EDU
TONEGAWA@MIT.EDU

What is an episode? An established view is that episodic experience is tracked by the brain as a sequence of moment to moment continuous variations but recently a neural correlate has also been demonstrated at the single cell level (Sun et al, 2019) which organizes the same episode as a sequence of discrete and unitary events (the “event code”). In this present study, we investigate the neural and computational basis for chunking experience into event units. We first suggest that the statistical structure of natural experiences may not be enough to permit event chunking and that an inductive bias may be necessary. We then show, using simultaneous calcium imaging and optogenetic inhibition, that hippocampal CA2 is mechanistically involved in event chunking. finally, we show the utility of the event code by demonstrating the event code is preserved even between two experiences that differ sensorily and spatially, suggesting its involvement in abstract knowledge transfer between experiences.

III-32. Geometric representation of abstract learned knowledge by neural manifolds in hippocampus

Edward H Nieh
 Manuel Schottdorf
 Nicolas W Freeman
 Sue Ann Koay
 Lucas Pinto
 Jeffrey L Gauthier
 Carlos D Brody
 David Tank
 Princeton University

ENIEH36@GMAIL.COM
 MSCHOTTDORF@PRINCETON.EDU
 NICOLASWFREEMAN@GMAIL.COM
 KOAY@PRINCETON.EDU
 LPINTO@PRINCETON.EDU
 JLGAUTHI@PRINCETON.EDU
 BRODY@PRINCETON.EDU
 DWTANK@PRINCETON.EDU

Hippocampal neurons encode physical and sensory variables such as space or auditory frequency in cognitive maps, but it is unknown how representations of abstract learned knowledge are instantiated by the coordinated activity of single neurons. Here, we recorded hippocampal activity in rodents performing an abstract evidence accumulation task in virtual reality. We found that individual hippocampal neurons jointly encoded evidence, a non-physical variable that the animal had to infer from the task structure, together with the animal's position in the maze. Nonlinear dimensionality reduction demonstrated that the activity of ~500 simultaneously recorded neurons could be accurately described by remarkably few dimensions, as few as ~3-4. Within this low-dimensional space, both physical and abstract variables were mapped in an orderly fashion, creating a geometric representation of multiple task variables as a neural manifold. The existence of conjoined cognitive maps suggests that the hippocampus performs a general computation – to create geometric representations of learned knowledge instantiated by task-specific low-dimensional manifolds.

III-33. Navigational uncertainty provides a unifying account of human navigational behavior and rodent grid field deformations

Yul Kang¹
 Daniel Wolpert²
 Mate Lengyel¹

YUL.KANG.ON@GMAIL.COM
 WOLPERT@COLUMBIA.EDU
 M.LENGYEL@ENG.CAM.AC.UK

¹University of Cambridge

²Columbia University and University of Cambridge

Spatial navigation relies on vision to sense remote landmarks. Although the role of uncertainty in visual perception has long been appreciated, its impact on navigational performance has not been studied systematically. Here, we develop an ideal observer model adapted from state-of-the-art probabilistic algorithms for Simultaneous Localization And Mapping (SLAM). Our model optimally combines visual and self-motion cues, and maintains a joint probabilistic representation of self location as well as the location of “landmarks” (e.g. the corners of an arena), of which only a subset may be visible at any time. We show that the fundamental constraints of the 3D-to-2D projection in vision lead to higher uncertainty about the location of a landmark along the line that connects it to the observer. In turn, our model accounts for a wide range of experimental observations. 1. When the task is to detect changes in the layout of objects when a scene is viewed from two different angles, subjects' performance varies non-monotonically with the angular difference. 2. Human navigation is less accurate in a trapezoid- compared to square-shaped environment, and in the narrow compared to the broad part of a trapezoid. 3. Assuming that grid cells implement a probabilistic population code for the ideal observer's posterior, the same model, without further parameter tuning, also explains how grid fields adapt to environmental geometry. Specifically, our model accounts for the deformation of grid fields recorded in rats navigating in trapezoid environments such that their hexagonal structure becomes more elliptical along the long axis of the trapezoid, and more so in the narrow than in the broad part of the trapezoid. These results suggest that visual landmark-derived uncertainty about one's location is a critical factor in navigation, and it is a major determinant of both behavior and neural representations.

III-34. Auditory activity is diverse and widespread throughout the central brain of drosophila

Diego Pacheco Pinedo¹
Eftychios Pnevmatikakis²
Stephan Thiberge¹
Mala Murthy¹

DPACHECO@PRINCETON.EDU
EPNEVMATIKAKIS@FLATIRONINSTITUTE.ORG
THIBERGE@PRINCETON.EDU
MMURTHY@PRINCETON.EDU

¹Princeton University

²flatiron Institute

Sensory pathways are typically studied by starting at the periphery – the receptor neurons - and then characterizing responses along postsynaptic partners in the brain. However, this leads to a bias in analysis of earliest stages of sensory processing in most model organisms. Here, we present new methods for unbiased volumetric neural imaging with precise across-brain registration, that allow us to characterize auditory activity throughout the entire central brain of *Drosophila melanogaster*, and to make comparisons across trials, individuals, and sexes (figure 1A). We image activity volumetrically using GCaMP6s (Chen et al., 2013) expressed pan-neuronally and fully automate segmentation of regions of interest (Pnevmatikakis et al., 2016). Using this approach, we discover that auditory activity is widespread across nearly all central brain regions (33 out of 36 major brain regions) (figure 1A and 1B). These regions include the mechanosensory pathway, and novel regions known to process other sensory modalities (including all levels of the olfactory pathway), or to drive motor behaviors (such as the central complex) (figure 2A and 2B). The auditory representations that we uncover are diverse in their temporal profiles, becoming more selective for particular aspects of the acoustic stimuli at higher stages of the pathway (figure 2B). We find that auditory responses are stereotyped across trials and animals in early mechanosensory regions, becoming more variable at higher layers of the putative pathway (figure 3A and 3B). Finally, we exhaustively map tuning to spectral and temporal features of acoustic signals within brain regions that carry the majority of auditory activity and find that responses are focused on aspects of the conspecific acoustic signals (figure 4A and 4B). This study leverages brain-wide imaging and computational techniques to provide the first brain-wide description of sensory processing and feature tuning in *Drosophila*.

III-35. Representation of 3D orienting movements in primary visual cortex

Grigori Guitchounts¹
Javier Masis¹
Steffen Wolff¹
David Cox²

G.GUITCHOUNTS@GMAIL.COM
JMESIS@G.HARVARD.EDU
STEFFENWOLFF@FAS.HARVARD.EDU
DAVID.D.COX@IBM.COM

¹Harvard University

²IBM

Movement-related neuronal activity has been found across many brain regions. In primary visual cortex (V1), locomotion signals in head-fixed animals are thought to reflect gain modulation, running speed, navigational variables, or prediction of visual flow direction. An important but unaddressed question is how the visual system distinguishes between self-generated sensory signals and those that come from the outside world. This is unknown, in part, because most previous studies have not examined the effects of unrestrained motion in 3D on sensory cortical dynamics. We found that neuronal dynamics in rat V1 during free movement are far richer than previously known, and relate in precise ways to particular movements. V1 neurons were suppressed before and during 3D orienting movements of the head on a sub-second timescale in the dark, excited in the light, and reflected the direction of movement even in complete darkness. Individual V1 neurons were tuned to turn direction in both lighting conditions, but formed largely non-overlapping populations: one that was tuned in the dark, and another tuned in the light. Responses to visual stimuli were suppressed during orienting movements compared to rest. Finally, lesions to secondary motor cortex (M2), a region that sends direct projections to V1 and has been implicated in controlling orienting movements⁸, greatly diminished head-turn-related responses and directional

information in V1. Our data show that the rodent visual system represents 3D orienting movements and suggest a mechanism for distinguishing between self-generated sensory signals and those that come from the outside world.

III-36. Understanding the emergence of hexagonal grid cells in networks trained to path integrate

Gabriel Mel
Surya Ganguli
Ben Sorscher
Samuel A Ocko
Stanford University

MELDEFON@GMAIL.COM
SGANGULI@STANFORD.EDU
BSORSCH@GMAIL.COM
SOCKO@STANFORD.EDU

The discovery of entorhinal grid cells engendered considerable interest in understanding how hexagonal firing patterns arise from neural circuits, and what functional role they might play in path-integration. There are two seemingly unrelated frameworks: Mechanistic models account for hexagonal firing through pattern-forming dynamics in a recurrent neural network (RNN) with hand-tuned center-surround connectivity. Normative models observe that grid-like firing fields emerge in RNNs trained to solve navigational tasks. However, across different normative models, such firing fields exhibit diverse lattice structures, including square, hexagonal, or highly heterogeneous. Here we provide an analytic theory that unifies the two frameworks by (1) explaining why grid patterns arise in both mechanistic and normative models through a single common mechanism, and (2) explaining when specific lattice structures arise. This theory enables us to obtain neural networks that robustly yield hexagonal (but not square) firing fields simply by training them to path-integrate with non-negative firing rates.

These normatively trained networks then yield a fundamental mechanistic challenge for theoretical neuroscience: despite full access to their connectivity and dynamics, it is difficult to understand precisely how they path-integrate. We develop algorithmic methods to disentangle the connectivity of these trained networks, finding an approximate, hidden translation invariant center-surround connectivity with velocity-driven shifter circuits that enable both grid cells and path-integration. Thus, remarkably, these trained RNNs can be understood as a generalization of prior hand-tuned grid cell models.

In summary, we (1) provide a unified theory for why grid cells and which lattice structures emerge across diverse normative and mechanistic models; (2) reveal that hexagonal grids arise robustly as a consequence of solving path-integration tasks with biologically plausible firing rate constraints; (3) develop methods for finding hidden connectivity patterns in RNNs that elucidate neural mechanisms for path-integration in trained networks; (4) show these mechanisms generalize those of prior hand-tuned networks.

III-37. Comparing physical inferences in humans, monkeys and task-optimized recurrent neural network models

Rishi Rajalingham
Aida Piccato
Mehrdad Jazayeri
Massachusetts Institute of Technology

RISHI.RAJALINGHAM@GMAIL.COM
APICCATO@MIT.EDU
MJAZ@MIT.EDU

From a glance, humans can richly parse a physical scene to infer the latent states of objects within the scene, predict plausible future states, and plan intervening actions. Cognitive science theories posit that such “intuitive physics” abilities are supported by internal simulations of the external world. Here, we test this theory by directly comparing the behavior of humans and monkeys to recurrent neural network (RNN) models constructed to mimic simulation, in a complex behavioral task. The objective of the task is to intercept a ball moving across a screen

at different velocities in the presence of reflectors and occluders. Critically, this task forces subjects to infer the dynamics of an unobserved external process. We first collected large datasets in humans and monkeys, and characterized their behavioral error patterns across task conditions. We reasoned that if flexible primate behavior in this task relies on simulation, then RNNs bestowed with the capacity to simulate should more accurately capture those behavioral patterns. To test this, we trained different RNNs to solve this task, either with or without simulation as an additional optimization goal, and compared the behavioral patterns of humans, monkeys and RNNs. We found that humans and monkeys exhibit remarkably similar behavioral patterns, suggesting common underlying neural circuits in primates. This primate behavior was not captured by all RNNs. RNNs optimized to mimic simulation rapidly learn to perform this task, and their output quantitatively resembles primate behavior. In contrast, RNNs trained without this inductive bias require much more training data to solve this task, and their behavior diverge from those of primates. Taken together, these results are consistent with the hypothesis that primates use simulation strategies to solve this task, and suggest that imposing suitable inductive biases in RNNs may help uncover how the primate brain implements internal models.

III-38. Modeling large-scale brain network dynamics in response to electrical stimulation

Yuxiao Yang¹
Shaoyu Qiao²
Omid Sani¹
Isaac Sedillo²
Breonna Ferrentino²
Bijan Pesaran²
Maryam Shanechi¹

YYX.YUKI@GMAIL.COM
SQIAO@NYU.EDU
OMID.GHASEMSANI@USC.EDU
JIS332@NYU.EDU
BREONNA.FERRENTINO@NYU.EDU
BIJAN@NYU.EDU
SHANECHI@USC.EDU

¹University of Southern California

²New York University

Modeling the dynamic response of large-scale brain networks to ongoing electrical stimulation has remained elusive to date. Achieving such modeling is critical for probing the functional organization of brain networks and for developing neuromodulation technology. Here, we develop novel dynamic input-output (IO) models that successfully predict the real-time response of brain networks to continuously-changing electrical stimulation in non-human primates (NHPs). Using a customized semi-chronic microdrive system, we deliver continuous time-varying electrical stimulation and simultaneously record local field potential (LFP) activity from multi-region brain networks. To obtain informative data for fitting the IO models, we apply a new stimulation pulse train to fully excite the brain network. The input in our IO models is the pulse train frequency and amplitude, which change in real time. The output in our IO models is the LFP power features. We evaluate the IO models in cross-validation, where we use the trained IO models to feed-forward predict the time-evolution (dynamics) of LFP power features in response to stimulation in the test set. We compare the predicted time-evolution to the observed time-evolution. We find that our IO models accurately predict the dynamic brain network response in cross-validation. Further, we study the dynamical characteristics of the brain network response and find that the response exhibits damping and oscillatory dynamics, showing that the dynamic structure of the IO models is essential for the prediction of the response. Finally, the relative strength of the dynamic stimulation effect at different brain regions can be predicted from at-rest functional connectivity within the network, which is computed from LFP recordings without stimulation. These models can lead to precise modulation of brain functions and dysfunctions.

III-39. Adaptive stimulus selection improves deep neural network models of macaque V4

Benjamin Cowley¹
 Patricia Stan²
 Matthew Smith³
 Jonathan Pillow¹

BCOWLEY@PRINCETON.EDU
 PLS23@PITT.EDU
 MATTSMITH@CMU.EDU
 PILLOW@PRINCETON.EDU

¹Princeton University

²University of Pittsburgh

³Carnegie Mellon University

A major goal in visual neuroscience is to understand how the brain transforms visual input into a neural code. A key test of our understanding is to accurately predict the responses of visual cortical neurons to novel stimuli, such as natural images. To date, the most accurate models rely on a linear mapping to predict responses from features of an intermediate layer in a pre-trained convolutional deep neural network (DNN). Ideally, one would train the entire DNN to directly map image to response, but this DNN would overfit to the limited amount of data currently obtainable from neuroscientific experiments. Indeed, we find that even in the case of fitting a DNN to map pre-trained features to responses, a linear mapping performs better. In this work, we propose a framework to train a DNN mapping between features and responses that incorporates two machine learning approaches: ensemble learning and adaptive stimulus selection (also known as active learning). We find that ensemble learning, which predicts responses by averaging the predictions of many networks, has better prediction than that of a linear mapping, even for the same amount of data. Next, to incorporate adaptive stimulus selection, we select the next session's training images such that the selected images have the largest disagreement among networks. In closed-loop experiments to predict responses of macaque V4 neurons to colorful natural images, we find models trained with adaptive stimulus selection perform better than randomly-selecting training images. Finally, we analyze our resulting model, which yields state-of-the-art predictions, and find that only a small number of the networks' embedding variables (~30) are needed to predict a large fraction of the responses' variance. Overall, our work provides a path forward to obtain highly predictive models of visual cortical neurons that would otherwise be untrainable with such limited recording time.

III-40. Why do neurons spike spontaneously?

Chaitanya Chintaluri
 Tim Vogels
 University of Oxford

CHAITANYA.CHINTALURI@DPAG.OX.AC.UK
 TIM.VOGELS@CNCB.OX.AC.UK

Spontaneous firing, observed in many neurons, is often attributed to ion channel or network level noise. Cortical cells during slow wave sleep exhibit transitions between so called Up and Down states. In this sleep state, with limited sensory stimuli, neurons fire in the Up state. Spontaneous firing is also observed in slices of cholinergic interneurons, cerebellar Purkinje cells and even brainstem inspiratory neurons. In such in vitro preparations, where the functional relevance is long lost, neurons continue to display a rich repertoire of firing properties. It is perplexing that these neurons, instead of saving their metabolic energy during information downtime and functional irrelevance, are eager to fire. We propose that spontaneous firing is not a chance event but instead, a vital activity for the well-being of a neuron. Neurons in anticipation of synaptic inputs, keep their ATP levels at maximum. As recovery from inputs requires most of the energy resources, neurons are ATP surplus and ADP scarce during synaptic quiescence. With ADP as the rate-limiting step, ATP production stalls in the mitochondria. This leads to toxic Reactive Oxygen Species (ROS) formation, which are known to disrupt many cellular processes. We hypothesize that spontaneous firing occurs at these conditions as a release valve to spend energy and to restore ATP production, shielding against ROS. By linking a mitochondrial metabolism model to a conductance-based neuron model, we show that spontaneous firing depends on baseline ATP usage and on ATP-cost-per-spike. From our model, emerges a mitochondrial mediated homeostatic mechanism that provides a recipe for different

firing patterns. Our findings, though mostly affecting intracellular dynamics, may have large knock-on effects on the nature of neural coding. Hitherto it has been thought that the neural code is optimised for energy minimisation, but this may be true only when neurons do not experience synaptic quiescence.

III-41. A dynamical model with E/I balance explains robustness to optogenetic stimulation in motor cortex

Lea Duncker¹
Daniel O'Shea²
Krishna Shenoy²
Maneesh Sahani¹

DUNCKER@GATSBY.UCL.AC.UK
DJOSHEA@STANFORD.EDU
SHENOY@STANFORD.EDU
MANEESH@GATSBY.UCL.AC.UK

¹Gatsby Computational Neuroscience Unit, University College London

²Stanford University

Targeted optogenetic perturbations are key to investigating functional roles of sub-populations within neural circuits, yet their effects in recurrent networks may be difficult to interpret. Previous work has shown that optogenetic stimulation of excitatory cells in macaque motor cortex creates large perturbations of task-related activity, but has only subtle effects on ongoing or upcoming behavior, or the future dynamical evolution of neural population activity [O'Shea et al., Cosyne14]. fitting dynamical models to data, Duncker et al. [Cosyne17] have shown that this robustness is consistent with a local nonnormal dynamical system whose nullspace is well-aligned with the optogenetic perturbation pattern. Here, we ask how such alignment might arise. We hypothesize that circuit-level features such as E/I balance might contribute crucially. To evaluate this hypothesis from neural recordings, we develop a fitting approach to identify a high-dimensional discrete-time balanced E/I network that expresses the low-dimensional and smooth dynamics observed in the recorded population responses. We fit our model to neural data recorded on nonstimulated trials of a center-out reaching task, and study the E/I network responses to perturbations of the excitatory sub-population. The simulated E/I network is robust to these perturbations and its response patterns closely match those of recorded neural data on stimulation trials. This suggests that E/I balance in cortical circuits may offer an explanation for the robustness to optogenetic perturbations observed in motor cortex. Ultimately, developing more explicit links between circuit-level properties and population-level dynamics will lead to a better understanding of the effects of perturbations on cortical dynamics.

III-42. System identification in the brain: clustering dynamical rules from sensory data

Tiberiu Tesileanu¹
Samaneh Nasiri²
Anirvan Sengupta³
Dmitri Chklovskii¹

TTESILEANU@FLATIRONINSTITUTE.ORG
SAMANEH.NASIRI.GH@GMAIL.COM
ANIRVANS.PHYSICS@GMAIL.COM
MITYA@FLATIRONINSTITUTE.ORG

¹flatiron Institute

²Emory University

³Rutgers University

Sensory information reaches the brain in a stream of data exhibiting non-trivial temporal correlations. These correlations arise from multiple effects, from dynamical laws acting in the environment to non-trivial receptor response profiles. Detecting changes in the correlation structure of sensory signals can thus shed light on important changes in the dynamics of the environment, as well as on internal changes.

Here we present a biologically plausible neural network for clustering time-series data based on their dynamical structure. Our method starts with a circuit that keeps a running estimate of the autocorrelation. The activations

of the output neurons in this system can be thought of as a vector of autocorrelations at several different lags. This autocorrelation vector is then used as input to a second neural network that acts as a clustering algorithm. The largest output of this network indicates the most likely identity of the generating process for the input series. The clustering neural network is derived from an optimization principle that aims to map signals with similar autocorrelation structure to similar outputs.

We test the performance of our network on signals that are drawn from autoregressive moving-average (ARMA) models with piecewise-continuous parameters. We find that the clusters identified using our method are in very good agreement with the ground truth. Comparing to a cepstral method used in control theory, we find that our method is just as effective, despite being simpler, faster, and using less memory.

We note that, unlike most work on change-point detection, our approach is sensitive to transitions in the temporal correlation structure of the signal, not only to changes in the instantaneous distribution of values.

III-43. Synaptic plasticity in balanced networks

Kresimir Josic¹

Alan Akil¹

Robert Rosenbaum²

KRESIMIR.JOSIC@GMAIL.COM

AAKIL@MATH.UH.EDU

RROSENB1@ND.EDU

¹University of Houston

²University of Notre Dame

The dynamics of local cortical networks is irregular, but correlated. While dynamic balance is a plausible mechanism that generates such stochastic activity both in asynchronous, and correlated regimes, it remains unclear how such balance is achieved and maintained in actual neural networks. In particular, it is not fully understood how plasticity induced changes in the network affect balance, and in turn, how correlated, balanced activity impacts learning.

How will the dynamics in balanced network change under different plasticity rules in the asynchronous or correlated state? How does the presence of correlations in the recurrent network impact learning? When and how do correlations change the evolution of weights, their eventual magnitude, and structure across the network?

To address these questions, we develop a general theory of learning in balanced networks. We show that, in general, balance is attained and maintained under plasticity induced weight changes both in the asynchronous and correlated state. We find that correlations in the input mildly, but significantly affect the evolution of synaptic weights. More importantly, under certain plasticity rules, we find an emergence of correlations between weights and rates. Under these rules, synaptic weights converge to a stable manifold in weight space with their final configuration dependent on the initial state of the network. Thus the network can retain a memory of its initial state even under learning.

III-44. Neural population dynamics in motor cortex during grasp

Aneesha K Suresh¹

James Goodman²

Elizaveta V Okorokova²

Matthew Kaufman²

Nicholas G Hatsopoulos²

Sliman Bensmaia²

ANEESHA.KRITHIKA@GMAIL.COM

J.MIKE.GOODMAN@GMAIL.COM

LIZOK@UCHICAGO.EDU

ANTIMATT@GMAIL.COM

NICHO@UCHICAGO.EDU

SLIMAN@UCHICAGO.EDU

¹The University of Chicago

²University of Chicago

Low dimensional linear dynamics are observed in neuronal population activity in primary motor cortex (M1) when

monkeys make reaching movements. This population-level behavior is consistent with a role for M1 as an autonomous pattern generator that drives muscles to give rise to movement. In the present study, we examine whether low-dimensional linear dynamics are also observed during grasping movements, which involve fundamentally different patterns of kinematics and muscle activations. Using a variety of analytical approaches, we show that M1 does not exhibit such dynamics during grasping movements. In fact, the grasp-related neuronal dynamics in M1 are similar to their counterparts in somatosensory cortex, which is driven primarily by inputs rather than by intrinsic dynamics. The basic structure of the neuronal activity underlying hand control is thus fundamentally different from that underlying arm control.

III-45. Temporal context invariance reveals a timescale hierarchy in human auditory cortex

Sam Norman-Haignere¹

Laura Long¹

Orrin Devinsky²

Werner Doyle²

Guy McKhann¹

Catherine Schevon¹

Adeen flinker²

Nima Mesgarani¹

¹Columbia University

²New York University

SN2776@COLUMBIA.EDU

LKL2124@COLUMBIA.EDU

ORRIN.DEVINSKY@NYULANGONE.ORG

WERNER.DOYLE@NYULANGONE.ORG

GM317@CUMC.COLUMBIA.EDU

CAS2044@CUMC.COLUMBIA.EDU

ADEEN.FLINKER@NYULANGONE.ORG

NIMA@EE.COLUMBIA.EDU

Natural stimuli, like speech and music, are structured at many different timescales from milliseconds (e.g. pitch) to seconds (e.g. words) and minutes (e.g. narrative meaning). How does the brain integrate information across such diverse timescales? Answering this question has been challenging in part because there is no general method for estimating sensory timescales in the brain. Here, we introduce a simple paradigm (the “temporal context invariance” or TCI paradigm) for inferring the temporal integration period of any sensory response, defined as the window in time during which stimuli alter the response. We present segments of natural stimuli in a pseudorandom order such that the same segment occurs in two different contexts (different surrounding segments), and we test how long segments need to be for the response to be unaffected by the surrounding context. We apply the TCI paradigm to map integration periods in human auditory cortex using intracranial recordings from epilepsy patients. We find that integration periods exhibit clear anatomical organization with substantially longer integration periods in non-primary (~400 ms) vs. primary (~100 ms) regions, providing support for hierarchical models. We also show that selectivity for sound categories (e.g. speech or music) first emerges at timescales of ~200 ms, suggesting selectivity for syllables or notes. Our method should be applicable to any sensory response, regardless of the recording method, stimulus or modality. We are currently taking advantage of this flexibility to test whether integration periods dynamically adapt to the timescale of the sound features being coded by a neural population.

III-46. Deep neural networks reveal adaptedness of human sound localization to real-world environments

Andrew Francl

Josh McDermott

Massachusetts Institute of Technology

FRANCL@MIT.EDU

JHM@MIT.EDU

Mammals localize sounds using information from their two ears. Decades of research has characterized auditory cues to spatial location. However, these cues are unreliable in real-world conditions, in which reflections provide erroneous information, and noises mask parts of target sounds. To better understand real-world localiza-

tion we equipped a deep neural network with simulated human ears and trained it to localize sounds in a virtual environment with simulated surface reflections and background noise. The resulting model accurately localized real-world recordings made in an actual classroom with noise and reverberation, outperforming standard microphone array algorithms popular in engineering. In simulated experiments the network exhibited many previously documented features of human spatial hearing: increased spatial acuity along the midline, frequency-specific sensitivity to interaural time and level differences, integration across frequency, and biases for sound onsets (the "precedence effect"). To explore the relationship between the learned strategy and the natural environment, we varied the training conditions. Specifically, hearing researchers have long hypothesized that the precedence effect exemplifies a localization strategy adapted to reverberant environments, in which erroneous location information from reflections is suppressed. We directly tested this idea by retraining our model in an anechoic virtual environment. The anechoically-trained model retained some signatures of human spatial hearing, but no longer showed a bias towards sound onsets. The results show how human hearing is adapted to the challenges of real-world environments and illustrate the use of artificial neural networks in place of traditional ideal observer models.

III-47. Working memory for temporal estimates via slow neural manifolds

Alexandra Ferguson¹

Manuel Beiran²

T Vincenza Parks³

Srdjan Ostojic²

Mehrdad Jazayeri¹

AFERGU@MIT.EDU

MANUEL.BEIRAN@ENS.FR

TVP6@PITT.EDU

SRDJAN.OSTOJIC@ENS.FR

MJAZ@MIT.EDU

¹Massachusetts Institute of Technology

²Ecole Normale Supérieure

³University of Pittsburgh

Decades of work have probed working memory (WM) as a substrate for holding information relevant for immediate future behavior. Neural circuits can readily hold discrete variables in WM through establishing stable fixed points. For continuous variables, however, the nature of the underlying representations is not well understood. Line attractors have been proposed as a candidate substrate. Theoretical considerations have questioned whether neural systems are able to establish robust and stable line attractors. Here, we set out to examine the nature and stability of WM for a continuous variable; i.e., an interval of time. We trained monkeys to measure a time interval drawn from a prior distribution, hold it in memory for a variable delay, and reproduce it afterwards. Consistent with the effect of noise on a line attractor, the animals' performance dropped progressively for longer memory delays. In contrast to predictions of a line attractor, this drop was largely due to an increase in bias (larger regression toward the mean), not variance. This finding suggests that the time interval might be memorized through a mechanism distinct from a perfect line attractor. To further probe candidate neural mechanisms, we trained low-rank recurrent neural networks (RNNs) to perform the same task. Based on previous recordings of neural activity in variants of this task ([1]-[3]), we looked for the minimal dimensionality implementation of this task. Analysis of dynamics revealed that RNNs relied on a hierarchy of slow manifolds to implement the task. In particular, the temporal interval was held in WM using a dynamical landscape close to, but distinct from a perfect line attractor. This deviation from a line attractor caused slow drifts of representation in WM creating biases similar to monkeys' behavior. Overall, guided by experimental results we propose a novel robust dynamical mechanism for WM based on slow manifolds.

III-48. Stimulus representation and inference by recurrently coupled neuronal populations

Wenhao Zhang¹
Si Wu²
Kresimir Josic³
Brent Doiron¹

WENHAO.ZHANG@PITT.EDU
SIWU@PKU.EDU.CN
JOSIC@MATH.UH.EDU
BDOIRON@PITT.EDU

¹University of Pittsburgh

²Peking University

³University of Houston

Conventional neural coding theory concludes that recurrent connections in a neural circuit degrade the available stimulus information because they introduce spurious correlations in population activity. This theory is at odds with the reality that recurrent interactions between neurons are widespread throughout the cortex. A fundamental, and still quite open, challenge in computational neuroscience is to elucidate the role of recurrent connections in cortical information processing. To shed light on this issue we consider a richer external world, one where an overarching context generates an associated stimulus structure. We formulate a novel population coding framework where recurrent circuitry is essential for inferring both a driving stimulus and the world context. In this framework, neural population activity and recurrent inputs combine to represent the joint posterior distribution over the stimulus and world context. We show that recurrent connections can encode the prior of the co-occurrence between a specific stimulus and context, and that stochastic network dynamics effectively implements sampling-based inference to approximate the posterior distribution. In this way our framework expresses the representation and inference of an abstract Bayesian theory through the structure of a biologically plausible, recurrent neural circuit. Our work also provides novel and testable pre/post-dictions; for example, it suggests that neuronal correlations from recurrent dynamics are a necessary consequence of the network reading out the prior, and are thus a reflection of the extra stimulus information contained in the prior. Moreover, the differential (information limiting) correlations which emerge from the recurrent circuit are a consequence of sampling-based inference on the stimulus manifold embedded in the high dimensional population code.

III-49. Continual learning, replay and consolidation in a forgetful memory network model

Nimrod Shaham¹
Jay Chandra¹
Kreiman Gabriel²
Haim Sompolinsky³

NSHAHAM@FAS.HARVARD.EDU
JAYCHANDRA@COLLEGE.HARVARD.EDU
GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU
HAIM@FIZ.HUJI.AC.IL

¹Harvard University

²Boston Children's Hospital

³The Habrew University of Jerusalem and Harvard University

The brain has a remarkable ability to deal with a continuous stream of information, storing new memories and learning to perform new tasks. This is done largely without losing previously learned knowledge, which can be stored throughout animal's lifetime. Understanding how the brain can continuously learn new information while avoiding catastrophic interference between memories is a major challenge in machine learning and neuroscience. Previous palimpsest models of recurrent neural networks assume Hebbian learning of new memories together with synaptic weight decay. This decay allows the networks to continually learn many new patterns of activity, while forgetting old memories. In these models, memories are almost perfectly retrievable up to a certain critical age and their retrievability decrease rapidly at older age, in contrast with behavioral observations. To overcome this deterministic forgetting of all old memories, we introduce replay to the system, in the form of stochastic nonlinear relearning of previously stored patterns. We show that this nonlinear replay mechanism gives rise to a stable consolidation effect - a fast initial increase in memory retrievability, followed by a forgetting process

which is orders of magnitude slower compared to the decay rate. Importantly, in our model, the retrievability of memories, and their consolidation is stochastic. Thus, our system exhibits a large capacity of learning new memories while a subset of memories is consolidated and long lived. Our model shows for the first time how a recurrent neural network can continuously learn and store selected information for lifelong timescales. We show that stochastic nonlinear replay of learned information results in an advantage for memory capacity. Finally, we generate predictions for retrieval probability as a function of a given memory's age, in healthy subjects as well as after stroke events.

III-50. A normative hippocampus model: robustly encoding variables on smooth manifolds using spiking neurons

Paxon Frady^{1,2}
Friedrich Sommer^{1,2}

EPAXON@BERKELEY.EDU
FSOMMER@BERKELEY.EDU

¹University of California, Berkeley
²Intel

Different brain regions extract and manipulate behavioral variables, which are often continuous and organized on smooth manifolds (Stringer et al., 2019), such as location in an environment. A useful encoding scheme for these variables has to fulfill two fundamental requirements: (1) Distinguishability: Points far apart on a manifold, should be encoded or “indexed” by dissimilar (independent or orthogonal) patterns so that the probability of confusion is low. (2) Smoothness: Points close to each other should be encoded by similar patterns such that the code is noise tolerant and points can be represented by interpolation of neighbors. Here we propose a normative encoding model with properties (1)-(2), leveraging timing of individual spikes. Our approach is based on complex-valued connectionist models for symbolic reasoning (Plate, 2003), and recent theoretical results on how complex state spaces can be mapped onto the dynamics of spiking neural networks through a phase-to-timing mapping (Frady and Sommer, 2019). The resulting model “tiles” a manifold such that each location is represented by a unique neural activity vector that can be used as an index (1), but also in a way that transitions between neighboring tiles are smooth (2). This is accomplished naturally in the complex domain, with the sinc function as a universal similarity kernel. Our approach is demonstrated in a model of hippocampus neurons encoding variables such as the location of the animal and also head direction, satisfying conditions (1) and (2). Interestingly, prominent coding properties in hippocampus, such as place cells and spike-phase precession are naturally reproduced by our normative model without building it a priori. Further, our model helps to explain recent observations of conjunctive receptive fields (Høydal et al., 2019), and how to relate computation with manifold structure observed in neural recordings (Low et al., 2018). Specifically, the model makes concrete predictions how the intrinsic structure of neural recordings is influenced by: (i) manifold dimensionality (number of variables), (ii) manifold smoothness and resolution (number of tiles), and (iii) redundant neural dimensions for robustness.

III-51. Phase transitions in vocal development are driven by energy-information balance

Thiago Tarraf Varella¹
Yisi Zhang¹
Daniel Takahashi²
Asif Ghazanfar¹

TVARELLA@PRINCETON.EDU
YZ9@PRINCETON.EDU
TAKAHASHIYD@GMAIL.COM
ASIFG@PRINCETON.EDU

¹Princeton University
²Federal University of Rio Grande do Norte

A better understanding of neuronal operations underlying vocal development requires a priori an understanding of what has changed in their vocalizations and why. Here, we studied vocal development in three different mammals:

humans, bats, and common marmosets. We found that all three undergo at least one sudden transition of the vocalization's acoustics during their development. This transition was best described with the balance of two cost functions estimated with a SoftMax function. Using data from marmoset monkeys, the natural dynamics of these two costs are consistent with the dynamics of the metabolic energy expenditure and the information transmission (as measured by the probability of infant vocalizations eliciting parental vocal responses). Using energy and the information as the two constraints that represent the costs, we were able to predict the differences in transition timing from immature to mature vocalizations. We tested these predictions with experimental data collected while either energy (infant vocalizations produced in helium-oxygen air versus normal air) or information (closed-loop contingent parental vocal playbacks) is manipulated. The model is supported by empirical findings about the vocalization dynamics. The model provides an important framework to test computational properties in neural systems. Understanding how the body and the social context influence vocal output at different ages should guide us in finding the brain mechanisms that are relevant for vocal learning, such as a way to balance different inputs from different sources (e.g. social and motor areas).

III-52. Laminar-specific sensorimotor integration of context in cortical dynamics

Hansem Sohn
Mehrdad Jazayeri

HANSEM@MIT.EDU
MJAZ@MIT.EDU

Massachusetts Institute of Technology

Recent advances in recording, analysis, and modeling of neural population activity have created new vistas for understanding how cortical dynamics perform behaviorally-relevant computations. For example, state-space analysis of neural population activity has provided an explanation of flexible context-dependent computations in terms of a contextual input driving cortical circuits to suitable regions of the state space with the desired latent dynamics. An implicit assumption of such population-level analyses, including recurrent neural network modeling, is that nearby neurons can be treated as samples from a homogeneous ensemble without regard to anatomical information in the microcircuitry. Here, we set out to test the validity of this assumption by analyzing the laminar structure of the geometry and dynamics of cortical population activity. We trained monkeys to reproduce different time intervals sampled from one of two prior distributions, a 'short' prior (480-800 ms) and a 'long' prior (800-1200 ms). We found that the contextual cue specifying the prior condition modulated signals in the dorsomedial frontal cortex (DMFC) and enabled animals to perform context-dependent Bayesian integration. Based on recent work implicating higher-order thalamus in the contextual modulation of cortical dynamics, we hypothesized that the prior-dependent modulation of population activity in DMFC may originate in the superficial layers where thalamo-cortical projections terminate. To test this hypothesis, we analyzed DMFC signals recorded simultaneously from neurons across the cortical laminae. Our initial analyses indicated that the response profile of individual neurons and the dimensionality of the population activity were similar across laminae. However, the key feature of DMFC dynamics that enabled context-dependent computations exhibited laminar specificity. In particular, we found that the context signals were initially present and subsequently amplified in the superficial – not deep – layers. These results suggest that the laminar organization of cortical microcircuits may provide a scaffold for further dissecting how cortical dynamics enable behaviorally-relevant computation.

III-53. Adjudicating between deep neural network models of biological vision with controversial stimuli

Tal Golan
Prashant C Raju
Nikolaus Kriegeskorte
Columbia University

TAL.GOLAN@COLUMBIA.EDU
PRASHANT.RAJU@COLUMBIA.EDU
N.KRIEGESKORTE@COLUMBIA.EDU

Deep neural networks (DNNs) provide the leading model of biological object recognition, but their power and flexibility come at a price: Different DNN models often make very similar predictions. To enable iterative testing and improvement of DNNs as scientific hypotheses about biological vision, we need to efficiently adjudicate between different candidate models.

Here, we suggest using synthetic controversial stimuli to achieve this aim. Controversial stimuli are inputs (e.g., images) whose classifications by two or more models are incompatible. We can efficiently compare DNN models of vision by their predictions of human (or animal) perceptual judgments of controversial stimuli. For each controversial stimulus, by definition, the models disagree about the correct label. The human or animal judgments, therefore, will be incompatible with at least one of the models, providing evidence against it.

To demonstrate our approach, we assembled a diverse set of nine models, all trained on MNIST, whose various solutions cannot all be good models of human vision. For each pair of models, we sampled random noise images and iteratively modified them to increase the incompatibility of their classifications by the two models. This resulted in a stimulus set that causes disagreement among the models. We tested these stimuli on 30 human observers, who rated the presence of each of the ten digits from 0% to 100% in each image. Contrasting each model's outputs with the judgments of each observer revealed that the two generative models we tested, a shallow Gaussian kernel-density estimate and the deep VAE-based Analysis-by-Synthesis model, were considerably and significantly better at predicting the human responses than all the discriminatively-trained DNNs we tested. Yet no model explained all of the explainable variability of the human responses.

We discuss controversial stimuli as a conceptual and practical generalization of adversarial examples that obviates the need for ground-truth labels.

III-54. Surprise learning: fast adaptation in volatile environments with spiking neural networks

Martin Louis Lucien Remy Barry^{1,2}
Carlos Stein N Brito¹
Wulfram Gerstner¹

MARTIN.BARRY@EPFL.CH
CARLOS.STEIN@EPFL.CH
WULFRAM.GERSTNER@EPFL.CH

¹EPfl
²EDNE

Surprise is a neurophysiological response to unexpected events. There is growing experimental evidence that surprise is a key process in learning; surprising information is more memorable and allows quick adaptation to a changing environment.

Recent developments in machine learning achieve high levels of performance in volatile environment tasks. Nevertheless, most of these models do not address biological constraints, such as local update rules, Dale's law and transmission of information via spikes.

Here, we develop a spiking neural network to investigate surprise-based learning in cortical circuits. Following principles of predictive coding, we implement a mismatch detection population learning to balance the excitatory/inhibitory (E/I) inputs. Biological constraints lead us to design a symmetric network capable of associating surprise to both under- and over-estimation of the prediction. The synaptic weights are updated through neo-Hebbian synaptic plasticity (three factor rule) composed of a pre/post local dependency modulated by a global neuromodulator, that depends on surprise. Yet, measuring surprise is not trivial. Motivated by experimental re-

sults showing that a surprising event tends to enhanced brain activity, we use the excess postsynaptic activity, generated by E/I balance violation, to extract a surprise signal from our network.

We applied our model to a task where an agent needs to estimate the time varying transition matrix of a discrete and volatile environment. Our network achieves similar levels of performance as non biologically constrained algorithms, with a relatively simple architecture. Furthermore, the necessity of detailed E/I balance for detecting surprising events suggests an additional purpose for balance in cortical circuits. The results show that an E/I balance violation mechanism combined with a surprise driven three factor learning rule can model how biological neural circuits learn to adapt in non-stationary environments.

III-55. Efficient decoding of large-scale neural population responses with Gaussian-process multiclass regression

C Daniel Greenidge¹
Jacob Yates²
Benjamin Scholl³
Jonathan Pillow¹

DANIEL@DANIELGREENIDGE.COM
JACOBY8S@GMAIL.COM
SCHOLL.BEN@GMAIL.COM
PILLOW@PRINCETON.EDU

¹Princeton University

²University of Rochester

³Max Planck Florida Institute for Neuroscience

In the high-dimensional regime, where the number of neurons exceeds the number of observations, decoding methods are the preferred tool for measuring a neural population's fisher information. An ideal decoder is accurate and computationally inexpensive, so it can be used for cross-validated information estimates. However, common decoders are challenging to fit to noisy, correlated, high-dimensional responses, tend to underestimate information, and can be computationally expensive. To address these issues, we introduce the variational Gaussian Process Multiclass Decoder (vGPMD). The vGPMD is a logistic regression model with Gaussian process smoothing priors over the model weights. The resulting framework encompasses probabilistic population codes for parametric stimuli, and is sensitive to non-diagonal correlation structure. Naive implementations cost $O(N^3K^3)$ for N neurons encoding K stimuli, which is infeasible for modern datasets. To scale the model, we developed a GPU-accelerated stochastic variational inference method using a spectral-domain weight representation that can handle $\sim 100k$ neurons and $\sim 100k$ observations. Our method recovers amplitude and smoothness parameters for each neuron, eliminating the need for ad-hoc preprocessing of untuned neurons. We evaluated performance using data from primary visual cortex in three different animals: macaque monkey (72 stimuli, ~ 150 neurons), ferret (180 stimuli, 320 neurons), and mouse (180 stimuli, $\sim 20k$ neurons). Our model works well even with huge numbers of parameters—the ferret model had 57k weights, and the mouse model had 3.6M. The vGPMD substantially outperformed previous methods, including those commonly-used in the literature: an independent Poisson decoder, a decoder based on support vector machines, logistic regression, and a nonlinear tuning curve decoder. Our model is computationally cheap: it trained in under three minutes for all datasets. Due to its performance and scalability, the vGPMD sets the state-of-the-art for characterizing information content in high-dimensional neural populations and examining the effects of population-level correlations.

III-56. Recurrent neural network dynamics are dependent on learning rule

Brandon McMahan
Michael Kleinman
Jonathan Kao

BMCMAHAN2025@UCLA.EDU
MICHAEL.KLEINMAN@UCLA.EDU
KAO@SEAS.UCLA.EDU

University of California, Los Angeles

Recurrent neural networks (RNNs) have been used as computational models to probe neural dynamics during

cognitive and motor tasks (Mante et al., 2013; Sussillo et al., 2015; Song et al., 2016). However, there exists multiple learning algorithms for training RNNs, raising the question of how training may influence conclusions and insights. To examine this question, we trained RNNs with different learning rules to discriminate noisy inputs, analogous to a random dots motion task (Williams et al., 2018). We trained RNNs using four learning rules: a genetic algorithm (GA), gradient descent via backpropagation-through-time (BPTT), first order reduced and controlled error (FORCE) (Sussillo & Abbott, 2009), and a Hebbian-inspired learning rule (Miconi, 2017).

first, we used tensor component analysis (TCA) to show that RNNs trained with different learning rules have differences in task representation and learning dynamics. Second, we observed that the recurrent weight matrix took on a structure suggestive of two neural populations that self-excite and cross-inhibit when trained with BPTT and GA, but not with the FORCE and Hebbian learning algorithms. finally, we found that the RNNs used different dynamical mechanisms. While the Hebbian trained RNN instantiated a "single attractor" mechanism, the GA RNN instantiated a "three attractor" mechanism comprising two stable and one unstable attractor for every input. Interestingly, the BPTT and FORCE trained networks appeared to instantiate both the single attractor mechanism and the three attractor mechanism depending on the magnitude of the input being fed to the network. We found that RNNs employing the three attractor mechanism were more robust to noisy inputs. In summary, careful consideration should be given to the choice of learning rule when training RNNs. Further, computational motifs that persist across models trained with different learning rules are good candidates for proposing novel computational mechanisms for in vivo neuronal systems.

III-57. Systematic tradeoffs between complexity and accuracy in human decision-making

Alexandre Filipowicz¹
 Eugenio Piasini¹
 Kamesh Krishnamurthy²
 Joseph Kable¹
 Joshua Gold¹

ALEXFILIP@PENNMEDICINE.UPENN.EDU
 EPIASINI@SAS.UPENN.EDU
 KAMESHKK@GMAIL.COM
 KABLE@PSYCH.UPENN.EDU
 JIGOLD@PENNMEDICINE.UPENN.EDU

¹University of Pennsylvania

²Princeton University

In noisy and uncertain environments, we often use past observations to build and update beliefs about environmental statistics to help guide future decisions. However, our cognitive resources are limited, highlighting the need to balance how much past information we use to form beliefs (complexity) with the accuracy of those beliefs (predictive accuracy). To better understand the nature of this balance, we developed a principled, non-parametric method using the information bottleneck to assess both the complexity and accuracy of human beliefs during decision-making. We compute belief complexity as the mutual information between past observations and a subject's responses, and predictive accuracy as the mutual information between the subject's responses and the future observations they are attempting to predict. We compare these behavioral metrics to an optimal bound computed using the information bottleneck, which describes the maximum possible accuracy, given the sequence of observations, for each value of belief complexity. This non-decreasing upper bound on accuracy as a function of belief complexity goes beyond standard ideal-observer approaches by providing a novel benchmark that is conditioned on the complexity of the strategy (or "mental model") that the subject actually used to solve the task. We applied these measures to three separate adaptive decision-making tasks and found a consistent trend: subjects who use more complex beliefs tend to have slightly higher accuracy but also use these beliefs less optimally than subjects with simpler beliefs. These results provide new insights into the information-processing tradeoffs that govern learning and decision-making, which can lead to systematic differences in the extent to which individual decision-makers are willing and/or able to use strategies that become increasingly difficult to optimize as they become more complex.

III-58. Cell-type specific inhibitory plasticity enhances within assembly cooperation and between assembly competition in spiking networks

Fereshteh Lagzi
 Martha Canto Bustos
 Anne-Marie Oswald
 Brent Doiron
 University of Pittsburgh

FERESHTEH.LAGZI@GMAIL.COM
 MARTHABCANTO@GMAIL.COM
 AMOSWALD@PITT.EDU
 BDOIRON@PITT.EDU

It is popular to ascribe distinct network functions to the different inhibitory neuron subtypes that makeup cortical circuits. Carving out functionally determined, cell-type specific circuit wiring is an essential component of this hypothesis. However, the plasticity rules for different interneuron subtypes, and how their interactions shape network dynamics are largely unexplored. We use in vitro patch clamp techniques paired with cell-specific optogenetic stimulation to measure the spike timing dependent plasticity (STDP) rules of the inhibitory inputs onto excitatory (E) pyramidal neurons from both Parvalbumin (PV) and Somatostatin (SOM) interneurons in mouse orbital frontal cortex (OFC). Consistent with past studies, PV inhibition shows a symmetric STDP that is often associated with the homeostatic control of excitatory firing rates. By contrast, SOM inhibition shows an asymmetric Hebbian STDP rule, so that recurrent SOM inhibition onto driving E neurons will ultimately depress. To understand the role of these different plasticity mechanisms in network dynamics we exploit large-scale network simulations of networks of spiking neuron model with plastic synapses, along with associated mean field theories of synaptic dynamics. We show that the asymmetric SOM plasticity rule promotes cross-inhibition between distinct E neuron assemblies, effectively providing a mechanism for competition between functionally grouped principle neurons. This competition will enhance computations where input comparisons must be made, as is often the case in decision task where the OFC is known to be essential. However, strong cross inhibition could lead to extreme (and unstable) winner-take-all assembly dynamics. Fortunately, the symmetric PV plasticity rule provides stability for the circuit, ensuring rich network dynamics. In sum, our work shows how distinct synaptic learning rules support a division of labor amongst the rich diversity of inhibition known to exist in cortex.

III-59. Constraints on the time course of neural population activity

Alan Degenhart¹
 Erinn Grigsby²
 Nicole McClain²
 Patrick Marino²
 Asma Motiwala¹
 Emily Oby²
 Aaron Batista²
 Byron Yu¹
¹Carnegie Mellon University
²University of Pittsburgh

ALANDEGENHART@GMAIL.COM
 ERINN.GRIGSBY@GMAIL.COM
 MCCLAINC91@GMAIL.COM
 PMARINO162@GMAIL.COM
 AMOTIWALA@CMU.EDU
 EMILYOBY@GMAIL.COM
 AARON.BATISTA@GMAIL.COM
 BYRONYU@CMU.EDU

The time course of neural population activity has provided tantalizing evidence of network-level mechanisms, such as point attractors and line attractors, underlying motor control, decision-making, and more. These mechanisms specify that neural activity evolves according to a flow field, which reflects the underlying network constraints. A key prediction suggested by these constraints is that it is difficult for the brain to produce neural activity that is inconsistent with this flow field. To test this prediction, we leveraged a brain-computer interface (BCI) to probe the ability of Rhesus monkeys to volitionally modulate their neural population activity within different 2D projections of a 10D latent space identified by Gaussian-Process Factor Analysis (GPFA). We first identified a 2D projection (an "intuitive mapping") designed to yield proficient control, and asked animals to move a BCI cursor between a pair of targets (A and B) in this 2D space. We observed that even though the A-> B and B-> A paths traversed by the cursor were highly overlapping in the animals' 2D workspace, neural trajectories were distinct in the 10D space,

consistent with the presence of an underlying flow field. We then sought to determine the extent to which animals could violate the empirically-observed flow field. To do this, we defined a "rotated" mapping, a 2D projection of neural activity designed to highlight the neural flow. Using this rotated mapping, animals were then required to attempt to generate cursor trajectories inconsistent with the flow field. We found that it was difficult for animals to violate the observed flow field. These results imply that the underlying network imposes strong constraints on the time course of population activity, and that the neural trajectories extracted from population activity during motor control, decision-making, etc. reflect underlying network mechanisms.

III-60. Metastable attractors in secondary motor cortex underlie self-initiated behavior

Stefano Recanatesi¹

Ulises Pereira²

Masayoshi Murakami³

Zachary Mainen⁴

Luca Mazzucato⁵

STEFANO.RECANATESI@GMAIL.COM

UPO201@NYU.EDU

MASAYOSHI.MURAKAMI@NEURO.FCHAMPALIMAUD.ORG

ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

LUCA.MAZZUCATO@GMAIL.COM

¹University of Washington

²New York University

³Champalimaud Centre for the Unknown, Lisbon, Portugal

⁴Champalimaud Foundation

⁵University of Oregon

The decision to choose specific self-initiated actions and execute them at certain times depends on several deterministic factors including the animal's goals and internal state, and external stimuli. Variability in self-initiated actions may have a stochastic origin as well, as suggested by the large variability in action timing¹. Here, we investigated the neural mechanisms underlying behavioral variability in a self-initiated waiting task. We recorded neural ensemble activity in secondary motor cortex (M2), which was previously shown to encode timing variability in the task². We found that M2 ensemble activity unfolded through sequences of metastable attractors. Attractor onset times showed large variability across-trials, yet reliably predicted upcoming self-initiated actions, allowing us to establish a dictionary between attractors and actions. To explain the mechanism generating reliable sequences of metastable attractors, we proposed a mesoscale circuit model, where M2 is reciprocally coupled to a subcortical area. Transitions between attractors in the model were driven by low-dimensional correlated variability originating from the M2-subcortical feedback loop. This mechanism generated low-dimensional noise correlations, aligned between attractors, which we evinced in the empirical data. While previous work argued that such differential correlations are detrimental for sensory processing³, our results suggest that they are an essential ingredient of motor generation. Our work establishes a framework for investigating the circuit origin of self-initiated behavior, based on recurrent networks supporting attractor dynamics. Our theory suggests a robust mechanism underlying self-initiated actions based on correlated neural variability.

III-61. Neural circuitry supporting approximate inference captures structure of correlations in hippocampus

Francesca Mastrogiuseppe

Naoki Hiratani

Peter Latham

Gatsby Computational Neuroscience Unit, University College London

FRAN.MASTROGIUSEPPE@GMAIL.COM

N.HIRATANI@UCL.AC.UK

PEL@GATSBY.UCL.AC.UK

Because the world is modular and hierarchical, percepts can often be described in terms of a few latent variables - such as the food you're eating (octopus) and the music you're listening to (fado). The co-occurrence of

such variables can be further summarized in terms of more abstract labels, such as the environment you're in (Portuguese restaurant). The brain is capable of reconstructing such a hierarchy of labels from sensory inputs, but how this is done in neural circuits is largely unknown. Indirect cues came from a recent study in rats [McKenzie et al, 2014], which investigated how different states of the world, and thus different sets of latent variables, are represented in hippocampus. It was found that neural activity displays structured signal correlations, which evolve asynchronously in time despite remaining overall stable as the environment becomes richer and new objects are learned. We leveraged this experimental finding to shed light on the algorithmic principles underlying inference of labels in neural circuits. To this end, we considered a hierarchical model of the world which mimics the experimental setup, and maps abstract and sensory labels into noisy neural responses. We then derived two alternative biologically-plausible circuit models which perform Bayesian inference of label values from sensory inputs. The two models implement different computational strategies: one is more conservative, while the other one performs inference in an approximate form, by assuming that the joint posterior distribution over different labels factorizes. We found that only the latter could capture the structure and temporal properties of correlations observed in experiments. It was less accurate, but the performance drop comes with computational advantages: the circuitry is more compact, and labels are represented in a format which better supports generalization.

III-62. From deep learning to mechanistic understanding in neuroscience: revealing computational mechanisms of retinal prediction via model reduction

Hidenori Tanaka¹
 Aran Nayebi¹
 Niru Maheswaranathan²
 Lane McIntosh¹
 Stephen Baccus¹
 Surya Ganguli¹

HDNRTNK8@GMAIL.COM
 ANAYEBI@STANFORD.EDU
 NIRUM@GOOGLE.COM
 LANEMCINTOSH@GMAIL.COM
 BACCUS@STANFORD.EDU
 SGANGULI@STANFORD.EDU

¹Stanford University

²Google Brain

Recently, deep feedforward neural networks have achieved considerable success in modeling biological sensory processing, in terms of reproducing the input-output map of sensory neurons. However, such models raise profound questions about the very nature of explanation in neuroscience. Are we simply replacing one complex system (a biological circuit) with another (a deep network), without understanding either? Moreover, beyond neural representations, are the deep network's computational mechanisms by which synaptic connections and nonlinearities generate behaviorally meaningful computations the same as those in the brain? Without a systematic approach to extracting and understanding computational mechanisms from deep neural network models, it can be difficult both to assess the degree of utility of deep learning approaches in neuroscience, and to extract experimentally testable hypotheses from deep networks. We develop such a systematic approach by combining dimensionality reduction and modern attribution methods for determining the relative importance of interneurons for specific visual computations. We apply this approach to deep network models of the retina, revealing a conceptual understanding of how the retina acts as a predictive feature extractor that signals deviations from expectations for diverse spatiotemporal stimuli. For each stimulus, our extracted computational mechanisms are consistent with prior scientific literature, and in one case yields a new mechanistic hypothesis. Thus overall, this work not only yields insights into the computational mechanisms underlying the striking predictive capabilities of the retina, but also places the framework of deep networks as neuroscientific models on firmer theoretical foundations, by providing a new roadmap to go beyond comparing neural representations to extracting and understand computational mechanisms.

III-63. Learning a texture model for representing cortical area V2

Nikhil Parthasarathy¹
Eero Simoncelli²

NP1742@NYU.EDU
EERO.SIMONCELLI@NYU.EDU

¹New York University

²New York University, Howard Hughes Medical Institute

Neurons in primary visual cortex (V1) respond selectively to oriented edge-like features, and this selectivity can be captured using oriented filters. Such filters can be learned from natural images using unsupervised techniques (eg., ICA or sparse coding), which provide a theoretical explanation for the experimental observations. Cortical area V2 is less well understood, but recent work shows that most V2 neurons respond selectively to 'visual texture'. Moreover, the responses of V2 populations can be used to distinguish between different textures, and to distinguish textures from their spectrally matched counterparts. Although several models have been proposed for V2, these are not learned from natural images, and they have not been shown to provide accurate predictions of V2 population responses to texture images. With these goals in mind, we develop a parametric functional model for V2, and choose its parameters by optimizing a novel self-supervised objective function over a dataset of homogeneous texture images. The model consists of a V1 layer implemented through a set of fixed convolutional filters followed by rectification (simple cells) and pooling (complex cells). These responses are provided as input to a V2 stage that consists of a set of learned convolutional filters followed by the same canonical simple and complex cell transformations. We optimize the filters in the V2 stage according to an objective function that seeks to simultaneously: 1) Minimize variability of V2 responses within individual homogeneous texture images 2) maximize variability of these responses across all images and 3) Minimize the error in locally reconstructing V1 responses from V2 responses. We find that the trained model captures the observed texture selectivities of V2 neurons. Additionally, the model responses provide much better linear predictivity of V2 population responses than is achieved with current state-of-the-art CNNs.

III-64. Evidence that recurrent pathways between the prefrontal and inferior temporal cortex are critical during core object recognition

Kohitij Kar
James DiCarlo

KOHITIJ@MIT.EDU
DICARLP@MIT.EDU

Massachusetts Institute of Technology

Using large-scale neural measurements, Kar and colleagues (2019) reported that behaviorally critical, recurrent computation dependent signals emerge in the late-phase responses of macaque inferior temporal (IT) neurons; ~30 ms after the feedforward pass. However, we do not yet know which brain circuits are most responsible for these recurrent computations. Given that Kar et al. precisely measured the time at which the object identity solutions emerged in IT, these estimated object solution times (OST) provide a strategy to test such critical recurrent circuits. Specifically, longer OSTs suggest additional (putatively recurrent) computations beyond what is achieved by the feedforward response of IT. Therefore, we hypothesize that deficits induced by disruption of recurrent circuit motifs should correlate with the OST. Here we used this strategy to test one such circuit hypothesis. Based on the presence of object-category selective neurons, anatomical projections to IT, and overall experimental feasibility, we focused our efforts to assess the role of "ventral PFC to IT recurrent circuit". We pharmacologically (via muscimol) silenced parts of ventral PFC and simultaneously measured IT population activity in two monkeys, while they performed a battery of object discrimination tasks on previously bench-marked images (n=1320). Our results show that muscimol injections in ventral PFC produced a significant behavioral deficit that is significantly correlated with the object solution times (larger deficits for images with longer OSTs). In addition, PFC inactivation also reduced the quality of the late-phase IT population decodes, commensurate with the behavioral changes. Interestingly, the late-phase IT neural activity after muscimol injection was better explained by a feed-forward ANN compared to the no-muscimol condition. Taken together, these results imply that PFC is a critical part of the recurrent circuitry underlying the production of explicit object representations in IT, used for core object recognition

behavior.

III-65. Parametric Copula-GP model reveals tuning of neuronal and behavioral relationships to visual stimuli

Nina Kudryashova
Theoklitos Amvrosiadis
Nathalie Dupuy
Nathalie Rochefort
Arno Onken
University of Edinburgh

NKUDRYAS@ED.AC.UK
T.AMVROSIADIS@ED.AC.UK
NATHALIE.DUPOUY@GMAIL.COM
N.ROCHEFORT@ED.AC.UK
AONKEN@INF.ED.AC.UK

Recent research in neuroscience shows that neural responses to stimuli are strongly modulated by behavioral context even as early as in sensory brain areas. Yet, the precise functional interactions underlying such modulation remain unknown. The analysis of this modulation proves challenging due to the different statistics and different timescales of the observed variables, e.g. discrete licks or neuronal spikes vs. continuous velocity or calcium signals. In order to overcome these hurdles, we propose parametric copula models with phase-dependent parameters, which separate the statistics of the individual variables from their dependence structure. For model fitting, we use a Bayesian framework with a Gaussian Process (GP) latent variable inference scheme, which is naturally suited to take into account the dependence of the copula parameters on the phase of the experiment (e.g. time or position) as well as the uncertainty of the parameter estimation given limited and variable amounts of data. We applied our model to calcium optical recordings from awake behaving mice performing a spatial navigation task in a virtual reality setup. We have found pairs of variables showing strong non-trivial changes in the dependence structure subject to the position in the virtual reality and related visual stimulus. This finding agrees with previous studies, showing that beyond independent responses, noise correlations carry additional visual information. Moreover, our results reveal that dependence structures are often non-Gaussian, both in noise correlations between neurons and in neurons vs. behavioural variables. By estimating information measures, we also show that modulations of these non-Gaussian tail dependencies can encode additional information about the stimuli. Therefore, contrary to other commonly applied methods, our Copula-GP approach makes stochastic (non-Gaussian) relationships explicit, representing them as the 'tuning curves' of the copula parameters, and generates interpretable models of dependencies between neural responses, visual stimuli, and behavior.

III-66. Cooling of medial septum reveals theta phase lag coordination of hippocampal cell assemblies

Peter Petersen
Gyorgy Buzsaki
New York University

PETERSEN.PETER@GMAIL.COM
GYORGY.BUZSAKI@NYUMC.ORG

Hippocampal theta oscillations coordinate neuronal firing to support memory and spatial navigation. The medial septum (MS) is critical in theta generation by two possible mechanisms: either a unitary 'pacemaker' timing signal is imposed on the hippocampal system or it may assist in organizing subcircuits within the phase space of theta oscillations. We used temperature manipulation of the MS to confront these models. Cooling of the MS reduced both theta frequency and power, was associated with enhanced incidence of errors in a spatial navigation task, but did not affect the spatial map. MS cooling decreased theta frequency oscillations of place cells, reduced distance-time compression but preserved distance-phase compression of place field sequences within the theta cycle. We suggest that reciprocal MS-hippocampal communication is essential for sustaining theta phase-coordination of cell assemblies and, in turn, supporting its role in spatial memory.

III-67. Identifying neural sequences in continuous time through point-process seqNMF

Alexander Williams¹
 Anthony Degleris¹
 Yixin Wang²
 Scott Linderman¹

ALEX.H.WILLIA@GMAIL.COM
 DEGLERIS@STANFORD.EDU
 YIXIN.WANG@COLUMBIA.EDU
 SCOTT.LINDERMAN@STANFORD.EDU

¹Stanford University
²Columbia University

Neural ensembles that successively fire in sparse, temporal sequences are posited to underlie aspects of working memory, motor production, and learning. Extracting and characterizing these sequences in a statistically rigorous manner is challenging, and has been extensively studied in the systems neuroscience community. Recent work by Peter et al. (2017) and Macevicius et al. (2019) proposed a convolutional nonnegative matrix factorization model (termed "seqNMF") to tackle this challenge. A key feature this model is its ability to leverage large numbers of simultaneously recorded neurons to detect sequences, even when individual neurons participate stochastically. Here, we describe several extensions to this model that enable finer scale characterizations of sequences, more rigorous analyses of uncertainty, a broader set of parameter estimation tools, and new strategies to address unaccounted forms of biological variability, such as variations in sequence duration.

Our approach, which we call point-process seqNMF (pp-seqNMF), is specialized for spiking data and utilizes a fully probabilistic, Poisson process modeling framework. This model circumvents the need to bin and smooth spikes times, and thus enjoys similar benefits to recent reformulations of Gaussian Process Factor Analysis (GPFA) in terms of point processes. More importantly, pp-seqNMF is highly flexible—we show that simple modifications allow the model to capture a rich repertoire of variations on each instantiation of a sequence. For example, we can incorporate a per-sequence time rescaling parameter that flexibly models sequences of varying duration. Such variations are critical in many biological contexts, such as hippocampal replay events, which unfold ~15-20 faster than lived experiences. Overall, this work provides a novel and more flexible formulation of a popular model in systems neuroscience, which will improve our ability to detect sequences in noisy spike train data and spur further modeling extensions.

III-68. BehaveNet: behavioral video embedding and neural analysis toolbox

Matthew R Whiteway¹
 Eleanor Batty¹
 Shreya Saxena¹
 Dan Biderman¹
 Taiga Abe¹
 Simon Musall²
 Winthrop Gillis³
 Jeffrey E Markowitz³
 Anne Churchland²
 John Cunningham¹
 Sandeep Robert Datta³
 Scott Linderman⁴
 Liam Paninski¹

MW3323@COLUMBIA.EDU
 ERB2180@COLUMBIA.EDU
 SS5513@COLUMBIA.EDU
 DB3236@COLUMBIA.EDU
 TA2507@COLUMBIA.EDU
 SIMON.MUSALL@GMAIL.COM
 WIN.GILLIS@GMAIL.COM
 JEFFREY_MARKOWITZ@HMS.HARVARD.EDU
 CHURCHLAND@CSHL.EDU
 JPC2181@COLUMBIA.EDU
 SRDATTA@HMS.HARVARD.EDU
 SCOTT.LINDERMAN@STANFORD.EDU
 LIAM@STAT.COLUMBIA.EDU

¹Columbia University
²Cold Spring Harbor Laboratory
³Harvard University
⁴Stanford University

A fundamental goal of systems neuroscience is to understand the relationship between neural activity and behav-

ior. Behavior has traditionally been characterized by task-related variables such as reach direction or response times. More recently, there has been a growing interest in the automated analysis of behavioral videos. One approach is the supervised tracking of body parts, although this requires labeling of training images and may not capture all of the useful information in the video. Another approach is the use of unsupervised linear dimensionality reduction, but this approach often requires a large number of dimensions to capture much of the video variance. Furthermore, neither of these approaches take into account the temporal structure of behavior. Here we introduce a probabilistic framework for the unsupervised analysis of behavioral video and neural activity, building upon previous work [Wiltschko et al. 2015, Johnson et al. 2016, Markowitz et al. 2018]. This framework provides tools for compression, segmentation, generation, and decoding of behavioral videos. We compress videos using convolutional autoencoders (CAEs), which yield a low-dimensional, continuous representation of behavior that requires fewer dimensions than linear methods to obtain similar reconstruction error. We then use an autoregressive hidden Markov model (ARHMM) to segment the CAE representation into discrete “behavioral syllables” based on similarity of dynamics. The resulting generative model can be used to simulate behavioral videos. Finally, we exploit this generative model to construct a novel Bayesian decoder which takes in neural activity and outputs probabilistic estimates of the full-resolution behavioral video. We demonstrate the use of this framework using two popular experimental paradigms and neural recording technologies: Neuropixel multielectrode array recordings from spontaneously behaving head-fixed mice, and widefield calcium imaging from task-engaged head-fixed mice. Our results demonstrate the usefulness of this framework as a path towards the joint modeling of behavior and neural population activity.

III-69. *Drosophila* small object detectors trigger stopping with a novel, displacement-sensitive algorithm

Ryosuke Tanaka
Damon Clark
Yale University

RYOSUKE.TANAKA@YALE.EDU
DAMON.CLARK@YALE.EDU

To survive and reproduce, animals need to be able to detect prey, predators, and conspecifics from a distance. These often appear as small moving objects, and as a consequence, neurons that respond selectively to small moving objects have been found in visual systems of diverse animal species, ranging from vertebrates to insects. However, it has been technically challenging to investigate their behavioral function, computational algorithm, and biological mechanisms – each of which constrains the others – simultaneously in a single system. Object-selective visual projection neurons (OSVPN) in fruit fly *Drosophila*, which is amenable to high-throughput behavioral assays and to cell-type specific neural measurement and manipulation, provide a unique opportunity to address these three functional aspects together. Using genetically targeted synaptic silencing and optogenetics, we found that a type of OSVPN is necessary and sufficient for previously undocumented stopping behavior induced by translating small objects, possibly paralleling the object-induced defensive freezing in rodents. Two-photon calcium imaging revealed that, unlike known object selective neurons in vertebrate retinas or insect brains, the *Drosophila* OSVPN is specifically sensitive to displacement of small objects in its receptive field. We constructed a computational model which features feed forward inhibition and nonlinear spatial pooling, which successfully recapitulated primary response properties of the OSVPN. In addition, by tracing signal transformation using voltage and neurotransmitter sensors, we revealed that center-surround antagonism in upstream neurons and a calcium nonlinearity in the OSVPN combine to implement the observed displacement sensitivity. We also identified an upstream partner of the OSVPN that provides size-tuned, excitatory inputs, using trans-synaptic labeling and optogenetics. Overall, our study illustrates an interesting example of how behaviorally relevant visual features are gradually elaborated across multiple layers of neurons.

III-70. Inter-area interactions encode real-time adaptive biasing of performance during motor skill learning

Lucas Tian¹
W Hamish Mehaffey²
Michael S Brainard²

LUCASYTIAN@GMAIL.COM
W.HAMISH@GMAIL.COM
MSB@PHY.UCSF.EDU

¹Massachusetts Institute of Technology

²University of California, San Francisco

Enhanced neural interactions between higher-order and primary motor areas have been observed during various types of motor learning, but the behavioral function of such interactions remains poorly understood. An intriguing possibility is that increased top-down interactions during learning reflect the development of signals that bias performance in real-time (i.e., by controlling moment-by-moment adaptive variation in motor output) in order to implement more adaptive behavioral variants. Here, we tested this possibility in adult songbirds during vocal learning, by measuring moment-by-moment interactions between a primary motor area (RA) - which controls the production of song acoustic features - and a higher order nucleus (LMAN) - which contributes to the learning of song modifications. We quantified LMAN-RA interactions by measuring the cross-covariance between LMAN and RA spiking activity. The cross-covariance function exhibited an LMAN-leading peak, consistent with a top-down influence of LMAN on RA. The strength of LMAN-RA cross-covariance increased during learning. This increase was specific to behavioral features targeted for learning, including both its precise timing (sub-second) and sequential context. Moreover, analysis of interleaved song renditions revealed a moment-by-moment correlation between the strength of enhanced LMAN-RA cross-covariance and the expression of adaptive pitch changes. These results indicate that enhanced interactions between higher-order (LMAN) and primary motor (RA) circuitry reflect a real-time, top-down motor bias that implements adaptive modifications of song during performance.

III-71. Neural state space geometry in human motor cortex underlying speaking different phonemes

Guy Wilson¹
Sergey Stavisky¹
Donald Avansino¹
Leigh Hochberg²
Jaimie Henderson¹
Krishna Shenoy¹
Shaul Druckmann¹

GHWILSON@STANFORD.EDU
SERGEY.STAVISKY@STANFORD.EDU
DAVANSIN@STANFORD.EDU
LEIGH_HOCHBERG@BROWN.EDU
HENDERJ@STANFORD.EDU
SHENOY@STANFORD.EDU
SHAULD@STANFORD.EDU

¹Stanford University

²Brown University

Neurological disorders such as stroke and amyotrophic lateral sclerosis impair speech in millions worldwide. For afflicted individuals, restoring speech through a brain-computer interface (BCI) is a promising therapeutic approach that could substantially improve quality of life. To access high-quality neural signals, speech BCIs have often used invasive electrocorticography (ECoG) measurements from distributed populations across the cortex (Mugler et al. 2014, Anumanchipalli et al. 2019, Angrick et al. 2019). In this work, we explored a complementary speech decoding approach: accessing higher spatial resolution neural activity using two 96-microelectrode arrays in "hand knob" area of human motor cortex. Using a linear decoder, we achieved 29% classification accuracy (chance = 6%) across 39 English phonemes and 73% accuracy for the best phoneme, comparable with ECoG-based performance on a similar corpus (Mugler et al. 2014). Performance did not saturate with increasing dataset size or channel count, indicating room for further improvement. Classification errors reflected latent phonetic structure, with higher confusion within phonetic groups. We analyzed the structure of phonetic representations at the neural population level. Our analyses reveal a distributed representation, with condition-invariant changes in activity during speech preparation and initiation, followed by divergence by phoneme group and individual

phoneme. Our results 1) suggest that multi-unit activity from nontraditional speech areas is a viable neural signal for speech decoding efforts, carrying implications for the design of future BCI systems, and 2) highlight the strength of natural speech as a lens for studying neural population dynamics during a complex motor behavior.

III-72. Real-time discovery of effective dynamics from streaming noisy neural observations

Yuan Zhao
Josue Nassar
II Memming Park
Stony Brook University

YUAN.ZHAO@STONYBROOK.EDU
JOSUE.NASSAR@STONYBROOK.EDU
MEMMING.PARK@STONYBROOK.EDU

New technologies for recording the activity of large neural populations during complex behavior provide exciting opportunities for investigating the neural computations that underlie perception, cognition, and decision-making. State-space models, that combine an intuitive dynamical system and a probabilistic observation model, provide an interpretable view of complex time series, and inferring such models can provide insights into neural dynamics, neural computation, and development of neural prosthetics and treatment through feedback control. It yet brings the challenge of learning both latent neural state and the underlying dynamical system because neither is known for neural systems a priori. While offline batch analyses, such as expectation maximization (EM) and variational autoencoders, are promising, their offline nature slow down the scientific cycle of experiment, data analysis, hypothesis generation, and further experiments. This highlights the need for new tools to analyze and manipulate neural activity and animal behavior in real time. We developed a novel streaming variational Monte Carlo (SVMC) framework for simultaneously inferring nonlinear latent dynamics and latent states that is applicable to a wide range of state-space models. Modeling unknown dynamics with sparse Gaussian processes allows for Bayesian inference on the dynamical system and efficient computation. Unlike previous approaches, our framework can incorporate non-trivial observational distributions. Constant time and space complexity makes our approach amenable to online learning scenarios and suitable for real-time applications. The aim of this work is to automate analysis and experimental design through feedback stimulation. It can potentially be used in the experiments to perturb neural activity in ways that testably affect behavior, and to track and modify behavior using stimuli designed to influence learning.

III-73. Counterintuitive effects of excitatory and inhibitory perturbations in spiking networks

William Podlaski
Nuno Calaim
Christian Machens
Champalimaud Centre for the Unknown

WILLIAM.PODLASKI@GMAIL.COM
NUNO.CALAIM@NEURO.FCHAMPALIMAUD.ORG
CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

Perturbation studies, in which neural activity in an area of interest is selectively manipulated, have become widespread in experimental neuroscience research. However, the interpretation of such experiments may not always be straightforward – e.g., the stimulation or suppression of a neural population may lead to indirect effects such as compensation from other areas. Here we present a theory of neural circuit responses to perturbations, and use it to explain robustness and sensitivity to different types of excitatory and inhibitory perturbations on neural populations. Using this theory, as well as accompanying spiking network simulations, we find that perturbation responses can change dramatically depending upon the amount and type of feedback in recurrent networks, and are also modulated by the proportion of perturbed neurons and the relationship between their receptive fields. On one extreme, networks with weak feedback, including strictly feedforward networks, exhibit biases in response to both excitatory and inhibitory perturbations, provided the perturbations are sufficiently large. On the other

extreme, in contrast, networks with strong and precisely balancing feedback show a remarkable robustness to inhibitory perturbations, but also a strong sensitivity to very sparse excitatory perturbations. We liken such effects to a tradeoff between robustness and sensitivity – brain circuits may tolerate fragility to some unlikely perturbations in order to ensure robustness to more commonly occurring perturbations. We hypothesize that such a theoretical perspective on robustness may be useful in designing perturbation experiments that exploit sensitivity in order to better distinguish between different neural coding schemes and dynamical regimes. Furthermore, our work could have important implications for the interpretation of recent causal manipulation studies, which have observed both robustness to inhibition and sensitivity to sparse excitation, providing evidence for precise feedback.

III-74. Deep Nose: Using artificial neural networks to represent the space of odorants

Ngoc Tran
Sergey Shuvaev
Daniel Kepple
Alexei Koulakov

Cold Spring Harbor Laboratory

NTRAN@CSHL.EDU
SSHUVAEV@CSHL.EDU
DKEPPLE@CSHL.EDU
KOULAKOV@CSHL.EDU

The olfactory system employs an ensemble of odorant receptors (ORs) to sense molecules and to derive olfactory percepts. We trained artificial neural networks to represent the chemical space of odorants and used that representation to predict human olfactory percepts. We hypothesized that ORs may be considered 3D spatial filters that extract molecular features and can be trained using conventional machine learning methods. We trained an autoencoder, called DeepNose, to deduce a low-dimensional representation of odorant molecules which were represented by their 3D spatial structure. Next, we tested the ability of DeepNose features in predicting physical properties and odorant percepts based on 3D molecular structures alone. We found that despite the lack of human expertise, DeepNose features led to perceptual predictions of comparable or higher accuracy to molecular descriptors often used in computational chemistry. Last, we showed that DeepNose predictions can be readily interpreted to give insights to what functional groups might give rise to certain activity. We propose that DeepNose network can extract molecular features predictive of various bioactivities and can help understand the factors influencing the composition of ORs ensemble.

III-75. Neural encoding characteristics match behaviorally measured prior expectations in visual speed perception

Lingqi Zhang
Alan Stocker

University of Pennsylvania

LINGQIZ@SAS.UPENN.EDU
ASTOCKER@PSYCH.UPENN.EDU

Bayesian inference provides an elegant normative framework for understanding the characteristic biases and discrimination thresholds in visual speed perception (Stocker & Simoncelli 2006). However, the framework is difficult to validate due to its flexibility and the fact that suitable constraints on the prior beliefs and the likelihood functions have been missing. Here we use assumptions of efficient coding to develop a better constrained Bayesian observer model (Wei & Stocker 2015). In the new model, the stimulus distribution links and jointly constrains both the likelihood function and the prior belief. We fit the model to existing psychophysical speed discrimination data, representing discrimination measurements over a broad range of speeds and stimulus uncertainties. Cross-validation confirms that the model fits the data as well as parametric approximations (Weibull fits) of each psychometric curve. The extracted prior beliefs are closely following power-law functions with an exponent of ~ -1 and are much more consistent across subjects than when extracted with the previous model. Furthermore, the efficient coding assumption of the model also makes the specific prediction that for a variable that follows a

power-law distribution, the neural encoding space should be logarithmic. We tested this prediction by analyzing the speed encoding characteristics of a large population of MT neurons (Nover et al. 2005). We show that the prior predicted by the encoding characteristics of the neurons (neural prior) very closely matches the slow-speed power-law prior extracted from the behavioral data (behavioral prior). Our results demonstrate that a Bayesian observer model constrained by efficient coding not only accurately accounts for the behavioral characteristics of visual speed perception, but also provides a normative account of the logarithm encoding of MT neurons (which gives rise to Weber's law) as the efficient neural representation of a power-law distributed perceptual variable.

III-76. Network principles predict motor cortex population activity structure across movement speeds

Shreya Saxena
Abigail Russo
John Cunningham
Mark Churchland
Columbia University

SS5513@COLUMBIA.EDU
RUSSOAA@GMAIL.COM
JPC2181@COLUMBIA.EDU
MC3502@COLUMBIA.EDU

We can often perform the same action at different speeds. Altering movement speed necessitates altering multiple aspects of muscle activity: frequency, amplitude, and overall pattern. How is neural activity structured to enable such changes? To predict the dominant structure of neural activity, we leveraged principles governing the dynamics of recurrent networks. Noise-robust network dynamics require neural trajectories to have 'low trajectory tangling'; small changes in neural state should never be associated with large changes in its derivative. For a rhythmic trajectory, the lowest-tangled geometry is a circle. This yields the first prediction: regardless of the trajectory of muscle activity, the neural trajectory should – across all speeds – be dominated by two-dimensional circular structure. More complex signals should be present in other dimensions but should be small, to avoid increasing tangling. A second prediction arises from the observation that traversing the same trajectory at different speeds would increase tangling. Thus, the near-circular trajectories (one per speed) should be separated by a translation and/or tilt into new dimensions. To test these predictions, we recorded motor cortex population activity and muscle activity as monkeys performed a cycling task. Cycling speed was instructed by visual cues. We first analyzed the data for each speed separately. In the top two principal components, muscle trajectories varied greatly across speeds. Neural trajectories did not, but instead remained consistently circular. Analyzing all speeds together revealed that these circular trajectories were separated by a translation (orthogonal to the main plane of rotation) and tilted modestly into new dimensions. As a result, neural trajectory tangling remained low despite high muscle trajectory tangling. Simulations confirmed that this strategy naturally emerges during network optimization. These results demonstrate that the dominant features of motor cortex activity can be anticipated and explained by considering how network activity should be structured to remain noise-robust.

III-77. Rapid confidence-based hierarchical and counterfactual reasoning

Mahdi Ramadan
Mehrdad Jazayeri
Massachusetts Institute of Technology

MRAMADAN@MIT.EDU
MJAZ@MIT.EDU

With each new spoken word, we make richer and more accurate inferences about what the speaker may wish to communicate. To do so, we retain a measure of confidence over our beliefs so that we can use new information to update and flexibly re-evaluate our inferences. Although many human behaviors depend on such flexibility, we do not know the computational algorithms that allow confidence across a sequence of inferences to inform hierarchical and counterfactual reasoning. A major roadblock in addressing this question has been that behaviors that involve such high-level reasoning are often complex and invoke idiosyncratic strategies that are difficult to

model. Here, we tackled this problem using a simple sensorimotor inference task. Subjects had to infer the path of an invisible ball within a maze based on a sequence of auditory cues that signaled when the ball changed direction. Importantly, the maze was structured hierarchically so that subjects could use each auditory cue to validate or revise their belief about the ball position. Our initial analysis of responses revealed systematic and similar errors across subjects, suggesting that the task was successful at invoking a common strategy. Next, we examined responses in relation to 1) an optimal model that relied on joint probabilities over multiple segments, 2) a model that solved the problem hierarchically, and 3) a hierarchical model that could additionally reason counterfactually. Results rejected the first model but could not distinguish between the second and third models. To distinguish between these models, we analyzed subjects' eye movements. Results provided strong evidence that subjects solved the problem hierarchically and used their confidence to consider counterfactual possibilities. Together, our results validate the role of confidence in hierarchical and counterfactual reasoning, and provide a platform for investigating the underlying computational algorithms.

III-78. A canonical correlation analysis model of the cortical microcircuit

David Lipshutz¹
 Siavash Golkar¹
 Yanis Bahroun¹
 Anirvan Sengupta²
 Dmitri Chklovskii^{3,4}

DLIPSHUTZ@FLATIRONINSTITUTE.ORG
 SGOLKAR@FLATIRONINSTITUTE.ORG
 YBAHROUN@FLATIRONINSTITUTE.ORG
 ANIRVAN@PHYSICS.RUTGERS.EDU
 DCHKLOVSKII@FLATIRONINSTITUTE.ORG

¹flatiron Institute

²Rutgers University

³flatiron Institut

⁴New York University

Cortical pyramidal neurons are part of complex, possibly hierarchical, networks with each neuron receiving feed-forward inputs via its basal dendrites and feedback inputs via its apical tuft. Pyramidal neurons include at least two integrative compartments that generate localized dendritic currents—one in the soma near the basal dendrites and the other in the apical tuft—and differences between localized dendritic currents and inhibitory signals produce calcium plateau potentials that drive synaptic plasticity. We propose a normative model of a cortical microcircuit as an implementation of an online canonical correlation analysis (CCA) algorithm, which yields an analytical solution to a multicompartamental neuron with learning rules that emulate plateau potential dependent synaptic plasticity. Given two datasets, CCA finds projections of the datasets that are maximally correlated. Previous neural network implementations of CCA are either not biologically plausible or rely on deflation.

In our model, two datasets are streamed as activity vectors connecting to the basal and apical dendrites of the pyramidal neurons. At each time step the activity vectors are multiplied by the synaptic weights to yield localized dendritic currents. The pyramidal neuron outputs a normalized sum of the dendritic currents, which is equal to the sum of the projections of the activity vectors onto their common subspace. The dendritic synaptic weight updates depend on differences between localized dendritic currents and inhibitory signals.

III-79. Investigating the ability of astrocytes to drive neural network synchrony

Gregory Handy¹
 Alla Borisyyuk²

GREGORYHANDY@PITT.EDU
 BORISYUK@MATH.UTAH.EDU

¹University of Pittsburgh

²University of Utah

Astrocytes are glial cells that make up 50% of brain volume, with each one wrapping around thousands of synapses. Recent experimental evidence suggests that astrocytes are involved in regulating extracellular con-

centrations of neurotransmitters, buffering of extracellular ions, and neurological disorders. However, the exact role astrocytes have in the governing the dynamics of the synapse and neuronal networks is still being debated. Previous computational modeling work has helped tease out possible mechanisms driving this interaction at the synapse level, with micro-scale models of calcium dynamics and neurotransmitter diffusion. Little computational work has been done to understand how astrocytes may be influencing spiking patterns and synchronization of large networks, partly because it is computationally infeasible to include the intricate details found in this previous work in such a network-scale model.

We overcome this issue by first developing an “effective” astrocyte that can be easily implemented to already established network frameworks. We do this by showing that the astrocyte proximity to a synapse makes synaptic transmission faster, weaker, and less reliable. Thus, our “effective” astrocytes can be incorporated by considering heterogeneous synaptic time constants, which are parametrized only by the degree of astrocytic proximity at that synapse. This parametrization makes sense in light of experimental evidence showing that the degree of astrocyte ensheathment varies by brain region and that it is a crucial component in certain disease states such as some forms of epilepsy. We then apply our framework to a network of 20,000 exponential integrate-and-fire neurons, similar to the one presented by Rosenbaum et al. (2017). Depending on key parameters, such as the number of synapses ensheathed, and the strength of this ensheathment, we show that astrocytes have the ability to push the network to a synchronous state and exhibit spatially correlated patterns.

III-80. Effects of structured neural correlations in population coding: beneficial or detrimental?

Seyedamin Moosavi¹

Magalie Tatischeff²

Bingyue Zhu²

Hideaki Shimazaki¹

MOOSAVI.SEYEDAMIN.7X@KYOTO-U.AC.JP

MAGALIE.TATI@GMAIL.COM

RENAISCO.Z@GMAIL.COM

HIDEAKI.SHIMAZAKI@GMAIL.COM

¹Kyoto University

²The University of British Columbia

In sensory cortices, information about external stimuli is conveyed by collective neural activities. Small neural populations often express perceptual sensitivity of psychophysical observers, suggesting that stimulus properties are redundantly coded by correlated neurons. This implies the existence of ‘information-limiting’ structures in noise correlations, which limit information as the population size increases. However, the existence of such correlations and their underlying mechanisms in cortical activities are yet largely unexplored. Information conveyed by neural activities can be signified as the linear fisher information that quantifies the sensitivity of population firing to stimulus changes. Here, from the viewpoint of information-geometry, we define detrimental correlations that decrease the linear fisher information, and theoretically underpin the information-limiting correlations as a subset of the detrimental correlations. This generalizes the so-called differential correlations that were originally suggested as an exclusive form of the information-limiting correlations. We then show that detrimental correlations appear in activity of orientation-selective neurons in the monkey V1 area responding to drifting gratings [Kohn & Smith CRCNS.org]. Interestingly, observed correlations of small populations (~10 neurons) enhance sensitivity to the stimulus orientation, rather than decreasing its information. This positive effect was observed for covariance structures of both spike counts and binary spikes in fine temporal windows (~10-100 ms), even after removing stimulus-driven co-variations caused by drifting of the sinusoidal gratings, using state-space models. This result implies that the correlations caused by the visual pathways act positively in coding the stimulus. Nevertheless, we find small structured correlations in coarse temporal-scales which have a detrimental role in population coding. The detrimental correlations appear as we increase the temporal window size of the analysis, indicating that they are caused by spatio-temporal correlations in longer time-scales. In summary, beneficial and detrimental correlations exist in respective time-scales of cortical activities, and the latter potentially limits the information in larger populations.

III-81. A RNN that runs forward and backward: A new approach to biologically plausible sequence learning

Huzi Cheng
Joshua Brown
Indiana University Bloomington

HUZCHENG@IU.EDU
JWMBROWN@INDIANA.EDU

The canonical training algorithm of widely used RNNs in computational neuroscience, backpropagation through time (BPTT), is biologically implausible partly because it requires neurons to store their complete activity history. Various approaches have been proposed to address this problem, but it remains unclear how real neurons might perform supervised sequence learning with purely localist information. One clue is seen in hippocampal research indicating that some neurons replay their temporal sequences in reverse. This phenomenon may represent an analog of backward temporal unfolding of neural dynamics in the BPTT and thus may play an essential role in learning. Inspired by these neurophysiological results, we developed a new RNN architecture that unfolds its trajectories like BPTT and learns tasks, but without violating biological constraints of local information. Our model is composed of a forward circuit (FC) coupled with a backward circuit (BC). The FC plays sequences forward. Later, the RNN uses its BC to reconstruct its complete forward dynamics in reverse, to compute error and activity signals for each previous step. Then the FC and BC processes proceed in an interweaving manner: the BC consolidates the updated FC and the FC offers new trajectories for the BC to learn. We tested this model on several tasks and found it can 1) reconstruct historical neural dynamics (reproduce the reverse play in the hippocampus); 2) memorize sequences based on the reconstructed dynamics and 3) learn tasks requiring temporal information integration. Remarkably, we found that reverse play in the network's resting period improves its performance in tasks, which matches a series of observations from the hippocampus. Our model offers new possibilities in building BPTT-like but plausible sequential learning models of the brain. Also, its behavior suggests a way to bridge the explanatory gap between observed phenomena in the hippocampus and the specific learning mechanisms behind them.

III-82. Generalized theory of dynamic balance with arbitrary low-rank structure

Itamar D Landau¹
Haim Sompolinsky²

ITAMAR.LANDAU@MAIL.HUJI.AC.IL
HAIM@FIZ.HUJI.AC.EDU

¹Hebrew University of Jerusalem

²Hebrew University of Jerusalem, Harvard University

The dynamic balance of excitation and inhibition is a paradigmatic theory for describing the activity of neocortical networks. We develop a general theory of dynamic balance in recurrent networks, extending the excitatory-inhibitory structure to arbitrary low-rank structure. The network is driven by strong feed-forward input which must be balanced by structured recurrent connectivity in order to prevent saturation. Recurrent connectivity consists of a strong structured, low-rank, component which yields potentially large input current to a linear subspace within the network and an additional random component. We derive conditions on the structured, low-rank, component of connectivity in order to balance the feed-forward input and enable the random component to propagate fluctuations and drive chaos. In order to achieve this we decompose the dynamics into the structured subspace and its orthogonal complement. We study the low dimensional effective recurrent connectivity within the structured subspace and find that in order to achieve the fluctuating chaotic activity of the dynamic balance regime this low dimensional connectivity must be properly aligned. We show analytically and numerically that dynamic balance is prevented by misaligned connectivity.

Our generalized theory with arbitrary low-rank structure can incorporate biological realism into the framework of dynamic balance, including different cell-types, layered/columnar anatomy, distance-dependence, functional connectivity, and heterogeneity between neurons. From another perspective, low-rank structure has been of

recent interest in designing networks that perform specific computations (e.g. Mastrogiuseppe and Ostojic 2017). Our work studies a new regime where low-rank connectivity is strong and creates a structured balance constraint, providing a link between networks exhibiting dynamic balance and those designed to perform computations. This may offer a broad framework for designing networks to perform dynamical input-output computations.

III-83. The complete synaptic-resolution connectome of a *Drosophila* brain center for navigation

Vivek Jayaraman¹
 Shin-ya Takemura¹
 Romain Franconville¹
 Hannah Haberkern¹
 Chuntao Dan¹
 Brad Hulse¹
 Marcella Noorman^{1,2}
 Ruchi Parekh¹
 Lou Scheffer¹
 Stephen Plaza¹
 Ann Hermundstad¹
 Sandro Romani¹
 Gerry Rubin¹
 Dan Turner-Evans¹

VIVEK@JANELIA.HHMI.ORG
 TAKEMURAS@JANELIA.HHMI.ORG
 FRANCONVILLER@JANELIA.HHMI.ORG
 HABERKERNH@JANELIA.HHMI.ORG
 DANC@JANELIA.HHMI.ORG
 HULSEB@JANELIA.HHMI.ORG
 NOORMANM@JANELIA.HHMI.ORG
 PAREKHR@JANELIA.HHMI.ORG
 SCHEFFERL@JANELIA.HHMI.ORG
 PLAZAS@JANELIA.HHMI.ORG
 HERMUNDSTADA@JANELIA.HHMI.ORG
 ROMANIS@JANELIA.HHMI.ORG
 RUBING@JANELIA.HHMI.ORG
 TURNEREVANS@JANELIA.HHMI.ORG

¹HHMI Janelia Research Campus

²Computation and Theory

The topology of a network is often informative about its function. Extracting detailed neural network structure has recently become possible for the fly brain through advances in electron microscopy (EM) and automatic reconstruction. Here we present the first complete EM-based connectome of a highly conserved insect brain region called the central complex (CX), including all its neurons and most of their synaptic connections. This highly recurrent central brain region, which is composed of ~3000 identified neurons enables flies to maintain an arbitrary heading over kilometers, form visually guided spatial memories, use internal models of their body size when performing motor tasks, and consolidate memories during sleep. CX neurons show clear signatures of ring attractor dynamics during navigation in visual environments and in darkness and are modulated by past experience, satiety, circadian rhythm and sleep state. Many conceptual and computational models have explored how the CX might perform some of these functions. However, these models have largely relied on anatomical overlap and functional connectivity to construct putative CX circuits. Our data shed new light on the region, revealing parallel sensory pathways to the CX, each of which appears separately normalized by gain control driven by mutual inhibition, a sub-network downstream of these inputs that provides the network substrate for a ring attractor, and structured recurrence that seems well-suited to perform vector-based computations, path integration, and heading-dependent action selection with contextual modulation. We also found locally recurrent motifs involving multiple neurons, each with mixed pre- and post-synaptic specializations within single compartments. Such recurrence may allow the fly brain significantly greater computational capacity than suggested by the small number of its neurons. Overall, the CX connectome reveals many functional motifs for navigation and demonstrates the power of analyzing anatomical structure to generate hypotheses for function in the brain.

III-84. Timescales of ongoing activity reflect local connectivity and are modulated during attention

Roxana Zeraati¹

Yanliang Shi²

Marc Gieselmann³

Nicholas Steinmetz⁴

Tirin Moore⁵

Alexander Thiele³

Tatiana Engel²

Anna Levina⁶

ROXANA.ZERAATI@TUEBINGEN.MPG.DE

SHI@CSHL.EDU

ALWIN.GIESELMANN@NCL.AC.UK

NICK.STEINMETZ@GMAIL.COM

TIRIN@STANFORD.EDU

ALEX.THIELE@NCL.AC.UK

ENGEL@CSHL.EDU

ANNA.LEVINA@UNI-TUEBINGEN.DE

¹University of Tübingen

²Cold Spring Harbor Laboratory

³Newcastle University

⁴University of Washington

⁵Stanford University

⁶University of Tübingen

Ongoing cortical dynamics unfold across different temporal scales. These timescales reflect the network's specialization for task-relevant computations. However, it is unknown how different timescales emerge from the spatial network structure and whether they can be flexibly modulated by cognitive demands, e.g., during attention. We developed a network model which consists of binary units representing local neural populations (mini-columns) with spatially structured connections among them. We find that activity of the mini-columns exhibits two distinct timescales arising from the network dynamics. The first timescale is induced by recurrent excitation within a mini-column (vertical connectivity), and the second timescale is induced by interactions among mini-columns (horizontal connectivity). The timescales depend on the network topology, and the second timescale disappears in networks with random connectivity.

To test model predictions, we analyzed spiking activity recorded from single cortical columns in the primate areas V1 and V4 during an attention task. We developed a novel method based on adaptive Approximate Bayesian Computations, which estimates the timescales from spiking activity and overcomes statistical biases due to finite sample size. We observed two timescales in both V1 and V4 population dynamics. Both timescales were significantly longer in V4 than in V1, which is explained by our model based on differences between V1 and V4 network properties. Moreover, the V1 and V4 timescales were longer when attention was directed toward neurons' receptive-fields. This result reveals how ongoing network dynamics is influenced during top-down attention even without measurable modulations of firing rates in the absence of visual stimuli. Based on our model, modulation of timescales arises from an increase in efficacy of vertical connections and a slight suppression of horizontal interactions. Our results suggest that timescales of local neural dynamics emerge from the spatial network structure and can flexibly change due to top-down influences according to task demands.

III-85. Probabilistic feedback updates priors in sequential decisions with accumulated evidence

Khanh Nguyen¹

Zachary Kilpatrick²

Kresimir Josic¹

KPNGUYEN@MATH.UH.EDU

ZPKILPAT@COLORADO.EDU

JOSIC@MATH.UH.EDU

¹University of Houston

²University of Colorado, Boulder

To make the best decisions organisms must flexibly accumulate information, accounting for what is relevant, and ignoring what is not. Many decision-making studies focus on sequences of independent trials in which

the evidence gathered to make a choice, as well as the resulting actions and feedback, are irrelevant to future decisions. To understand decision-making under more natural conditions, we propose and analyze models of ideal observers who accumulate evidence to freely make choices across a sequence of correlated trials and receive noisy feedback.

Two-alternative forced choice tasks (2AFC) are often used to characterize strategies subjects use to make decisions. Normative theories have been developed for such tasks when rewards provide the sole evidence (e.g., two-armed bandit tasks). Less is known about how observers should integrate probabilistic rewards interspersed with noisy evidence, to inform their decisions in future, correlated trials. To understand decision-making in such situations, we describe the behavior of an ideal observer and ask how actual subjects' decision-making strategies could deviate from optimality.

We extend drift-diffusion models to obtain the normative form of evidence accumulation in a series of 2AFC trials with the correct choice evolving as a two-state Markov process. We analyze 3 different feedback cases: no trial-to-trial feedback, probabilistic trial-to-trial feedback, and probabilistic reward as feedback. Ideal observers integrate noisy evidence within a trial until reaching a decision threshold and bias their initial belief depending on the evidence accumulated, and feedback received on previous trials. Ideal observers optimize their reward rate across trials by adjusting their sequence of decision thresholds to deliberate longer on early decisions and respond more quickly in later trials. We show how conflicts between feedback and accumulated evidence are resolved, and under which conditions one of them dominates. Our findings are consistent with experimentally observed response trends.

III-86. Correlation-based spatial layout of deep neural network features generates ventral stream topography

Eshed Margalit¹
Hyodong Lee²
Tiago Marques^{2,3}
James J DiCarlo²
Dan Yamins¹

ESHED.MARGALIT@GMAIL.COM
HYO@MIT.EDU
TMARQUES@MIT.EDU
DICARLO@MIT.EDU
YAMINS@STANFORD.EDU

¹Stanford University

²Massachusetts Institute of Technology

³Brain and Cognitive Sciences

The primate visual system is organized into functional maps, including pinwheel-like arrangements of orientation-tuned neurons in primary visual cortex (V1) and patches of category-selective neurons in higher visual cortex. Recent work has demonstrated that deep convolutional neural networks (DCNNs) trained for object recognition are good descriptors of neural representations throughout the ventral pathway, with early, intermediate, and late cortical brain areas best predicted by corresponding layers of the DCNN. Despite this success, DCNNs have no inherent spatial layout for features at a given retinotopic location, and thus, make no predictions regarding many of the characteristic topographic phenomena observed in the brain beyond retinotopy itself, e.g., pinwheels and patches. Cortical map formation has been modeled using self-organizing maps that leverage principles of wiring-length minimization and local correlations of unit responses to produce topographic structure. However, these methods rely on simplified feature parameterizations that limit their ability to accommodate more realistic descriptions of neuron response properties, especially in higher visual areas. Here, we augment DCNNs by assigning model units spatial positions in a 2D "cortical sheet" and introduce a novel algorithm to arrange units so that local response correlations are maximized. Applying this algorithm to a categorization-optimized DCNN, we find that layouts generated from earlier layers recapitulate core features of V1 orientation, spatial frequency, and color preference maps, while those generated from later layers naturally exhibit category-selective clusters. Because this wide range of apparently disparate phenomenology is produced by the same underlying principle, our results suggest that the functional architecture of the visual system can be explained by two fundamental constraints: the need to perform visual tasks and the pressure to minimize biophysical costs such as wiring length. Our framework for spatially mapping DCNNs integrates biophysical and representational phenomenology,

allowing a more unified understanding of the visual system's functional architecture.

III-87. Spectral cues are essential for the auditory azimuthal map in the mouse superior colliculus

Shinya Ito
Yufei Si
David Feldheim
Alan Litke

University of California, Santa Cruz

SHIXNYA@GMAIL.COM
YUFEI.SI@UCSC.EDU
DFELDHEI@UCSC.EDU
ALAN.LITKE@CERN.CH

Sound localization plays a critical role in animal survival. To compute the incident sound direction, the animal can use three cues: interaural timing differences (ITDs), interaural level differences (ILDs) and the direction-dependent spectral filtering of the sound by the head and pinnae (spectral cues). Compared to ITDs and ILDs, little is known about how spectral cues contribute to the neural encoding of auditory space. Here we report on auditory space encoding in the mouse superior colliculus (SC). We show that the mouse SC contains neurons with spatially-restricted receptive fields (RFs) that form a topographic map of azimuthal auditory space. By freezing each sound localization cue in the stimuli, we found that nasal RFs require spectral cues and temporal RFs require ILDs. Therefore, the lack of either cue results in the disruption of the azimuthal topographic map. These results demonstrate an unexpected role of spectral cues in azimuthal sound localization and indicate a novel computational aspect of the auditory system of the mouse.

III-88. Efficient coding in large networks of spiking neurons with tight-balance

Jonathan Timcheck
Jonathan Kadmon
Surya Ganguli
Stanford University

TIMCHECK@STANFORD.EDU
KADMONJ@STANFORD.EDU
SGANGULI@STANFORD.EDU

Cortical neurons often exhibit irregular activity and unreliable response patterns, raising the fundamental question of how accurately such circuits can encode signals. Classically, one expects that for a population of N neurons with weakly correlated Poisson spiking, coding error will scale as $1/\sqrt{N}$. However, a recent efficient coding framework demonstrated spiking neural networks with coding error decreasing linearly with N . The key idea is that whenever the decoding error exceeds a threshold, an individual neuron is always available to spike, correct this error, and then rapidly inhibit other neurons from over-correcting. This framework raises several fundamental questions: (1) can we achieve such superclassical error scaling in rate networks, or is the spiking mechanism essential; (2) how do delays in spike transmission and stochasticity in spike generation affect decoding error; (3) under what conditions does superclassical scaling survive; (4) how does error scale with the dimensionality D of the input signal?

We answer these questions through both analytics and matching simulations. First, we demonstrate analog rate networks with thresholding nonlinearities (i.e. the ReLU nonlinearity) can exhibit superclassical scaling, thereby greatly expanding the generality of the efficient coding framework. Second, we derive an analytic theory for transmission delays and noise in spiking networks. We find for any given delay an optimal noise level that minimizes decoding error; it arises from two competing effects: the need to introduce some heterogeneity to prevent delay-induced synchronization, but not so much so as to destroy coding fidelity. Third, our theory reveals that decoding error scales as the square-root of delay divided by N . Thus, if delays scale inversely with any power of N , superclassical scaling survives in large networks. Fourth, we generalize our theory to higher dimensional inputs, finding that decoding error scales linearly with the dimensionality.

III-89. Hippocampal replay rapidly and repeatedly adapts to reconfigurations of barrier wall structure in a changing complex maze

John Widloski
David Foster

JOHNWIDLOSKI@GMAIL.COM
DAVIDFOSTER@BERKELEY.EDU

University of California, Berkeley

Hippocampal replay has been proposed as an important substrate for path planning and learning because it encodes behaviorally relevant trajectories. However, there is little work on how replay responds to environments with barriers. Barriers change the topological structure of the environment and can have profound effects on navigation. To address this, we developed a new task to elicit replay through barriers. We adapted an open field, goal-directed task to include transparent, jail bar barriers that impeded the rat's path but minimized interruption of visual and olfactory cues. The 6 barriers could be fitted into 12 different positions, giving a total of 924 unique configurations; new configurations were chosen pseudorandomly for each session, and each rat received between 30-70 training sessions before recording. All other features of the environment remained fixed. The task consisted of goal directed trials to a fixed unmarked "home" well (which changed across sessions) alternating with cued trials to random wells backlit by a flashing light which were designed specifically to lure the replay through the barriers. We recorded up to 376 place cells simultaneously in dCA1 using high tetrode-density hyperdrives. Remarkably, across barrier configurations, awake replays rarely crossed the barriers. To avoid being biased toward smooth trajectories that might not have been as readily detected because of sampling issues at the barriers, we developed a new quantitative method to measure and visualize replay-like representational transitions in all candidate replay events and confirmed the finding. Since multiple sessions were recorded each day, we were able to assess the relationship between place field remapping and barrier-respecting replay. We found that most place cells do not remap between configurations, suggesting that the same place cells with the same place fields rapidly reconfigure their relationships to one another to support awake replay that obeys the topological constraints of each environment.

III-90. Top-down attention as efficient perceptual inference

Wiktor Młynarski
Gasper Tkacik

WMLYNARS@IST.AC.AT
GASPER.TKACIK@IST.AC.AT

Institute of Science and Technology Austria

Top-down attention is hypothesized to be a process which dynamically allocates limited neural resources to task-relevant computations. Such strategy could enable the brain to perform multiple tasks without exceeding its energetic and computational constraints. Consistently with this idea, sensory neurons are driven not only by external stimuli but also by feedback signals from higher brain areas, allowing sensory code to be adapted to the goals of the organism and its belief about the state of the environment. While this hypothesis inspired multiple case-specific models of top-down attention, a unified theoretical perspective which derives top-down attentional effects from a general design principle has been lacking. In particular, it remains unclear how a perceptual observer operating under uncertainty and resource constraints should dynamically reorganize sensory representations to make accurate inferences about the changing environment.

Here we propose a candidate normative theory of top-down attention, which builds on a recent synthesis of probabilistic inference and efficient coding. As an example, we use a canonical model of V1 neurons extended with nonlinear threshold functions. A downstream observer decodes the image from V1 activity and uses it for Bayesian inference of a task-relevant latent variable. The observer can dynamically modulate firing thresholds of individual neurons in V1 via feedback connections. We demonstrate that key attentional processes such as object, feature and spatial attention emerge from the same design principle: maximization of inference accuracy at minimal metabolic cost. The model generates quantitative hypotheses about the task specificity of neural correlations and about the impact of perceptual uncertainty on sparsity and dimensionality of the population code. Our results provide a theoretical framework for analysis of dynamic, attentionally-modulated sensory codes.

III-91. Subpopulation coding reveals a mechanism for improved information flow through cortical circuits

Matthew Getz
Chengcheng Huang
Brent Doiron
University of Pittsburgh

MPG39@PITT.EDU
HUANGC@PITT.EDU
BDOIRON@PITT.EDU

Cortical state is modulated by myriad cognitive and physiological mechanisms [3]. Yet it is still unclear how changes in cortical circuit activity relate to changes in neuronal processing. Previous studies have reported state dependent changes in response gain or population-wide shared variability, motivated by the fact that both are important determinants of the performance of any population code. However if the operating regime of cortex is well-described by a linear response function (as is often the case), then the linear fisher information (fi) about a stimulus available to a decoder is invariant to changes in the linearization about the operating point. This is because any state-dependent changes in response gain are neutralized by changes in population variability. In this study we show that by contrast, when one restricts a decoder to a subset of a cortical population, information within the subpopulation can increase through modulation of cortical state. A clear example of such a subpopulation code is one in which decoders only receive projections from excitatory cells in a recurrent excitatory/inhibitory (E/I) network. We demonstrate the counterintuitive fact that when decoding only from E cells, it is exclusively the I cell response gain and connectivity which governs how information flow changes. In particular, increased fi is achieved through an interplay between inhibitory response gain and inhibition cancelling shared variability or reducing independent noise. Our results raise the possibility that inhibitory circuitry forms a key component in modulating information flow in recurrent cortical networks.

III-92. Population coding and network dynamics during OFF responses in auditory cortex

Giulio Bondanelli¹
Thomas Deneux²
Brice Bathellier²
Srdjan Ostojic¹

GIULIO.BONDANELLI@ENS.FR
DENEUX@UNIC.CNRS-GIF.FR
BRICE.BATHELLIER@CNRS.FR
SRDJAN.OSTOJIC@ENS.FR

¹Ecole Normale Supérieure
²Université Paris-Sud

Across sensory systems, complex spatio-temporal patterns of neural activity arise following the onset (ON) and offset (OFF) of simple stimuli. While ON responses have been widely studied, the mechanisms generating OFF responses in cortical areas have so far not been fully elucidated [Kopp- Scheinpfu 2018, Scholl 2010]. Recent studies have argued that OFF responses reflect strongly transient sensory coding at the population level [Mazor 2005], suggesting they may be generated by a collective, network mechanism. We examine here the hypothesis that OFF responses are single-cell signatures of network dynamics and propose a simple model that generates transient OFF responses through recurrent interactions. To test this hypothesis, we analyze the responses of large populations of neurons in auditory cortex (AC) recorded using two-photon calcium imaging in awake mice passively listening to auditory stimuli. Adopting a population approach, we find that the OFF responses evoked by individual stimuli define trajectories of activity that evolve in low-dimensional neural subspaces. A geometric analysis of the population OFF responses reveals that, across different stimuli, these neural subspaces lie on largely orthogonal dimensions that define low-dimensional transient coding channels. We show that a linear network model with recurrent interactions provides a simple mechanism for the generation of the strong OFF responses observed in auditory cortex, and accounts for the low- dimensionality of the transient channels and their global structure across stimuli. We finally compare the network model with a single-cell mechanism and show that the single-cell model, while explaining some prominent features of the data, does not account for the structure across stimuli captured by the network model.

III-93. A motor cortical model of brain-machine interface learning, fast and slow

Jorge Menendez¹
 Jay Hennig²
 Matt Golub³
 Emily Oby⁴
 Aaron Batista⁴
 Steven Chase²
 Byron Yu²
 Peter Latham¹

J.AUDI11@GMAIL.COM
 JHENNIG@ANDREW.CMU.EDU
 MGOLUB@STANFORD.EDU
 EMO22@PITT.EDU
 AARON.BATISTA@PITT.EDU
 SCHASE@ANDREW.CMU.EDU
 BYRONYU@CMU.EDU
 PEL@GATSBY.UCL.AC.UK

¹Gatsby Computational Neuroscience Unit, University College London

²Carnegie Mellon University

³Stanford University

⁴University of Pittsburgh

An extraordinary property of mammalian motor systems is their capacity to flexibly and quickly adapt to novel environments. A remarkable demonstration of this is primates' ability to acquire proficient control of a brain-machine interface (BMI). What are the algorithms underlying this learning process? An important clue is offered by the finding that primates can learn to use certain motor cortical (M1) BMI decoders with just a few hundred trials of practice (Sadler et al. '14, Nature), while other decoders require many thousands (Oby et al. '19, PNAS). Here, we sought to build a mechanistic model of the changes underlying these two timescales of learning. We argue that the short timescales of learning are consistent with a "re-aiming" learning strategy whereby upstream M1 inputs are modified within a task-relevant low-dimensional space, while the local M1 circuit is left unchanged. The low dimensionality of the search space can allow for efficient optimization, but also imposes constraints on the BMI decoders that can be learned. By explicitly modelling "re-aiming" solutions, we show that these constraints are consistent with the results of Sadler et al. We additionally find that the population activity generated by these solutions maintains its distribution after learning, as is observed empirically (Golub et al. '18, Nat Neuro). The model also makes a novel experimental prediction: an asymmetry in the set of achievable BMI velocities, such that BMI reaching ability should depend strongly on the direction of the reach. This asymmetry follows from the nonlinear dynamics inherent in neural circuits. Upon re-analysis of the data of Sadler et al., we indeed find that this prediction holds. Finally, we show that a similar mechanism is consistent with longer timescales of learning, whereby learning comprises optimizing upstream M1 inputs within a higher-dimensional space encompassing the full space of natural motor commands.

III-94. A novel method to encode sequences in a computational model of speech perception

Meropi Topalidou
 Gregory Hickok
 University of California, Irvine

MTOPALID@UCI.EDU
 GREG.HICKOK@UCI.EDU

The ability to sequence at the phoneme, syllable, and word level is essential to speech production. Here, we present a computational model of speech perception that learns the mapping between sound sequences and representations of individual words. It contains two bidirectionally connected structures that onto the cortical regions of mid-posterior superior temporal sulcus/middle temporal gyrus (pSTS/pMTG) for the lexical component, and posterior superior temporal gyrus (pSTG) for the auditory-phonological one. The basic idea of the model is that the "word" at the lexical level, and its "phonemes" at the auditory level are connected by synaptic weights, where the first element of the sequence is more strongly connected with the word than the second one and so on. This is essentially what Karl Lashley proposed in 1951; the serial order of the sequence is encoded into the activity level of each unit. This eliminates the need for buffers or position slots used by other models

([2], [3]). Another advantage of our model is that it doesn't include a separate working memory to explicitly store symbolic information. Furthermore, we demonstrate how the proposed sequence encoding in the weights emerges as a result of initial learning of auditory-lexical association. This results from (i) the different maximum activity of the auditory-phonological units on each trial, emerging from a soft winner-take-all mechanism, and (ii) Hebbian-learning, where the more active a pre- and a post-synaptic unit are, the stronger they will be connected. A limitation of the model is that cannot represent a sequence with duplicated elements; although, this can be overstepped by adding hierarchical organization at the lexical layer. To conclude, our model proposes a new method for encoding sequences in speech perception, that can be embedded in a broader model of sensorimotor planning for speech perception and production.

III-95. Bayesian inference through attractor dynamics in medial entorhinal cortex

Malcolm Campbell
Alexander Attinger
Samuel Ocko
Surya Ganguli
Lisa Giocomo
Stanford University

MALCOLMC@STANFORD.EDU
ATTIALEX@STANFORD.EDU
SAMOCKO@GMAIL.COM
SGANGULI@STANFORD.EDU
GIOCOMO@STANFORD.EDU

While it is widely believed that the brain performs approximate Bayesian inference to reconstruct properties of the world from noisy or ambiguous sensory input, how these algorithms are implemented in brain circuits is not fully understood. Here, using spatial maps in medial entorhinal cortex (MEC) as a model system, we explored whether attractor networks can implement Bayesian inference to mediate conflicts between path integration and landmarks during navigation. first, we showed that when conflicts are small, an attractor model can implement a Kalman filter to resolve disagreement between the two position estimates. In this "Kalman approximation," conflicts between path integration and landmarks cause spatial maps to coherently shift without remapping, reflecting a weighted average that combines the two position estimates according to their relative certainties. Using large-scale recordings of hundreds of MEC neurons while mice experienced conflicts between path integration and landmarks in virtual reality, we confirmed that MEC spatial maps often coherently shifted without remapping when conflicts were small. However, we observed significant variability across individual sessions, with many sessions showing partial remapping—a breakdown of the Kalman approximation. Reducing visual contrast led to increased shifts but less remapping during cue conflicts, indicating that the system moved closer to the Kalman approximation but with stronger weighting of path integration relative to landmarks. These results point to a two-step inference process in which MEC first infers the degree of cue conflict, remapping if the conflict is large enough, and then uses a Kalman filter to infer position within the chosen map. finally, we recorded from additional brain regions important for visually-guided navigation—visual cortex, retrosplenial cortex, and hippocampal CA1—and found much lower levels of shifting and remapping. This suggests that the two-step inference process may be a unique specialization of MEC.

III-96. Separable Manifold Geometry in Macaque Ventral Stream and DCNNs

SueYeon Chung¹
Joel Dapello*²
Uri Cohen³
James J DiCarlo⁴
Haim Sompolinsky⁵
¹Columbia University
²Harvard University

SUEYEON.CHUNG@COLUMBIA.EDU
DAPELLO@G.HARVARD.EDU
URI.CO@MAIL.HUJI.AC.IL
DICARLO@MIT.EDU
HAIM@FIZ.HUJI.AC.IL

³Hebrew University of Jerusalem

⁴Massachusetts Institute of Technology

⁵The Hebrew University of Jerusalem and Harvard University

The mixed selectivity of a population of sensory neurons to different stimuli conditions (such as distance, pose, background, etc.) gives rise to an “object manifold”, defined as the set of population response vectors to the different stimuli that does not change perceived object identity. In this work, the amount of object identity information in a population of neurons is quantified through “classification capacity” defined by the maximum number of manifolds-per neuron, that can be classified into two classes by a hyperplane. Manifold classification capacity was shown to depend on the geometry of these manifolds in the population state space, in particular their radius and dimension. Other important statistical properties of the manifolds are their correlation structure, such as correlations between their centers or axes of variation. Together, these properties characterize how disentangled object representations are geometrically and computationally. Here we apply the theory to analyze the geometry and capacity of object manifolds in different regions of the macaque ventral stream, in response to objects with a variety of orientations and backgrounds, and compare these to the object manifolds from an array of ImageNet-trained deep convolutional neural networks (DCNNs) in response to the same stimuli. Our analysis shows that separable geometry emerges in both the tested DCNNs and macaque ventral stream. While overall trends are qualitatively similar between different models, the rates and trajectories of disentangling differs. Furthermore, correlation measures between manifold centers and manifold axes indicate that the population responses in the neural data is far more correlated than those in the deep networks. Finally, we present how neural noise can be thought of as an imposed margin on the classification task, and present how the population response’s dependence on the level of noise can be reflected in the classification capacity by activating a finite size margin.

III-97. Saccades enhance perception via interactions between visual and motor signals in area V1

Seth Akers-Campbell¹

James Niemeyer²

Michael Paradiso¹

SETH_AKERS-CAMPBELL@BROWN.EDU

NIEMEYJ@GMAIL.COM

MICHAEL_PARADISO@BROWN.EDU

¹Brown University

²Weill Cornell Medical Center

Several times per second, macaques and humans make rapid saccadic eye movements that direct the fovea to objects of interest. There is increasing evidence that saccades and a corollary discharge (CD) signal (associated with the motor command to move the eyes) both play important roles in perception. However, we know little about the relationship between perception and peri-saccadic neural activity because the two are usually not studied concurrently. To investigate this relationship, we recorded neural activity in primary visual cortex (V1) while the macaques performed a contrast discrimination task with stimuli presented during a saccade or simulated saccade (no CD signal). Consistent with human psychophysical data, in macaques we find that a saccade reduces contrast sensitivity at low spatial frequencies; importantly, this suppression is accompanied by reduced single-neuron spike rates. The correlation of neural and perceptual effects suggests that a CD signal in V1 is involved in saccadic suppression, which is thought to be critical for the lack of perceived blur during saccades and the apparent stability of perception despite jumps in gaze. At higher spatial frequencies, we find that saccades enhance contrast sensitivity, with a corresponding increase in firing rates. Increases in neural activity and perceptual sensitivity may enhance object recognition, as significant form information is present at higher spatial frequencies. In addition to these frequency-specific effects, at all spatial frequencies, each new fixation is accompanied by a non-visual signal that resets and synchronizes local field potential oscillations and spiking. Using data-driven stochastic process models, we find that the modulations of firing rate can be explained by interactions between visual and motor signals within V1 neurons. Animals’ perceptual decisions can also be predicted from V1 spikes; to do this the classifier makes use of the spatio-temporal pattern of spikes across the recording array (rather than a simple rate code).

III-98. Neural correlations in a cortical circuit that generates behavioral variability

Galen Lynch¹
Joseph Scherrer¹
Harbaljit Sohal²
Natasha Denisenko¹
Ed Boyden¹
Michale Fee¹

GALEN.LYNCH@GMAIL.COM
JOEBOB@MIT.EDU
HARBI.SOHAL@GMAIL.COM
NDENISEN@MIT.EDU
ESB@MEDIA.MIT.EDU
FEE@MIT.EDU

¹Massachusetts Institute of Technology

²Paradromics Inc

The acquisition of many learned behaviors begins with a phase of highly variable exploration. This variability plays a crucial role in reinforcement learning and other models of learning. In songbirds, the premotor brain region lateral magnocellular nucleus of the anterior nidopallium (LMAN) actively generates behavioral variability and injects it into the motor system [1]. But how does LMAN generate this behavioral variability? Theoretical studies propose that interconnected neurons within LMAN exhibit chaotic dynamics, thereby producing seemingly random patterns of activity. One such model posits that LMAN is a balanced excitatory-inhibitory (EI) network, and predicts that pairs of neurons in LMAN would have uncorrelated activity [2]. Another model posits that LMAN may act as an excitable media producing locally propagating waves of activity, and predicts that all nearby pairs of neurons would be highly correlated. To test these models and to understand how LMAN actively generates behavioral variability, we built a miniature lightweight microdrive to simultaneously record from multiple neurons, as well as a lightweight endoscope to perform functional calcium imaging of ensembles of LMAN neurons. With these new technologies, we observed the simultaneous activity of pairs of single units in singing juvenile and adult birds. We find that most pairs of neurons with small separation ($<250 \mu\text{m}$) are completely uncorrelated, which is incompatible with the wave model. However, a small subset of pairs have strikingly large correlations, with correlation coefficients of up to 0.81. Intriguingly, these correlated pairs of neurons can be separated by up to $400 \mu\text{m}$. The existence of such highly correlated neurons within LMAN is inconsistent with LMAN being a balanced excitatory-inhibitory network with uniformly random connectivity. These results suggest that new models of variability generation are required to explain how LMAN generates exploratory behavioral variability

III-99. Functional stripes in V2 are a byproduct of interleaved central-peripheral visuotopy

Madineh Sedigh-Sarvestani¹
Kuo-sheng Lee²
Juliane Jaepel²
Rachel Satterfield²
Nicole Shultz²
David Fitzpatrick²

MADINEH.SARVESTANI@MPFI.ORG
KUO-SHENG.LEE@MPFI.ORG
JULIAN.JAEPEL@MPFI.ORG
RACHEL.SATTERFIELD@MPFI.ORG
NICOLE.SCHULTZ@MPFI.ORG
DAVID.FITZPATRICK@MPFI.ORG

¹Max Planck Florida Institute of Neuroscience

²Max Planck Florida Institute for Neuroscience

Seminal studies of primate secondary visual cortex (V2) reported 'functional stripes' for various visual features: Different striped regions of V2 showed sensitivity to different features such as orientation. These findings led to the theory of parallel processing streams in visual cortex. We wondered whether such specialization existed in V2 of the tree shrew, a small diurnal animal with a smooth brain, decent visual acuity, and cortical columns. We performed wide-field and 2p imaging in V1/ V2 of awake animals. We found functional stripes for orientation, disparity, and spatial-frequency in tree shrew V2, similar to primate. However, we also found that this functional specialization was highly correlated with the underlying visuotopy, unlike any previously reported organizational framework. Tree shrew V2 contains interleaved stripes that oscillate between central-peripheral visual space. These central-

peripheral stripes have the same width and pattern as the functional stripes for orientation and disparity. Furthermore, the visuotopic map is predictive of the feature maps, with regions representing central and peripheral visual space exhibiting distinct orientation and disparity properties. To make sense of this central-peripheral visuotopy, we developed a self-organizing-map model of V2. Our model suggested that the central-peripheral stripes arise to maximize smooth visual field coverage given an elongated cortical area. Therefore, elongated visual areas in any animal should produce similar central-peripheral stripes. We confirmed this prediction in the elongated area 19 of the mature ferret, where we observed similar visuotopy and functional maps. Our findings suggest a novel and simple mechanism for map organization in higher visual areas: specialized functional maps arise as a byproduct of specialized visuotopy. We believe this challenges the parallel processing streams theory in favor of a single continuous stream with varying central-peripheral specialization. Our hypothesis applies equally to primates and non-primates, and produces falsifiable predictions for future experimental and computational studies.

III-100. Dynamics and computation with anti-leaky integrate-and-fire neurons

Paul Manz
Sven Goedeke
Raoul-Martin Memmesheimer
University of Bonn

MANZ@TH.PHYSIK.UNI-BONN.DE
SGOEDEKE@UNI-BONN.DE
RM.MEMMESHEIMER@UNI-BONN.DE

Networks in the brain consist of different types of neurons. We investigate the influence of neuron diversity on the dynamics, phase space structure, and computational capabilities of spiking neural networks. We consider inhibitory networks of leaky (LIF) and “antileaky” (XIF) integrate-and-fire neurons that generalize irregularly spiking nonchaotic LIF neuron networks. Endowed with simple conductance-based synapses for XIF neurons, our networks can generate a balanced state of irregular asynchronous spiking as well. We determine the voltage probability distributions and self-consistent firing rates assuming Poisson input with finite-size spike impacts. We show that introducing a single XIF neuron in a network of LIF neurons renders the entire network dynamics unstable and chaotic as indicated by a positive largest Lyapunov exponent. To further characterize the dynamics we consider the full spectrum of Lyapunov exponents and the related covariant Lyapunov vectors (CLVs), which indicate the directions in which infinitesimal perturbations grow or shrink with a rate given by the Lyapunov exponents. We find that for each XIF neuron in a network there will be approximately one positive Lyapunov exponent. In addition we find that the stable CLVs are concentrated in the subspace of perturbations to LIF neurons whereas the unstable CLVs are concentrated in the subspace of perturbations to XIF neurons. A simple mean-field approach, which can be justified by properties of the CLVs, provides a good approximation for the Lyapunov spectrum. As an application, we propose a spike-based computing scheme where our networks serve as computational reservoirs. The different stability properties bestow them with specific computational capabilities, where mixed networks can combine those of networks with only one neuron type.

III-101. Hierarchical tensor partitioning and tiling for learning coupled multi-scale neuronal dynamics

Gal Mishne¹
 Hadas Benisty^{2,3}
 Maria Lavzin⁴
 Shahar Levy⁴
 Ronen Talmon⁴
 Ron Meir⁴
 Jackie Schiller⁴
 Ronald Coifman²

GMISHNE@UCSD.EDU
 HADAS.BENISTY@GMAIL.COM
 MARIA.LAVZIN@GMAIL.COM
 SHAHAR.LEVY86@GMAIL.COM
 RONEN@EE.TECHNION.AC.IL
 RMEIR@EE.TECHNION.AC.IL
 JACKIE@TECHNION.AC.IL
 RONALD.COIFMAN@YALE.EDU

¹University of California, San Diego

²Yale University

³Neuroscience

⁴Technion

We present a geometry learning approach for data-driven organization of neurons into subgroups of mutual functionality and of joint global and local temporal dynamics. We harness the inter-trial variability and model the data as a tensor of three domains – neurons, time and trials and apply tensor analysis to organize the data into hierarchical partition trees along each domain. The combined partition trees define a multiscale decomposition of the tensor which leads to an overcomplete representation across coarse to fine resolutions. To find an optimal compact partitioning, we develop a novel tiling method that segments the tensor into disjoint sub-volumes (tiles) having similar dynamical regimes. The identified tiling reveals meaningful co-dependencies between the observed domains. Our tiling approach is analogous to finding the best basis in classical wavelet analysis. However, our tiles define the support of a set of data-driven filters in the observed domains as opposed to predefined wavelet filters in the time-frequency domain. We apply our approach to synthetic spike trains governed by a latent Lorenz system. The obtained tiling pattern reveals a phase transition between stability and chaos thus providing an empirical bifurcation map of the dynamic regimes. We also apply our analysis to an experimental setting combining a head-fixed hand-reach motor task with chronic optical calcium imaging from M1. The extracted tiling patterns identify both the global separation of the trials based on task outcome, as well as local temporal behavioral events. Furthermore, they reveal the fundamental difference between the dynamics of the neuronal activity in layer 2/3 and layer 5 during the hand-reach motor task. The time scales of layer 5 correspond directly to the motor movement, whereas the time scales of layer 2/3 neurons are both movement and sensory related.

III-102. Dynamic resource allocation during reinforcement learning accounts for ramping and phasic dopamine activity

Minryung Song
 Sang Wan Lee
 KAIST

RYUNG20@GMAIL.COM
 SANGWAN@KAIST.AC.KR

For an animal to learn about its environment full of diverse stimuli with limited motor and cognitive resources, it should focus its resources on potentially important stimuli. However, too narrow a focus is disadvantageous for adaptation to environmental changes. Midbrain dopamine neurons are excited by potentially important stimuli, such as reward-predicting, intense or novel stimuli, and allocate resources to these stimuli by modulating how an animal approaches, exploits, explores, and attends. The current study examined the theoretical possibility that dopamine activity reflects the dynamic allocation of resources for learning. Dopamine activity may transition between two patterns: (1) phasic responses to cues and rewards, and (2) ramping activity arising as the agent approaches the reward. Phasic excitation has been explained by prediction errors generated by experimentally inserted cues. However, when and why dopamine activity transitions between the two patterns remain unknown. By parsimoniously modifying a standard temporal difference (TD) learning model to accommodate a mixed pre-

sensation of both experimental and environmental stimuli, we simulated dopamine transitions and compared them with experimental data from four different studies. The results suggested that dopamine transitions from ramping to phasic patterns as the agent focuses its resources on a small number of reward-predicting stimuli, thus leading to task dimensionality reduction. The opposite occurs when the agent re-distributes its resources to adapt to environmental changes, resulting in task dimensionality expansion. This research elucidates the role of dopamine in a broader context, providing a potential explanation for the diverse repertoire of dopamine activity that cannot be explained solely by prediction error.

III-103. Dynamic network representation of invariant external context in medial entorhinal cortex

Isabel Low¹
Alex Williams¹
Malcolm Campbell²
Tucker fisher¹
Scott Linderman¹
Lisa Giocomo¹

ILOW@STANFORD.EDU
AHWILLIA@STANFORD.EDU
MALCGCAMP@GMAIL.COM
TGFISHER@STANFORD.EDU
SCOTT.LINDERMAN@STANFORD.EDU
GIOCOMO@STANFORD.EDU

¹Stanford University

²Harvard University

Decades of experimental research has functionally dissected navigational circuits in terms of single neuron coding properties, leading to the identification and characterization of grid cells, place cells, and landmark responsive cells, amongst others. In contrast, leading theoretical frameworks for spatial navigation (e.g. grid cell attractor networks) operate at a population level, assuming recurrent connectivity and autonomous network dynamics. Here, we interrogate network-level dynamics in navigating animals, using Neuropixels silicon probes to simultaneously record from ~100-300 cells in medial entorhinal cortex (MEC) as mice explore a tightly controlled virtual linear track. While MEC is known to exhibit subtle remapping with environmental changes, we were surprised to find dramatic shifts in coding properties in an invariant environment. In many sessions and in nearly all animal subjects, we observed spontaneous and coordinated shifts in navigational coding across the full dorsal-to-ventral axis of MEC. We characterized these functional circuit reconfigurations through a series of population-level analyses. First, principal components analysis (PCA) and nonnegative matrix factorization (NMF) models did not significantly outperform a k-means clustering model of single-trial dynamics, suggesting that shifts in population coding represent discrete remapping events, rather than blends of distinct tuning curves or a deterioration of navigational coding. Second, we fit hidden Markov models with generalized linear model observations (HMM-GLMs) to precisely identify the timing of latent state transitions in a trial-agnostic manner. Using these models, we uncovered signatures of behavioral and cognitive states difficult to observe with single neuron analyses. Both frameworks revealed that discrete shifts in navigational coding often follow periods of slow running speeds, suggesting that internal state may induce spontaneous remapping of network-wide spatial representations. Overall, this work provides an early experimental characterization and statistical analysis of population-level MEC dynamics in awake, behaving animals and in addition indicates that behavioral state changes may drive population-wide remapping events.

III-104. The neural code for face memory in macaque IT and perirhinal face patches

Liang She¹
 Marcus Benna²
 Stefano Fusi²
 Doris Tsao¹

LIANGSHE@CALTECH.EDU
 MKB2162@CUMC.COLUMBIA.EDU
 SF2237@COLUMBIA.EDU
 DORTSAO@CALTECH.EDU

¹California Institute of Technology

²Columbia University

The role of face-selective regions (“face patches”) in face memory remains unclear: how is a new face remembered, and how is this memory re-activated during perception? To address this, we recorded single neurons from two face patches, face patch PR in perirhinal cortex and face patch AM in anterior IT cortex. In both PR and AM, cells were modulated by both face identity and familiarity, and exhibited stronger, prolonged responses to unfamiliar compared to familiar stimuli. Further measurement of neuronal responses to thousands of parametrized unfamiliar faces revealed that for representing unfamiliar faces, a substantial population of PR cells inherit the “axis code” previously reported in IT face patches (Chang and Tsao, Cell 2017), which is a low dimensional, linear representation. Strikingly, we observed that in both AM and PR, the preferred axis computed using responses to familiar faces was similar to that for unfamiliar faces at short latency, but diverged at long latency. Finally, in PR, we found a small population of very sparse, ‘grandmother cell’-like cells, responding only to a single or a few faces, often personally familiar. These cells might contribute to the memory storage of individual instances of faces. Together these results suggest that face memories are represented in both PR and AM, and PR cells may facilitate association between semantic and perceptual memories.

III-105. Maintenance of visual working memory by a visuomotor cortical loop

Ivan Voitov^{1,2}
 Tom Mrsic-flogel¹

I.VOITOV@UCL.AC.UK
 T.MRSIC-FLOGEL@UCL.AC.UK

¹University College London

²Sainsbury Wellcome Centre

Connections between cortical areas are almost always reciprocal. One function of cortical loops may be to maintain coherent representations between functionally distinct cortical areas, and corresponding theoretical frameworks of perception as inference have stressed the interdependent nature of feedback and feedforward representations. We therefore hypothesized that task-dependent, internally generated feedback representations will be sensitive to ongoing feedforward activity. We examined this hypothesis in a novel head-fixed visual working memory paradigm in mice, in which mice alternated performing blocks of a delayed (non)match-to-sample task and a simple Pavlovian discrimination task. Critically, the two tasks did not differ in stimuli, reward, or movement, such that we could experimentally isolate the internal representations of visual working memory in our neural recording and perturbation experiments. Using optogenetic silencing we identified the contributions to visual working memory of several cortical regions in different phases of each trial, and identified a highly distributed role of the neocortex in supporting working memory at the start of the inter-stimulus delay periods. We next investigated neural activity in higher visual area AM and premotor area M2, which are reciprocally connected, and found robust population representations of task variables, including working memory engagement and strength. As a more direct test for the instantaneous necessity of this visuomotor cortical loop for memory maintenance, we transiently silenced either area AM or M2 (‘feedforward’) while imaging incoming axons from the reciprocally connected area (‘feedback’) at the start of the delay. While the magnitude of this optogenetic perturbation on neural activity was similar between the two tasks, the population representation of working memory in the axons was selectively disrupted. This effect was specific to the memory-dependent task and persisted until the end of the delay. Our results support the role of reciprocal corticocortical interactions in maintaining distributed cognitive representations during visual working memory.

III-106. Rapid sensorimotor adaptation through cortical input control

Nicolas Meirhaeghe^{1,2}
Hansem Sohn¹
Mehrddad Jazayeri¹

NMRGHE@MIT.EDU
HANSEM@MIT.EDU
MJAZ@MIT.EDU

¹Massachusetts Institute of Technology

²Harvard-MIT Division of Health Sciences and Technology

Humans can efficiently adapt to changes in their environment. The neural basis of such adaptability, however, is not understood. In particular, the logic behind the circuit-level plasticity required to drive adaptation remains poorly characterized. Recent advances suggest that cortical areas involved in the control of sensorimotor function may be viewed as implementing a dynamical system whose activity can be adjusted in two complementary ways, either through adjustments of inputs or adjustments of recurrent dynamics (i.e., synaptic couplings). Synaptic changes offer countless degrees of freedom, but may be costly and occur on long timescales, making them possibly inadequate for rapid forms of adaptation. Inputs, on the other hand, are typically low-dimensional and might provide an efficient way to quickly drive coherent changes across cortical units to support fast behavioral adjustments. Here we tested the hypothesis that the brain employs an input-control strategy to achieve rapid sensorimotor adaptation. We built on a time reproduction task requiring monkeys to reproduce measured time intervals. In this task, responses are biased toward the mean of previously encountered intervals indicating that animals learn the underlying interval distribution. We leveraged animals' reliance on the distribution to elicit adaptation by covertly switching the range of intervals. Animals adjusted relatively quickly to this manipulation within a single behavioral session. Recordings of population activity in dorsomedial frontal cortex during adaptation revealed a concomitant change of cortical responses. Consistent with our prediction, this adjustment appeared to be mediated by an input displacing neural trajectories to a region of the state space compatible with the new behavioral demands. These results provide evidence that rapid sensorimotor adaptation may not involve changes of synaptic couplings within cortical areas; instead, inputs to cortex are adjusted giving access to suitable regions of the cortical latent dynamics.

III-107. Functional connectivity graph estimation from nonsimultaneous calcium imaging recordings

Giuseppe Vinci¹
Genevera Allen¹
Gautam Dasarathy²

GV9@RICE.EDU
GALLEN@RICE.EDU
GAUTAMD@ASU.EDU

¹Rice University

²Arizona State University

Neuronal functional connectivity is the statistical dependence structure of neurons' activities. Functional connectivity is typically inferred from data recordings in the framework of graphical models, where a neuron is represented by a node, and an edge connects two nodes if the two neurons' activities share covariability conditionally on all other sources of variability in the network. Estimating functional connectivity is an extremely relevant task that helps us understand how neurons interact with one another while they process information under different stimuli and other experimental conditions; studying the relationship between functional and anatomical connectivity is also of great interest. Functional connectivity estimation becomes a compelling statistical problem when based on calcium imaging data. Some ambitious projects involve the calcium imaging recording of the activities of 100,000 neurons in a 1mm³ portion of mouse brain. However, typically only small subsets of neurons are recorded at once to guarantee a good temporal resolution. In such settings, the joint activities of several pairs of neurons remain unobserved for which even the simplest metric of covariability, the sample covariance, is unavailable. In the Gaussian graphical model framework this situation translates into the impossibility of recovery of the functional connectivity graph of all 100,000 neurons. We call this extremely challenging problem "graph quilting" problem. We propose a novel regularized graph estimator that, in high-dimensions with high probability, can correctly re-

cover not only the functional connections of the observed pairs of neurons, but also a superset of the connections among the neurons that are never recorded jointly. The estimator can be easily embedded in more general frameworks, such as GLMs and dynamical systems. An extensive simulation study and the analysis of large scale calcium imaging data revealed that our approach can provide a functional connectivity graph recovery similar to the case of full data.

III-108. Context-dependent representation of the decision-making process in the monkey prearcuate gyrus

Mohamad Saleh Esteki¹
 Marcus Benna²
 Stefano Fusi²
 Roozbeh Kiani¹

ME1168@NYU.EDU
 MKB2162@COLUMBIA.EDU
 SF2237@COLUMBIA.EDU
 RK97@NYU.EDU

¹New York University
²Columbia University

To make successful decisions in natural environments, the brain must gather sensory information, retrieve memories of past actions as well as their consequences in contexts similar to the present, and accurately combine sensory information with expected reward to make the most beneficial choices. Past studies focused largely on the isolated effects of sensory, reward, or contextual information on decisions, leaving integration of the three underexplored. Here, we investigate the neural mechanisms for this integration by simultaneously recording from a large population of neurons in the prearcuate region of monkeys performing a variant of a direction discrimination task with random dots. In our task, correct rightward and leftward choices provided unequal rewards, and the reward magnitude of the two choices switched every few trials without any explicit cue. Monkeys adapted by choosing the high-paying option faster and more frequently. Choices and reaction times were more strongly influenced by reward for weaker (more ambiguous) motion stimuli, as expected for optimal integration of sensory and reward information. The neural population systematically represented the gradual integration of sensory information over time, while manifesting stimulus- and time-dependent effects of reward information, compatible with behavior. Further, monkeys proved capable of abstract representation of contexts. After each reward switch, they needed to experience the change of reward for only one of the choices to adjust their reaction times for both choices in subsequent trials. A subgroup of recorded neurons represented the reward context; they maintained a nearly constant firing rate in each block and changed their rate following a switch. These context-representing neurons were intermixed with a larger group that simultaneously represented the context with integration of sensory and reward information. Together, the prearcuate neural population encoded both the inference of contextual changes and the computations that govern integration of sensory, reward, and contextual information for decision-making.

III-109. The role of thalamus in maintenance of sequential cortical activity

Nader Nikbakht
 Michale Fee

NIKBAKHT@MIT.EDU
 FEE@MIT.EDU

Massachusetts Institute of Technology

Performing learned behavior requires animals to produce precisely timed motor sequences. Birdsong is an excellent natural example of a complex, learned and precisely timed motor behavior. Here we used zebra finch as a model organism to ask how thalamic input maintains and controls the dynamics of the learned cortical sequences underlying song production. In birds, the production of the male courtship song is controlled by a brainstem-thalamocortical feedback loop that exclusively controls singing behavior. A critical component of this network is the nucleus Uvaeformis (Uva), a small thalamic region that is necessary for adult singing. Despite its importance,

single unit recordings from Uva during singing have never been reported and consequently, the exact cellular and circuit-level function of this nucleus is unknown. We developed a lightweight (~1 g) microdrive for juxtacellular recordings and with it performed the very first extracellular single unit recordings in Uva during song production. The population of Uva neurons that project to the song premotor cortical region HVC (used as a proper noun), fire bursts of action potentials prior to syllable onsets, produce rhythmic and stereotyped bursts within syllables, and are suppressed prior to syllable offsets. We additionally traced the axon terminals of single Uva neurons that project to HVC, and found that they form a focal arborization within HVC such that each HVC neuron receives input from only one Uva neuron. Our physiological and anatomical results suggest that Uva maintains sequential cortical activity during song, but does not provide unambiguous timing information. Our observations are more consistent with a model in which the brainstem-thalamocortical feedback loop acts at the syllable time-scale (~100 ms) and is less consistent with a recent model in which the brainstem-thalamocortical feedback loop acts at fast time-scale (~10 ms) to generate sequences within cortex.

III-110. A common cause for multiple suboptimalities in perceptual decision-making

Diksha Gupta
Carlos D Brody
Princeton University

DIKSHAG@PRINCETON.EDU
BRODY@PRINCETON.EDU

The Drift Diffusion Model (DDM) has been immensely successful in describing choices and reaction times (RTs) during perceptual decisions. However there are several reports of deviations from this reward-rate maximizing computation (Holmes & Cohen 2013) e.g. asymmetric correct & error RTs, post-error slowing, trial history effects and stimulus-independent errors. These deviations are often studied in isolation from each other, attributed to multifarious biases or variability in DDM parameters and therefore lack a parsimonious explanation. Here, we demonstrate that all the above suboptimalities arise naturally in the DDM with the simple addition of a previously proposed Bayesian algorithm in which agents assume that observations are drawn from a nonstationary prior with Markovian dynamics (Yu & Cohen 2008). We applied this idea to rats trained on a new RT task that we developed. Consistent with the DDM, rats had higher RTs on harder trials, however they also exhibited the aforementioned suboptimalities. Their choices were strongly influenced by recent history of stimulus category (assessed with stepwise GLM). Moreover, patterns of repetition/alternation of stimulus category modulated their RTs and accuracy. This led us to hypothesize that rats have a nonstationary prior over repetitions/alternations of stimulus category. A DDM with initial point set by such a belief (Zhang et al., 2014) not only captures the effects of stimulus category history, but surprisingly, also produces lower error RTs, stimulus-independent biases, and post-error slowing. Due to belief in nonstationarity, spurious sequences of stimulus category transiently bias the agent's prior leading to fast errors, however following a break in this pattern the agent believes that the prior is likely to have been redrawn, giving rise to slower post-error responses. This model offers a concise explanation for widely reported deviations from DDM predictions and makes clear prescriptions about how to manipulate these behaviors.

III-111. Fast gating and slow amplification of relevant evidence in frontal cortex during flexible decision-making

Marino Pagan
 Vincent Tang
 Mikio Aoi
 Jonathan Pillow
 Carlos D Brody
 Princeton University

MPAGAN@PRINCETON.EDU
 VDTANG@PRINCETON.EDU
 MAOI@PRINCETON.EDU
 PILLOW@PRINCETON.EDU
 BRODY@PRINCETON.EDU

Our ability to flexibly select, based on context, the relevant information to guide our decisions is a fundamental cognitive process, yet its neural underpinnings are still largely unknown. To address this issue, we trained rats to perform a task requiring context-dependent selection and integration of sensory information (adapted from Mante et al., Nature, 2013). In our task, rats are presented with a train of randomly-timed auditory pulses, where each pulse varies in its location (right or left) and its tone pitch (high or low). In separate blocks of trials, rats are cued to report either the prevalent location of the pulses, or their prevalent pitch. Using high-throughput training, we trained 46 rats to perform this task with high accuracy. Optogenetic perturbation of the Frontal Orienting fields (FOF) impaired rats' ability to form the correct decision, consistent with a causal role of this area in the task. Electrophysiological recordings in FOF revealed a mixed encoding of multiple task variables, including task context, pulses of relevant and irrelevant evidence, and the upcoming decision. Using linear regression, we were able to express the heterogeneous responses of individual units as a sum of task-related components. This approach allowed us for the first time to quantify the contribution of single pulses of evidence to neural population dynamics. We found that selection of relevant evidence is supported by two distinct mechanisms. First, pulses of relevant evidence are immediately projected onto a "choice axis" which determines the rats' subsequent choices, whereas irrelevant evidence projects mostly onto orthogonal "null axes". Second, the projection of relevant evidence over the choice axis is slowly amplified over time, whereas irrelevant evidence is not. Our results suggest that the brain implements a combination of fast feedforward and slow recurrent computations to flexibly select relevant evidence during decision-making.

III-112. Reorganization of neuronal population codes in the ACC underlies long-term memory retrieval

Meytar Zemer
 Noa Brande-Eilat
 Liron Sheintuch
 Alon Rubin
 Ofer Yizhar
 Yaniv Ziv
 Weizmann Institute of Science

MEYTARZEMER@GMAIL.COM
 N.BRANDE.EILAT@GMAIL.COM
 LIRON.SHEINTUCH@WEIZMANN.AC.IL
 ALON.RUBIN@WEIZMANN.AC.IL
 OFER.YIZHAR@WEIZMANN.AC.IL
 YANIV.ZIV@WEIZMANN.AC.IL

The contribution of prefrontal circuits to episodic memory retrieval increases with time following learning, suggesting that in these circuits, memory representations evolve in a way that benefit retrieval. How memory representations in the prefrontal cortex evolve with time and what aspects of this evolution contribute to remote memory retrieval remain unknown. We used Ca²⁺ imaging to longitudinally record from hundreds of neurons in the anterior cingulate cortex (ACC; part of the prefrontal cortex) as mice underwent contextual fear conditioning (FC), and again when placed in the same context after two days (recent recall) or 35 days (remote recall). Contrary to predictions made by recent studies that tagged active neurons using immediate-early genes, we found no differences in the number of active neurons or their firing rates between mice in the recent and remote recall groups. However, ACC neurons that were active during remote recall exhibited higher pairwise correlation than neurons active during recent recall. The increased pairwise correlation during remote recall was specific to the FC context and was not observed during exploration of a neutral environment, suggesting the observed time-dependent changes

in population activity were specific to the learned information. finally, we found that population activity in the ACC better differentiated between the FC context and a neutral context during remote recall, and that the difference between these contextual representations correlated with the amount of time spent freezing during remote recall. Overall, our results reveal that population activity in the ACC becomes more organized and more contextually informative with time following learning, pointing to specific aspects of the neural code that could underlie the emergent time-dependent involvement of prefrontal circuits in memory retrieval.

III-113. Brain-inspired replay in artificial neural networks

Gido van de Ven
Andreas Tolias
Baylor College of Medicine

GIDOVANDEVEN@GMAIL.COM
ASTOLIAS@BCM.EDU

Current state-of-the-art deep neural networks can solve almost any task they are trained on. But when such a network is trained on a new task, the previously learned task is quickly forgotten. Importantly, this 'catastrophic forgetting' is not due to limited capacity of the network, as the same network could learn both tasks when trained in an interleaved fashion. In the real world, however, training examples are not presented interleaved but typically appear in sequences. A straight-forward solution would be to store the encountered examples from previously learned tasks and revisit them when learning new tasks. Although such 'replay' or 'rehearsal' solves catastrophic forgetting, in the deep learning community replay is typically believed not to be a scalable solution as constantly retraining on all previous problems is very inefficient and the amount of data that would have to be stored becomes unmanageable very quickly. Yet, in the brain – which clearly has implemented an efficient and scalable algorithm for continual learning – the replay of previous experiences is important for stabilizing new memories. Inspired by this, here we revisit the use of replay as a tool for continual learning with artificial neural networks. We find that: (1) fully replaying previously learned problems is not needed, as a handful of replayed examples could be enough; (2) a perfect memory (i.e., storing all encountered examples) is not required, as a low capacity generative model could suffice; and (3) brain-inspired modifications enable generative replay to scale to complicated problems with many tasks (≥ 100) or complex inputs (natural images), resulting in state-of-the-art performance on challenging continual learning benchmarks. Moreover, when incrementally learning new classes (as opposed to new tasks), we find that replay might actually be necessary. This last result suggests a specific, so far unappreciated, computational goal for replay in the brain.

III-114. Stimulus encoding in the lateral intraparietal parietal cortex during categorization depends on training history

Kenneth Latimer^{1,2}
David Freedman¹

LATIMERK@UCHICAGO.EDU
DFREEDMAN@UCHICAGO.EDU

¹The University of Chicago
²Department of Neurobiology

Different tasks within a single environment can require performing distinct computations on the same stimuli. For example, discrimination between two stimuli involves comparing specific stimulus features, while categorization implies grouping together sets of stimuli. Here, we compare neural activity during a visual categorization task between pairs of monkeys with different prior experience to reveal how learning affects cortical representations which support performance of the categorization task. The first pair of monkeys were trained exclusively on a delayed-match-to-category (DMC) task in which the animals were required to categorize directions of moving dot stimuli. The second pair of animals had instead been trained extensively on a fine direction discrimination task, using a delayed-match-to-sample (DMS) paradigm, prior to learning the DMC task. Importantly, the discrimination and categorization task used the same structure and stimuli. Previous work has found correlates of visual motion

categorization in the lateral intraparietal (LIP) cortex during the DMC task and stimulus encoding during the DMS task. We find that LIP activity consistently exhibits low-dimensional encoding of category information during the DMC task, but the dynamics of the encoding depend on the monkeys' training. LIP activity in monkeys trained on the DMS task strongly reflects motion, and specific motion directions and category identity are represented in orthogonal dimensions in the neural population. Additionally, within-category direction encoding is weaker and more transient in monkeys trained only on categorization. Long-term training history can therefore impact the dynamics and representations of task-relevant sensory signals in LIP.

III-115. A computational model for border ownership

Miguel Lazaro-Gredilla
Rajeev Rikhye
Dileep George
Wolfgang Lehrach
Vicarious AI

MIGUEL@VICARIOUS.COM
RAJEEV@VICARIOUS.COM
DILEEP@VICARIOUS.COM
WOLFGANG@VICARIOUS.COM

figure-ground segregation of unknown objects is a challenging perception problem that our brains solve with ease. In particular, solving this problem requires integrating the global context to drive local decisions. Border-ownership cells observed in the visual cortex signal the foreground side of contours and are believed to be important for figure-ground segregation. However, there are no coherent computational theories on how the visual system might learn these cells, or how they are utilized during the inference process. Here, we propose a biologically inspired structured probabilistic network called Cloned Markov Random field (CMRF) that learns border ownership as an emergent effect of learning contour-surface-background models of images. An observation from biology about clonally related neurons – neurons that are genetically wired to share the same bottom-up receptive field – was a crucial inductive bias for CMRF to learn the right representations. Inference, performed using a biologically plausible local message passing algorithm, solves many challenging figure-ground segmentation problems and its emergent temporal dynamics of border-ownership signaling matches physiological observations. We believe the CMRF is a first step towards a model of objectness compatible with neurobiological evidence

III-116. High-capacity pattern labeling and robust retrieval with internally modular attractor networks

Ila fieta¹
Rishidev Chaudhuri²

FIETE@MIT.EDU
RCHAUDHURI@UCDAVIS.EDU

¹Massachusetts Institute of Technology

²University of California, Davis

Robust stable states of neural population activity ('attractors') are thought to underlie noise-tolerant representation and memory. Traditional neural attractor models focus on networks with a small number of attractors whose structure is entirely determined by the inputs to be stored. Recently, several lines of work have identified attractor networks that contain a huge number of internally structured attractor patterns (up to exponential in network size). The attractors are well-separated, making interference low and downstream learning easy, and highly noise-robust, with network dynamics correcting a large number of errors. These features make them appealing candidates for cognitive representations. Here we show that such high-pattern number (HPN) attractor networks can perform a number of important tasks, in a single architecture whose structure is reminiscent of the hippocampus. We use HPN networks as a module in a high-capacity pattern labeling network. The pattern labeling network can rapidly learn to map a large number of inputs onto the attractors of a much-smaller HPN network. The attractors serve as abstract labels for the input states. The network can then retrieve the attractors from noisy or partial input, maintain them without further input and despite ongoing noise, and use them to drive

downstream states and actions. The pattern-labeling network can perform a variety of computations including nearest-neighbor search, robust classification, locality-sensitive hashing and robust action-selection in response to a noisy example from an input category. The network can also be used to rapidly allocate conjunctive memory representations and to allow those memories to later become independent of the HPN network, similar to theories of the hippocampal formation. finally, the network can rapidly index large numbers of inputs and then make judgments about familiarity or novelty during testing, compatible with a postulated role for medial temporal cortex in the prodigious recognition memory found in human psychophysics.

III-117. AutoLFADS: A large-scale neural network training framework for generalized estimation of single-trial population dynamics

Mohammad Reza Keshtkaran^{1,2}
 Raghav Tandon^{1,2}
 Diya Basrai³
 Sarah L Nguyen²
 Raees H Chowdhury⁴
 Hansem Sohn⁵
 Mehrdad Jazayeri⁵
 Lee E Miller⁴
 Chethan Pandarinath^{1,2}

MKESHTK@EMORY.EDU
 RAGHAV.TANDON@EMORY.EDU
 DBASRAI@UCSD.EDU
 SNGUYEN.GATECH@GMAIL.COM
 RAEED.CHOWDHURY@GMAIL.COM
 HANSEM@MIT.EDU
 MJAZ@MIT.EDU
 LM@NORTHWESTERN.EDU
 CHETHAN@GATECH.EDU

¹Emory University

²Georgia Tech

³University of California, San Diego

⁴Northwestern University

⁵Massachusetts Institute of Technology

Large-scale recordings of neural activity are becoming ubiquitous, providing new opportunities to study network-level dynamics in diverse brain areas and during increasingly complex, natural behaviors. However, the sheer volume of data and its dynamical complexity are critical barriers to uncovering and interpreting these dynamics. Deep learning methods are a particularly promising approach due to their ability to uncover meaningful relationships from large, complex, and noisy datasets. One such method, latent factor analysis via dynamical systems (LFADS[1]), uses recurrent neural networks to infer latent dynamics from high-D neural spiking data. When applied to motor cortical (M1) activity during stereotyped behaviors, LFADS substantially improved the ability to uncover dynamics and their relation to subjects' behaviors on a moment-by-moment, millisecond timescale. However, applying LFADS to less-structured behaviors, or in brain areas that are not predominantly driven by intrinsic dynamics, is far more challenging. This is because LFADS, like many deep learning methods, requires careful hand-tuning of complex model hyperparameters (HPs) to achieve good performance. Here we demonstrate AutoLFADS, a large-scale, automated model tuning framework that can characterize dynamics in diverse brain areas without regard to behavior. AutoLFADS uses distributed computing to train dozens of models simultaneously while using evolutionary algorithms to optimally tune HPs in a completely unsupervised way. AutoLFADS required 10-fold less data to uncover dynamics from macaque M1, with better generalization to unseen behavioral conditions than previous LFADS models. We then tested data from the somatosensory and dorsomedial frontal cortices, areas with very different dynamics from M1. AutoLFADS produced precise estimates of population dynamics without any prior knowledge of the areas' dynamics, tasks, or subjects' behaviors, outperforming any individually-trained LFADS model obtained through random HP searches. finally, we present a cloud software package and comprehensive tutorials that enable new users to apply the method without needing dedicated computing resources.

III-118. Understanding the functional and structural differences across excitatory and inhibitory neurons

Sun Minni¹
 Li Ji-An¹
 Theodore Moskovitz¹
 Grace Lindsay²
 Kenneth Miller¹
 Mario Dipoppa¹
 Guangyu Robert Yang¹

SUNMINNI1031@GMAIL.COM
 JIAN.LI.ACAD@GMAIL.COM
 TED@GATSBY.UCL.AC.UK
 GRACEWLINDSAY@GMAIL.COM
 KENDMILLER@GMAIL.COM
 MARIO.DIOPPA@GMAIL.COM
 GYYANG.NEURO@GMAIL.COM

¹Columbia University

²Gatsby Computational Neuroscience Unit, University College London

Deep neural networks have become powerful tools to model brains. They have been successfully used to model various aspects of the sensory, cognitive, and motor systems. Despite many achievements of applying deep networks to model brains, standard architectures deviate from the brain in fundamental ways, one of them is the separation of excitatory (E) and inhibitory (I) neurons. E and I neurons differ in many important aspects, such as postsynaptic effects, population size, projection range, stimulus selectivity, and connection density. It is unclear whether these differences across E and I neurons serve computational purposes, and it is unknown whether they are mutually independent. To answer these questions, we first trained deep recurrent convolutional neural networks on object classification tasks with several built-in structural principles of the brain, including Dale's law, the abundance of E neurons, and the role of E units as principal neurons. We found that differences in stimulus selectivity and connection density across E and I neurons naturally emerged through training. Not hardwired, these emerged properties could only have evolved under the pressure of performing the tasks. We further explored the necessary conditions for the networks to develop distinct selectivity and connectivity across cell types. We found that the number of E, but not I, channels have a strong influence on both image selectivity and connection density of E neurons; neurons projecting to higher-order areas will have greater stimulus selectivity, regardless of whether they are excitatory or not; sparser connectivity is required for higher selectivity, but only when the recurrent connections are excitatory. These findings demonstrate that the differences observed across E and I neurons are not independent, and can be explained using a smaller number of factors.

III-119. Antagonism normalizes input to the olfactory system in freely breathing animals

Joseph Zak
 Venkatesh Murthy
 Harvard University

JZAK@FAS.HARVARD.EDU
 VNMURTY@FAS.HARVARD.EDU

An odor landscape is a complex blend of discrete molecules that each activate unique, yet highly overlapping, populations of olfactory sensory neurons (OSNs). Despite a rich diversity of ORN subtypes (~1000 in mouse and ~400 in humans), the overlapping nature of odor inputs may saturate neural responses at the early stages of stimulus encoding. To prevent stimulus information loss, normalizing mechanisms are necessary. One mechanism, independent of neural processing, may operate through antagonism at the level of receptor-ligand interactions. We first measured population responses from OSNs by imaging odor mixture responses their axon terminals in olfactory bulb glomeruli. Our results were in close agreement with our theoretical work and showed that mixture responses in OSNs, at the population level, are non-linear and favor antagonism. We then developed a new method that allows for direct optical access to OSN cell bodies within the olfactory epithelium, in freely breathing mice. We used this new approach to investigate odor mixture interactions in individual sensory cells. Our results indicate that antagonism is a remarkably common feature of odor mixture encoding in olfactory sensory neurons at both the population and single-cell levels, and that both the prevalence and magnitude of this interactions increase with odor mixture complexity. We show for the first time, within the constraints of naturalistic

respiration, odor mixtures suppress the activity of OSNs. The interplay between our theoretical and experimental work suggests that antagonism plays a fundamental role in how odor mixtures are encoded at the earliest stages of olfaction.

III-120. Spatio-temporal dynamics of cortical cholinergic signaling

Sweyta Lohani
Andrew Moberly
Mike Higley
Jessica Cardin
Yale University

SWEYTA.LOHANI@YALE.EDU
ANDREW.MOBERLY@YALE.EDU
MIKE.M.HIGLEY@YALE.EDU
JESS.CARDIN@YALE.EDU

During wakefulness, animals transition rapidly between discrete behavioral states, such as arousal and locomotion. These state transitions are associated with alterations in ongoing cortical activity to influence perception and cognition. Neuromodulators, such as acetylcholine, may serve as a key link between behavioral state and cortical function. Ascending cholinergic projections were long thought to homogeneously impact cortical activity. However, more recent evidence for spatial diversity in cholinergic axonal arborizations suggests the potential for heterogeneous modulation across the cortex. Importantly, the spatial and temporal dynamics of cholinergic signaling across cortical networks during behavioral state transitions are unknown.

To address this question, we used a recently developed genetic tool— fluorescent G-protein coupled receptor based sensors (GRABs)¹—that enables measurement of acetylcholine levels with excellent spatial and temporal precision. Using a recently reported approach for widespread viral expression following neonatal sinus injection², we simultaneously monitored green fluorescent GRABs signaling and red fluorescent neuronal activity via jRCaMP1b across the entire cortex using widefield, mesoscopic imaging. These experiments were performed in head-fixed mice freely running on a wheel, transitioning between bouts of arousal, quiescence and locomotion.

Our results demonstrate that cholinergic signaling is remarkably heterogeneous across cortical regions, with arousal, as measured by pupil diameter, and locomotion, as measured by wheel motion, associated with distinct spatial and temporal cortical patterns of cholinergic activation. In particular, arousal (induced by airpuff) in the absence of locomotion homogeneously increases cholinergic levels in the cortex, whereas locomotion onset predominantly increases activity in the frontal cortex. Furthermore, locomotion onset is associated with decreased correlation of the cholinergic signal across diverse cortical areas. These findings are also associated with distinct patterns of activity in cortical excitatory neurons. This state-of-the-art dual imaging approach allows us to quantitatively dissect the relationship between heterogeneous cholinergic signals across the cortex and local cortical dynamics.

III-121. Ongoing correction of temporally perturbed spike sequences by fluctuating scalar modulation

Rich Pang
Adrienne Fairhall
University of Washington

RPANG.CONTACT@GMAIL.COM
FAIRHALL@U.WASHINGTON.EDU

In birdsong, a standard model for complex learned behaviors, precise neural activity sequences in premotor area HVC are thought to drive stereotyped song production. While the mechanisms behind precision HVC timing are unknown, lesion and electrophysiology studies suggest that upstream inputs from the thalamic Uvaeform nucleus (Uva) are critical, potentially activating individual sequence-generating networks in HVC that represent different song syllables. Neural recordings, however, show that Uva activity is strongly spatially correlated and exhibits little syllable-specific structure, challenging its ability to selectively activate distinct HVC subnetworks at different times.

Using an excitatory-inhibitory chain-structured spiking network model, we considered an alternative role for Uva as a global modulatory input that actively corrects timing errors produced by a full sequence-generating network in HVC. We found that if a spike pulse's position along the chain has been erroneously shifted by several chain links relative to a hypothetical target pulse, a time-varying but scalar-valued modulatory input to the inhibitory population can reestablish its correct position. Crucially, inhibition must be nonuniformly spread over the chain so as to reflect the modulation timecourse spatially, yielding effective spatiotemporal "correction zones" that attract modestly temporally perturbed sequences back toward stable ones. This mechanism could coordinate left and right HVC during song, which share no direct connections, and sheds light on a potential generic method to suppress timing errors in neural activity.

III-122. The caudate nucleus controls coordinated, reward-dependent adjustments to perceptual decisions

Takahiro Doi
Yunshu Fan
Joshua Gold
Long Ding
University of Pennsylvania

DOI.TAKAH@GMAIL.COM
YUNSHUF@PENNMEDICINE.UPENN.EDU
JGOLD@PENNMEDICINE.UPENN.EDU
LDING@PENNMEDICINE.UPENN.EDU

Our decisions often need to balance what we observe and what we desire. However, our understanding of how and where in the brain such decisions are made remains limited. The basal ganglia (BG) pathway is known to make separate contributions to: 1) perceptual decisions that require the interpretation of uncertain sensory evidence, and 2) value-based decisions that select among outcome options. This pathway is thus a prime candidate for mediating decisions that integrate sensory observations and desired rewards. We examined the roles of the caudate nucleus, a key input station in the BG pathway, in this type of decision using a combination of modeling, behavior, and electrophysiology (recording and microstimulation) in monkeys. On a random-dot visual motion discrimination task with asymmetric choice-reward associations, monkeys used reward context-dependent biases that allowed them to obtain near-maximal rewards (Fan, et al., eLife, 2018). Consistent with causal involvements of the caudate nucleus in these reward-biased decisions, we found that, during motion viewing, 1) many caudate neurons (44 out of 142) showed joint modulation by sensory evidence (motion strength corresponding to at least one of the two choices) and reward (either the lateralized reward context or reward size) at the single-neuron level; and 2) electrical microstimulation in the caudate nucleus evoked reward context-dependent effects in perceptual sensitivity and bias (n=55 sessions). Within a drift-diffusion framework, the monkeys' voluntary reward bias strategy involved coordinated asymmetries in the drift rate and relative bound heights. Strikingly, caudate microstimulation induced similarly coordinated, reward context-dependent changes. These coordinated microstimulation effects further depended on how neural activity at the stimulation sites were modulated by visual and reward-related factors. These results suggest that the caudate nucleus plays a causal role in mediating a decision process that balances external evidence and internal preferences to guide adaptive behavior.

III-123. Inhibitory inputs onto premotor neurons regulate vocal timing in an avian model of vocal turn-taking

Jonathan Benichov*¹
Philipp Norton*²
Susanne Schreiber³
Daniela Vallentin¹

JONATHAN.BENICHOV@GMAIL.COM
PHILIPP.NORTON@HU-BERLIN.DE
S.SCHREIBER@HU-BERLIN.DE
DANIELA.VALLENTIN@ORN.MPG.DE

¹Max Planck Institute for Ornithology

²Humboldt University of Berlin and Bernstein Center Berlin

³Humboldt-Universität zu Berlin

During conversations we rapidly switch between listening and speaking, which often requires withholding or delaying our speech in order to hear others and avoid overlapping. This capacity for vocal turn-taking is exhibited by non-linguistic species as well, however the neural circuit mechanisms that enable us to regulate the precise timing of our vocalizations during interactions are unknown. Here we examine how zebra finches control their vocal response timing in an interactive setting. Birds presented with call playbacks at a regular rate will typically respond with stereotyped latencies of ~260ms, but will adjust their call timing to avoid overlapping with an additional social partner. Using an intracellular microdrive, we record from premotor neurons and interneurons within the cortical vocal-motor nucleus HVC (proper name) and find premotor bursting associated with call production, preceded by increased interneuron activity. Based on this experimental data, we developed a leaky integrate-and-fire model of HVC neurons, by first simulating the observed firing patterns of interneurons and premotor neurons. With those neurons, as well as simulated auditory-evoked excitation counterbalancing inhibition, as inputs onto a model premotor neuron, we were able to reproduce the experimentally observed premotor bursts that are time-locked to call production. In this model, reducing inhibitory weights translates to earlier and stronger premotor bursts, predicting a reduction in call response latency when inhibition within HVC is decreased. Indeed, pharmacological blocking of GABAergic transmission results in faster and stronger premotor drive and faster response times during vocal interactions. The model further predicts a sustained auditory-evoked excitatory drive onto HVC that gets delayed by local inhibition to produce behaviorally relevant response latencies. In the proposed model for vocal turn-taking, appropriately timed vocal responses to auditory signals can be generated through a descending forebrain pathway in which local inhibitory interneurons regulate the timing and intensity of premotor drive.

III-124. Neural circuit strategy for categorical selection in the midbrain

Nagaraj Mahajan
Shreesh Mysore
Johns Hopkins University

NAGARAJ.MAHAJAN@JHU.EDU
SHREESH.MYSORE@JHU.EDU

The ability to select the most important stimulus among multiple competing stimuli is critical for survival and adaptive behavior. Such selection must be robust to neural response variability and sensory noise. How does the brain implement selection with such robustness? One brain area that plays an important role in stimulus selection across space is the midbrain selection network in vertebrates. It comprises of the superior colliculus (SC in mammals; optic tectum (OT) in birds), and two isthmic nuclei located near the SC/OT in the lateral tegmentum. In this work, we show that the isthmi pars magnocellularis (Imc), a GABAergic nucleus in the network, exhibits a specially structured 'donut-like' pattern of net functional inhibitory projections onto the other components of the network. Using a computational model of the network, along with extracellular recordings and focal iontophoretic inactivation of Imc neurons, we demonstrate that the donut-like pattern of connectivity results in increased noise-robustness and in a categorical representation of competing stimuli. Additionally, we find that the donut-like pattern of connectivity has multiple holes, needed to implement Imc's combinatorial encoding properties that were reported recently. Furthermore, using focal inactivation of neurons in the isthmi pars parvocellularis (Ipc), a cholinergic amplifier in the network, to mimic the effect of introducing self-inhibition, we show that disrupting the donut-like pattern of net functional connectivity results in decreased noise robustness. These results demonstrate that the donut-like pattern of connectivity is required for noise-robust, categorical stimulus selection. Finally, we also compare our model and experimental results with the divisive normalization model² which involves uniform inhibition, discuss the differences between the two, and highlight important advantages of the donut-like motif for categorical decision making.

Author Index

- Aazhang B., 132
Abbasi-Asl R., 173
Abbaspourazad H., 120
Abbott L., 35, 150, 161, 167, 176, 195, 203
Abe E., 108
Abe T., 231
Abrams J., 77
Adachi J., 44
Adam Y., 142
Adesnik H., 118, 119
Adoff M., 197
Adolphs R., 94
Agarwal G., 48
Agnes E. J., 58, 75
Agrochao M., 40
Aguillon V., 130
Ahmadabadi M. N., 136
Ahmadian Y., 108, 172
Ahrens M., 59, 68
Akbari H., 102
Akers-Campbell S., 248
Akil A., 156, 217
Akitake B., 179
Akkad A., 194
Albrecht D., 80
Aldarondo D., 117, 203
Allen G., 254
Allen K., 210
Ambrose E., 140
Amvrosiadis T., 230
Andalman A., 45
Andalman A. S., 199
Andrei A., 156
Annayini E., 161
Ansuini A., 161
Aoi M., 257
Apodaca-Montano D., 141
Arkhipov A., 184
Arlt C., 57, 207
Aschner A., 82, 188
Ashhad S., 89
Ashwood Z., 130
Asokan M., 111
Atilgan H., 38
Attinger A., 247
Avansino D., 193, 233
Axel R., 48, 69, 161, 203
Ba D., 128
Baccus S., 228
Bade M., 104
Baek S., 157
Bahroun Y., 237
Bai Y., 108, 208
Bair W., 62
Baker C., 39, 41
Balasubramanian H., 115
Balasubramanian V., 103
Ballesta S., 186
Banerjee D., 60
Banyai M., 165
Bao P., 47
Barahona M., 109
Barak O., 149
Barak T., 78
Barendregt N., 187
Bari B., 129, 153
Barnett A., 65
Baron J., 86
Barrett-Jolley R., 93
Barri A., 79
Barroso-Luque R., 57, 207
Barry M. L. L. R., 223
Basrai D., 260
Basu R., 55
Bathellier B., 245
Batista A., 36, 226, 246
Batra K., 77
Batty E., 97, 231
Baumann M., 147
Baumbach A., 159
Bauza M., 94
Beck A., 118
Becker M., 43
Becker S., 114
Behnam R., 154
Behnia R., 176, 195
Beiran M., 83, 127, 219
Ben-Shushan N., 162
Bengio Y., 74, 156, 184
Benichov J., 263
Benisty H., 251
Benna M., 253, 255

- Bennett J., 127
Bensmaia S., 217
Berens P., 74
Berman G., 203
Bernacchia A., 39
Bernaez Timon L., 57
Berners-Lee A., 140
Besserve M., 150
Bethge M., 68
Bhalla U., 138
Bharioko A., 194
Bhatia A., 138
Bhattacharyya S., 163
Biane J. S., 141
Bicknell B., 191
Biderman D., 231
Bigelow A., 62
Bill J., 130
Billaudelle S., 159
Billeh Y., 184
Bindas S., 106
Bishop W., 68
Biswas T., 110
Bittner K., 96
Bittner S., 176
Blancas Reyes E., 97
Boccaro C., 96
Bodor A., 135
Boehm U., 142
Bolkan S., 201
Bollu T., 197
Bondanelli G., 245
Bondy A., 84, 142
Bonnachi N., 48
Bonnen T., 190
Boorman E., 183
Bordelon B., 123
Borisjuk A., 237
Born G., 200
Bossaerts P., 53
Botvinick M., 29
Bouchard K., 163
Boundy-Singer Z., 126
Boutin V., 155
Boyden E., 249
Bradbury S., 199
Brain Laboratory I., 130
Brainard M. S., 233
Brande-Eilat N., 257
Braun W., 188
Bredenberg C., 120
Breitwieser O., 159
Brennan C., 177
Brincat S., 63
Brinkman B., 145
Brody C. D., 84, 128, 142, 176, 211, 256, 257
Bromberg-Martin E., 61
Brown J., 239
Brown M., 172
Brown S., 70
Brunel N., 169, 178
Bruno R., 67, 192
Brunton B., 69
Buchanan J., 135
Buice M., 84, 166
Bukhari Q., 160
Bumbarger D., 135
Buonocore A., 147
Burak Y., 162
Burge J., 186
Burgess N., 62
Buschman T., 82, 198
Busse L., 200
Buzsaki G., 47, 90, 93, 230

Cain N., 113
Calaim N., 234
Calame D., 43
Calhoun A., 45
Calipari E., 175
Campbell M., 247, 252
Can T., 126
Canto Bustos M., 226
Carandini M., 50
Carcea I., 106
Cardin J., 92, 262
Carey M., 31
Carlson D., 97
Carter E., 45
Celik N., 93
Chalk M., 109
Chandra J., 220
Chandrasekaran C., 67
Chang Y., 114
Chapochnikov N., 131
Charles A., 121
Charlton J., 108
Chase S., 36, 68, 246
Chau H. Y., 85
Chaudhuri R., 181, 259
Chaudhuri-Vayalambrone* P., 94
Chavane F., 155
Che D., 82
Chen C., 106
Chen G., 34
Chen R., 199
Chen W., 168
Chen Y., 160
Chen Z., 143
Cheng H., 239

- Cherian C., 139
Chettih S., 207
Chichilnisky E., 97
Chin M., 117
Ching S., 80, 87
Chintaluri C., 215
Chklovskii D., 131, 216, 237
Choi H., 113
Chong E., 49
Choudhary M., 120
Chow C., 89
Chowdhury R. H., 260
Christenson M., 195
Chua R., 184
Chubykin A., 141
Chung S., 65, 247
Churchland A., 130, 162, 231
Churchland M., 35, 236
Ciocchi S., 116
Cisek P., 184
Clark D., 40, 232
Climer J., 197
Clopath C., 60, 109
Coen P., 50
Coen-Cagli R., 82, 98, 188
Coenen F., 93
Cohen A., 142
Cohen J., 129, 153
Cohen M., 112, 163
Cohen U., 247
Coifman R., 251
Cone J., 104
Confavreux B., 58
Constantinople C., 111
Conzelmann K., 194
Costa N., 135
Costa R. P., 116
Cowley B., 45, 215
Cox D., 192, 212
Crombie D., 200
Csillag M., 159
Cueva C., 117
Cunningham J., 30, 35, 176, 231, 236
- D'Andola M., 161
Dąbrowska P. A., 179
da Silveira R. A., 194
Dabagia M., 122
Dahmen D., 179, 180, 191
Dan C., 240
Daniel C., 159
Dapello* J., 247
Daptardar S., 91
Darshan R., 121
Dasarathy G., 254
- Dasgupta S., 207
Datta S. R., 154, 160, 231
David S., 99
Davidson J., 177
Davis M., 209
Davoudi H., 197
Daw N. D., 66
Dayan A., 41
de Vries S., 84, 166
Degenhart A., 226
Degleris A., 231
Deisseroth K., 45, 199
Delis I., 174
Deneux T., 245
Denisenko N., 249
Denman D., 165
DePasquale B., 84
Derdikman D., 105
Dethe N., 97
Deutz L., 179
Devinsky O., 218
DiCarlo J., 229
DiCarlo J. J., 58, 242, 247
Dickinson E., 69
Diesmann M., 179
DiGregorio D., 79
Ding L., 263
Dipoppa M., 52, 261
Dobs K., 112
Doeller C., 134
Doi E., 88
Doi T., 186, 263
Doiron B., 43, 52, 143, 182, 220, 226, 245
Dolen G., 32
Domanski A., 149
Dombeck D., 197
Dorkenwald S., 135
Doroc K., 53
Dorrell W., 72
Dougherty M., 163
Dovrolis C., 70
Doyle W., 218
Dragoi V., 156
Druckmann S., 34, 233
Drugowitsch J., 130, 158
Duan C., 176
Dubreuil A., 149, 152
Dudman J., 88
Duncker L., 216
Dunn T., 203
Dupuy N., 230
Durango A., 55
Durstewitz D., 149
Dyer E., 122

- Ebitz B., 72, 106
 Echavarria C., 192
 Economo M., 133
 Egger S., 140
 Einevoll G., 200
 Einstein M. C., 51
 Ekelmans P., 103
 Engel T., 100, 122, 162, 241
 Engelken R., 167
 English D., 47
 Erisken S., 200
 Ernoult M., 156
 Escola S., 199, 202
 Essig J., 37
 Esteki M. S., 255
- Fabre G., 41
 Fahey P., 135
 Fairhall A., 153, 262
 Fan J. L., 52
 Fan L., 142
 Fan Y., 263
 Fang C., 92
 Feather J., 55, 81
 Fee M., 249, 255
 Feldheim D., 243
 Feldman J., 89
 Felsen G., 37
 Feng T., 140
 Feng Y., 61, 86
 Ferguson A., 219
 Ferguson K., 92
 Ferrentino B., 214
 Festa D., 82
 field G., 79, 146
 Fiete I., 113
 fiete I., 77, 116, 171, 259
 filipowicz A., 225
 fink A., 48, 69
 finkelstein A., 133
 finkle E. A., 114
 fisher T., 252
 fitzgerald J., 110
 fitzpatrick D., 249
 flavell S., 41
 flinker A., 218
 Fontolan L., 133
 Forseth K., 132
 Foster D., 140, 244
 Frady P., 221
 Franciosini A., 155
 Francl A., 81, 218
 Franco P., 53
 Franconville R., 240
 Fraser K., 129
- Freedman D., 34, 104, 107, 258
 Freedman J., 203
 Freeman N. W., 211
 Friedrich R., 33
 Froemke R., 106
 Froudarakis E., 135
 Fry A., 44
 Fukai T., 61
 Fumarola F., 119
 Fusi S., 67, 94, 192, 253, 255
- Gabriel K., 220
 Gandhi S., 60
 Ganguli S., 64, 213, 228, 243, 247
 Gao M., 141
 Garrett M., 113, 138, 184
 Garvert M., 134
 Gauthier J. L., 211
 Geerts J., 62
 Geisler W. S., 77
 Genkin M., 100
 George D., 259
 Gershman S., 32
 Gershman S. J., 130
 Gerstner W., 223
 Getz M., 245
 Geva N., 76
 Ghazanfar A., 221
 Ghazizadeh E., 87
 Ghosh S., 189
 Giaffar H., 159
 Gieselmann M., 241
 Gigante S., 121
 Gill J., 198
 Gillett M., 169
 Gillis W., 231
 Gillon C., 74
 Gilmer J., 174
 Gilra A., 76
 Giocomo L., 33, 148, 247, 252
 Gire D., 137
 Gjorgjieva J., 38, 56, 107, 151, 202
 Glaser J., 35
 Goedeke S., 76, 188, 250
 Goetz G., 97
 Goffinet J., 54
 Gokcen E., 41
 Golan T., 223
 Gold J., 225, 263
 Goldbach H., 179
 Goldberg J., 197
 Golkar S., 237
 Golshani P., 51, 180
 Goltz J., 159
 Golub M., 36, 246

- Gonzalez R., 55
Goodman J., 217
Gopal V., 137
Goris R., 108, 126, 208
Greeman E., 192
Greenidge C. D., 224
Gremel C., 101
Greydanus S., 132
Grienberger C., 96
Grigsby E., 226
Grimaud J., 72
Grissom N., 106
Grobowski P., 138
Grollier J., 156
Grosmark A., 209
Gross J., 105
Grossman C., 153
Grover P., 189
Grun S., 179
Gu K., 180
Gu Q., 53
Gu-Lamberg K., 132
Guitchounts G., 212
Guo D., 90
Guo S., 199
Gupta D., 142, 256
- Haber A., 71
Haberker H., 240
Hady A. E., 177
Hafed Z., 147, 181
Haga T., 61
Hajnal M., 51
Hancock K., 111
Handy G., 237
Harris K., 50, 65
Harvey C., 33, 57, 124, 207
Harvey S., 64
Hatsopoulos N. G., 217
Hausmann S., 77
Hausser M., 191
Hayden B., 72
Heath S., 195
Heinze S., 209
Helias M., 64, 179, 180, 191
Heller C., 99
Hemberger M., 151
Henaff O., 126, 208
Henderson J., 193, 233
Hennequin G., 39
Hennig J., 36, 246
Henning S., 146, 183
Hermundstad A., 46, 108, 240
Herrero J., 102
Hertaeg L., 183
- Hickok G., 246
Higley M., 262
Hillier D., 194
Hillman E., 203
Hiratani N., 205, 227
Hirokawa J., 169
Histed M. H., 179
Hochberg L., 193, 233
Hocker D., 111
Hong W., 180
Hosseini E., 194
Huang C., 52, 182, 245
Huang L., 166
Huang S., 180
Huang Z., 74
Hulse B., 240
Hunt J., 37
Huszar R., 47
- Ih D., 135
Ireland E., 133
Israely S., 124
Istrate N., 100
Ito B., 197
Ito H., 55
Ito S., 243
- Jaepel J., 249
Janak P., 129
Janz R., 156
Jayakumar S., 134
Jayaraman V., 46, 240
Jazayeri M., 29, 127, 185, 194, 213, 219, 222, 236, 254, 260
Ji Ma W., 32
Ji N., 41, 52
Ji-An L., 261
Jia X., 42
Johnston W. J., 107
Jones M., 149
Joshua M., 162
Josic K., 156, 187, 217, 220, 241
Jun J., 65
Jun N. Y., 146
Jung M., 164
- Kable J., 225
Kadmon J., 243
Kaleb K., 60
Kalle Kossio Y. F., 76
Kang B., 34
Kang L., 110
Kang Y., 211
Kanwisher N., 112
Kao J., 67, 224

- Kapoor V., 150
Kar K., 229
Karashchuk P., 69
Kardon B., 197
Karpas E., 98
Karpova A., 154
Kass R., 63
Kato I., 138
Kaufman M., 148, 217
Kawashima T., 59, 68
Keijser J., 146
Kell A., 112
Kepecs A., 169, 206
Kepple D., 235
Kerekes P., 94
Kersen D., 103
Keshishian M., 102
Keshtkaran M. R., 260
Keup C., 180
Keyes L., 77
Khalighinejad B., 102
Khanna S., 43
Khatib D., 105
Kheirbek M., 141
Khona M., 113
Kiani R., 98, 120, 255
Kilpatrick Z., 177, 187, 241
Kim C., 89
Kim J., 91
Kim K., 47
Kim R., 208
Kim T., 62, 129
Kingsbury L., 180
Kinsella I., 97, 135, 142
Kira S., 124
Kirchner J. H., 38, 202
Kivikas S., 138
Klein A., 200
Klein N., 63
Kleinman M., 67, 224
Klos C., 76
Knep E., 106
Koay S. A., 128, 211
Koh T. H., 68
Kohn A., 41, 82
Kohn J., 176
Koulakov A., 206, 235
Kramer L., 163
Krause E., 158
Kraynyukova N., 103, 118
Kremkow J., 115
Kreutzer E., 205
Kriegeskorte N., 223
Kriener L., 159
Krishnamurthy K., 126, 225
Krishnan S., 139
Krishnaswamy S., 121
Krstulovic M., 94
Krupic J., 94
Kudryashova N., 230
Kuhn T., 64, 180
Kuleshova E., 154
Kungl A. F., 159
Kutlu M., 175
Kvitsiani D., 206
Kwan A., 38
Kwon M., 91
L'Heureux Q., 138
Lagzi F., 226
Lainscsek C., 51
Lajoie G., 168, 184
Lak A., 169
Lam N., 53
Landau I. D., 239
Lao C. L., 200
Lappalainen J., 125
Larsen R., 113, 138
Latham P., 147, 205, 227, 246
Latimer K., 258
Laurent G., 151
Lavzin M., 251
Layer M., 179
Lazar A. A., 54
Lazaro-Gredilla M., 259
Lecoq J., 74, 81, 84, 138
Ledochowitsch P., 84, 166
Lee A., 171
Lee D. D., 200
Lee H., 242
Lee J., 97
Lee J. H., 114
Lee J. S., 171
Lee K., 135, 249
Lee S. W., 251
Lee T. S., 85
Leedy C., 179
Lehrach W., 259
Leibold C., 190
Lengyel M., 211
Lerman G., 198
Letai K., 82
Levenstein D., 90
Levina A., 109, 122, 241
Levine A., 89
Levy S., 251
Lewis S., 137
Li J., 196
Li K., 97
Li N., 34, 133

- Li Q., 172
Li W., 203
Li X., 133
Li Y., 96
Lillicrap T. P., 74
Lin T. W., 160
Linderman S., 88, 160, 231, 252
Linderman S. W., 154
Lindsay G., 78, 261
Lindsey J., 167
Lipshutz D., 237
Lisberger S., 140
Litke A., 243
Litwin-Kumar A., 161, 167, 206
Liu M., 197
Liu T., 54
Livezey J., 163
Loewenstein Y., 78
Logothetis N., 109, 150
Lohani S., 262
Lohmann C., 38
Long L., 218
Lopez Espejo M., 99
Losey D., 36
Losonczy A., 209
Louie K., 171, 185
Lovett-Barron M., 199
Low I., 252
Lu R., 135
Lubin E. E., 114
Luo T., 84, 142
Luppi A. I., 75
Lustig A., 154
Lynch G., 249
- MacDowell C., 82
Mace E., 194
Machens C., 41, 234
Macrina T., 135
Madan G., 41
Maes A., 109
Magee J., 96
Magland J., 65
Mahajan N., 264
Mahalingam G., 135
Maheswaranathan N., 132, 228
Mainen Z., 48, 227
Maldonado P., 38
Mallik S., 80
Mamassian P., 98
Mamelak A., 94
Mamidanna P., 68
Manakov M., 154
Manneschi L., 127
Manz P., 250
- Maoz O., 98
Marchesi P., 174
Margalit E., 242
Marino P., 226
Markowitz J. E., 231
Marques T., 58, 242
Marshall J., 117, 203
Marshall N., 35
Martinez J., 112
Masis J., 192, 212
Masse N., 104
Masset P., 169
Mastrogiuseppe F., 149, 152, 227
Mathis A., 68, 134
Mathis M., 68
Maunsell J., 104, 189
Mazzucato L., 100, 175, 227
McClain N., 226
McDermott J., 55, 81, 144, 218
McGill M., 47, 125
McGrory C., 85
McIntosh L., 228
McKenzie S., 47
McKhann G., 218
McMahan B., 224
McNaughton B., 60
McWalter R., 144
Meding K., 126
Mehaffey W. H., 233
Mehrpour V., 104
Mehta A., 102
Meier K., 159
Meir R., 251
Meirhaeghe N., 194, 254
Mel G., 213
Melugin P., 175
Memmesheimer R., 76, 188, 250
Menendez J., 246
Merel J., 117
Mesgarani N., 102, 218
Metz L., 132
Meyer T., 104
Meyniel f., 90
Michaelis C., 68
Miehl C., 107
Mikell C., 44
Miller D., 183
Miller E., 63
Miller K., 52, 118, 119, 261
Miller L. E., 260
Millman D., 84
Milstein A., 96
Minni S., 261
Minxha J., 94
Mishne G., 121, 251

- Mishra N., 203
Mistic B., 168
Mitelut C., 65, 97
Mittal D., 145
Mizes K., 44
Mlynarski W., 108, 244
Mobarhan M., 200
Moberly A., 262
Mofakham S., 44
Mohan K., 34
Mohinta S., 116
Mokashe S., 178
Monosov I., 61
Monsalve Mercado M. M., 190
Mooney R., 54
Moore T., 241
Moosavi S., 238
Moroni M., 49
Morris G., 105
Moskovitz T., 261
Mossing D., 118, 119
Motiwala A., 226
Moza S., 138
Mrsic-flogel T., 78, 253
Murakami M., 227
Murawski C., 53
Murphy C., 38
Murray J., 199, 202
Murray J. D., 53, 185
Murthy M., 45, 133, 196, 212
Murthy V., 72, 134, 261
Musall S., 231
Muscinelli S., 206
Mysore S., 264
- Nagy D., 165
Najafi F., 138
Nakayama H., 204
Narayan S., 59
Narayanan R., 145
Nashef A., 124
Nasiri S., 216
Nassar J., 44, 65, 234
Naveh G., 99
Navratilova Z., 60
Nayebi A., 228
Nejatbakhsh A., 85
Neupane S., 185
Newman-Boulle A., 154
Nguyen K., 241
Nguyen S. L., 260
Nieh E. H., 211
Niell C., 108
Niemeyer J., 248
Nienborg H., 31
- Nikbakht N., 255
Niu X., 85
Nogueira R., 67, 192
Nold A., 114
Noorman M., 46, 240
Norman-Haignere S., 218
Norton P., 263
Nowotny T., 127
Nwabudike I., 41
- O'Leary T., 94
O'Shea D., 36, 216
O'Brien F., 93
O'Connor D., 114
Obeid D., 40
Oby E., 36, 226, 246
Ocker G., 84, 166
Ocko S., 247
Ocko S. A., 64, 213
Ogasawara T., 61
Okorokova E. V., 217
Olcese U., 174
Oliver M., 81, 84, 166
Ollershaw D., 138
Olsen S., 42, 138
Olveczky B., 44, 117, 203
Onken A., 230
Orban G., 51, 165
Orellana J., 63
Orlova N., 81, 138
Ostojic S., 127, 149, 152, 219, 245
Oswald A., 226
Ott T., 169
Ottenheimer D., 129
oude Lohuis M., 174
- Pacheco Pinedo D., 133, 212
Padilla-Coreano N., 77
Padoa-Schioppa C., 186
Pagan M., 257
Page E., 104
Paik S., 157
Pakman A., 97
Pallasdies F., 188
Palmigiano A., 118, 119
Pan D., 52
Panagiotaropoulos F., 150
Pandarinath C., 260
Pang R., 262
Panichello M., 198
Paninski L., 59, 97, 135, 142, 203, 231
Panzeri S., 49, 124
Paradiso M., 248
Parekh R., 240
Parga N., 83

- Park I. M., 65, 234
Park S., 183
Park Y., 157
Parker P., 108
Parks T. V., 219
Parthasarathy N., 229
Pasarkar A., 135
Patel K., 203
Pearson J., 54, 146
Pedrosa V., 60
Pehlevan C., 40, 72, 123, 131, 152, 172
Pennartz C., 174
Pereira T., 196
Pereira U., 227
Perich M., 45
Perkins J., 74
Perkins S., 35
Perrinet L., 155
Person A., 43, 174
Pesaran B., 120, 170, 214
Petersen P., 230
Petrovici M. A., 159, 205
Pettine W., 185
Peyrache A., 164
Pfeiffer B., 140
Piasini E., 225
Pica G., 124
Piccato A., 213
Piet A., 176, 184
Pillow J., 45, 88, 130, 173, 201, 215, 224, 257
Pina J., 74
Pinto L., 211
Piray P., 66
Pisokas I., 209
Pitkow X., 91
Plaza S., 240
Plitt M., 148
Pnevmatikakis E., 212
Podlaski W., 234
Pogodin R., 147
Pojoga S., 156
Polack P., 51
Polley D., 111
Ponvert N., 184
Poo C., 48
Popovkina D., 62
Portes J., 176
Pospisil D., 62
Pouget A., 130
Prakhya S., 125
Proekt A., 177
Proskurin M., 154
Prut Y., 124
Puelma Touzel M., 184
Purdon P., 118
Pyle R., 95
Qiao S., 214
Qin S., 152
Querlioz D., 156
Quick K., 36
Rager D., 43
Rajalingham R., 213
Rajan K., 45
Raju P. C., 223
Ramadan M., 236
Raman B., 80
Ramirez-Villegas J. F., 150
Ramos Traslósheros Lopez L. G., 73
Rangel L. M., 86
Rashid Shomali S., 136
Rast L., 158
Rasuli S. N., 136
Ravindranath S., 196
Recanatesi S., 191, 227
Redd J., 197
Reid C., 84, 135, 165
Reimer J., 135
Reinagel P., 49
Rhee J., 192
Ricci M., 164
Richard J., 129
Richards B., 74
Richter L., 56
Rikhye R., 259
Rinberg D., 49, 198, 204
Rinzel J., 90
Riquelme J. L., 151
Ritt J., 70
Riviere P., 86
Rivkind A., 121
Roach J., 162
Rocheftort N., 230
Rodgers C., 67, 192
Roemschied F., 133
Roman R., 83
Romani S., 46, 96, 133, 240
Romo R., 83
Rosenbaum R., 39, 95, 156, 217
Roska B., 194
Rozsa B., 194
Rubin A., 76, 257
Rubin G., 240
Ruda K., 79
Rule M., 94
Rullan C., 173
Rungratsameetaweemana N., 51
Rupp K., 69
Russo A., 236

- Rust N., 104
Rutishauser U., 94
Rutten V., 39
Rysakova M., 154
Ryu S., 36
- S Cash S., 51
Sabatini D., 148
Sabrin K., 70
Sachdeva P., 163
Saddler M., 81
Saderi D., 99
Sadtler P., 36
Sadun L., 116
Safavi S., 109, 150
Sahani M., 78, 216
Salazar-Gatzimas E., 40
Samuel A., 86
Sandbrink K., 68
Sani O., 170, 214
Sarno S., 83
Satterfield R., 249
Savin C., 111, 120
Sawtell N., 196
Saxena S., 231, 236
Scellier B., 156
Schaffer E., 203
Scheffer L., 240
Schemmel J., 159
Scherrer J., 249
Schevon C., 218
Schiller J., 251
Schneider F., 200
Schneidman E., 71, 98
Scholl B., 224
Schoonover C., 48, 69
Schottdorf M., 211
Schrater P., 91
Schreiber S., 263
Schreiner D., 101
Schrimpf M., 58
Schuck N., 134
Schuessler F., 149
Schwab D., 126
Sedigh-Sarvestani M., 249
Sedillo I., 214
Seid S., 74, 138
Sejnowski T., 51, 160, 208
Semedo J., 41
Seminara A., 137
Sengupta A., 216, 237
Senn W. M., 205
Serre T., 164
Shaevitz J., 196
Shaham N., 162, 220
- Shanechi M., 120, 170, 214
She L., 47, 253
Shea-Brown E., 102, 191
Sheffer B., 160
Sheffield M., 139
Sheintuch L., 76, 257
Shenoy K., 36, 193, 216, 233
Shi W., 186
Shi Y., 122, 241
Shimazaki H., 136, 238
Shlizerman E., 91
Shoham S., 49, 198
Shokri H., 97
Shultz N., 249
Shusterman R., 172
Shuvaev S., 206, 235
Si G., 86
Si Y., 243
Sibille J., 115
Siegler J., 42, 81
Silies M., 73
Silva D., 140
Simoncelli E., 104, 229
Sinha N., 124
Sirota A., 168
Slepukhin V., 89
Smith M., 43, 123, 182, 215
Snyder A., 182
Snyder L., 87
Sohal H., 249
Sohl-Dickstein J., 132
Sohn H., 222, 254, 260
Sokol P. A., 65
Sokoloski S., 188
Soltesz I., 96
Sommer F., 221
Sompolinsky H., 99, 200, 220, 239, 247
Song A., 128
Song M., 251
Soni A., 164
Sorscher B., 200, 213
Soulat H., 118
Spacek M., 200
Spalla D., 96
Sparks F., 209
Stachenfeld K., 62
Stan P., 215
Starosta S., 206
Stavisky S., 233
Stefanini F., 141
Stein N Brito C., 223
Steinmetz N., 241
Stephen E., 118
Stern M., 100, 102, 137
Stetler B., 198

- Stock C., 64
Stocker A., 235
Stone I., 201
Suarez E., 168
Sullivan D., 138
Sun C., 210
Sun X., 36
Sun Y., 161
Suresh A. K., 217
Sutlief E., 129
Svoboda K., 133
Szathmary E., 159
- Tabris N., 196
Tafazoli S., 82
Taillefumier T., 116, 171
Takahashi D., 221
Takemura S., 240
Talmon R., 251
Tanaka H., 228
Tanaka R., 40, 232
Tandon N., 132
Tandon R., 260
Tang V., 257
Tank D., 128, 211
Tariq M., 137
Tarrat Varella T., 221
Tassa Y., 117
Tatischeff M., 238
Tavoni G., 103
Tchumatchenko T., 57, 103, 114, 118
Temereanca S., 128
Tervo D., 154
Tesileanu T., 216
Theurel D., 143
Thiberge S., 128, 212
Thiele A., 241
Thompson C., 74
Tian L., 233
Tian X., 181
Tikidji-Hamburyan A., 97
Timcheck J., 243
Tingley D., 93
Tkacik G., 244
Tocker G., 105
Tolias A., 135, 258
Tolooshams B., 128
Tonegawa S., 210
Topalidou M., 246
Torres R., 135
Toyoizumi T., 110
Tran D. T., 51
Tran N., 235
Trautmann E., 36
Trenholm S., 194
- Treves A., 96
Tsao D., 47, 253
Tschopp F., 125
Tseng S., 207
Tsybouski D., 138
Tu C., 72
Turaga S., 125
Turner-Evans D., 240
Tuthill J., 69
Tye K., 77
Tyler-Kabara E., 36
Tyulmankov D., 150
- Umakantha A., 123
Urai A., 130
- Vacher J., 98
Valente A., 152
Valenza R., 74
Vallentin D., 263
van de Ven G., 258
van den Heuvel M., 70
van Meegen A., 64
Varol E., 85
Vasilaki E., 127
Venkatesh P., 189
Vergara J., 83
Viejo G., 164
Vinci G., 254
Vogels T., 58, 75, 215
Voges N., 179
Voitov I., 253
Voletti V., 203
von Papen M., 179
- Wagle M., 199
Wagner A., 190
Wakeman W., 138
Wallach A., 196
Walshe R. C., 77
Wang H., 38, 153
Wang J., 180, 194
Wang P., 117, 161
Wang W., 203
Wang X., 185
Wang Y., 97, 231
Waters J., 166
Wayne G., 117
Webb B., 209
Wei X., 117
Wei Y., 70
Wei Z., 59
Weinreb C., 154
Weisenhaus M., 195
Weiss O., 82

- Wells M., 50
Wertz A., 194
White H., 80
White J. K., 61
Whitehead S., 197
Whiteway M., 203
Whiteway M. R., 231
Widloski J., 244
Wiechert M., 79
Wilbrecht L., 30
Willett F., 193
Williams A., 231, 252
Williamson R., 111, 123
Wilson C., 49
Wilson G., 233
Wilson M., 143
Witten D., 102
Witten I., 201
Wolff S., 212
Wolpert D., 211
Wong Y., 120
Woods N., 141
Wosniack M., 38
Wu E., 180
Wu J., 135
Wu S., 97, 182, 220
Wu X., 140
Wyrick D., 175
- Xia N., 57
Xie M., 142
Xu L., 137
Xu Q., 113
Xue C., 163
- Yadav N., 53
Yamins D., 190, 242
Yang G. R., 150, 161, 261
Yang S., 130
Yang X., 143
Yang Y., 52, 214
Yates J., 88, 224
Yatsenko D., 135
Yeh C., 54
Yellapantula S., 132
Yim M. Y., 116, 171
Yizhar O., 257
Yonehara K., 194
Yoon E., 47
Yu A., 90
Yu B., 36, 41, 68, 123, 173, 182, 226, 246
Yu P., 52
- Zabouri N., 38
Zachry J., 175
- Zak J., 261
Zandvakili A., 41
Zatka-Haas P., 50
Zemer M., 257
Zeng H., 166
Zenke F., 58
Zeraati R., 122, 241
Zhang K., 61
Zhang L., 235
Zhang T., 181
Zhang W., 85, 220
Zhang Y., 164, 221
Zhao H., 198
Zhao Y., 234
Zheng Y., 93
Zhong W., 200
Zhou D., 135, 142
Zhou P., 135
Zhu B., 238
Zhu O., 34
Zhu V., 39
Ziemba C., 126, 208
Ziv Y., 76, 257
Zlatic M., 30
Zoccolan D., 161
Zoltowski D., 88
Zyllberberg J., 74, 79