COSYNE

CONFERENCE

Denver CO, *Mar 1–4*

2018

cosyne.org

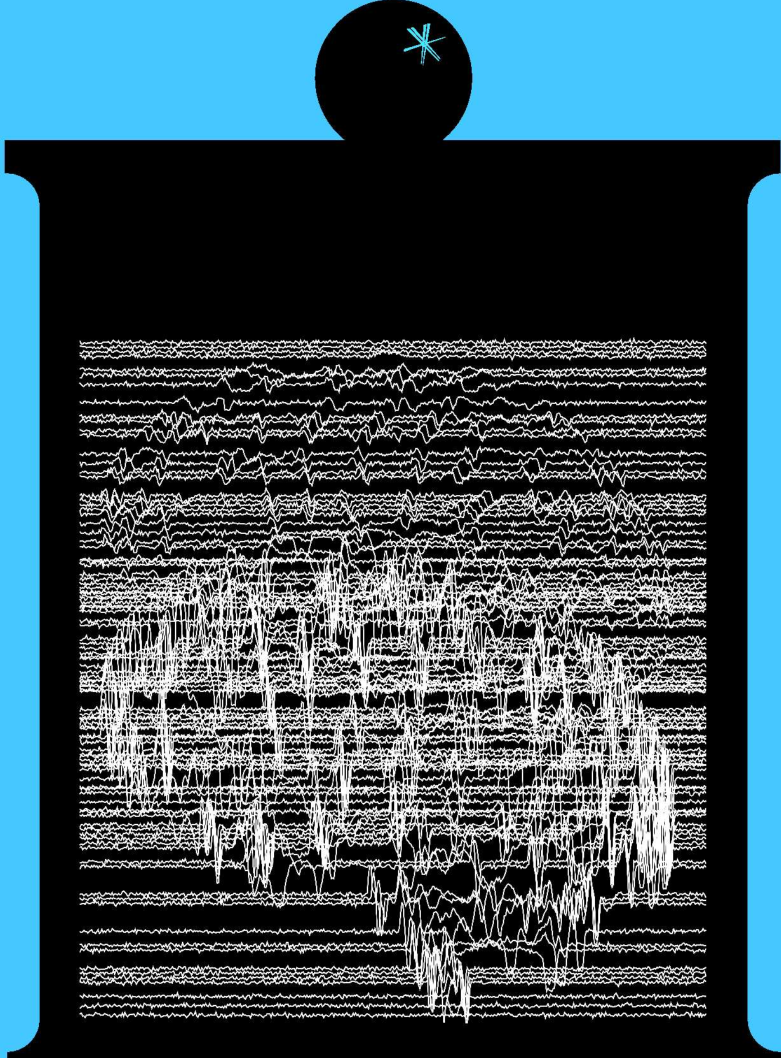

GATSBY

Google

SIMONS FOUNDATION

BURROUGHS WELLCOME FUND

Program Summary

Thursday, 01 March

4.00p	**Registration opens**
5.00p	**Welcome reception**
5.45p	**Opening remarks**
6.00p	**Session 1: Opening session** Invited speakers: Tiago Branco, Iain D. Couzin
8.00p	**Poster Session I** *Refreshments are sponsored by the Simons Foundation*

Friday, 02 March

7.30a	**Breakfast**
8.30a	**Session 2: Listen and learn** Invited speaker: Claudia Clopath; 3 accepted talks
10.30a	**Session 3: Reinforcement learning: listen again** Invited speaker: Timothy Behrens; 2 accepted talks
11.45p	**Lunch break**
2.00p	**Session 4: Basal ganglia and dopamine** Invited speaker: Joshua Berke; 4 accepted talks
4.15p	**Session 5: Functional aspects of neural variability** Invited speaker: Marlene Cohen; 2 accepted talks
5.30p	**Dinner break**
8.00p	**Poster Session II** *Refreshments are sponsored by the Simons Foundation*

Saturday, 03 March

7.30a	**Breakfast**
8.30a	**Session 6: Controlling the brain** Invited speaker: Byron Yu; 3 accepted talks
10.30a	**Session 7: The doors of perception** Invited speaker: Jessica Cardin; 2 accepted talks
11.45p	**Lunch break**
2.00p	**Session 8: Dynamical neural codes** Invited speaker: Mate Lengyel; 4 accepted talks
4.15p	**Session 9: Complexity in simple networks** Invited speaker: Carina Curto; 2 accepted talks
5.30p	**Dinner break**
8.00p	**Poster Session III** *Refreshments are sponsored by the Simons Foundation*

Sunday, 04 March

7.30a	**Breakfast**
8.30a	**Session 10: Decision making and navigation** Invited speaker: Ann M. Graybiel; 3 accepted talks
10.30a	**Session 11: New vistas in neural integration** Invited speaker: Vivek Jayaraman; 2 accepted talks
11.45p	**Lunch break**
2.00p	**Session 12: Pathways to decisions** Invited speaker: Joni Wallis; 2 accepted talks


PLEXON®
Your provider for Neuroscience technology needs
Featured Products
OmniPlex® D
Neural data acquisition
N-Form® Array
U.S. Patent No. 9,095,267 and patents pending
Customizable 3D neural probe
Plexon Probes
U-Probe
V-Probe
S-Probe
The original, multi-use, multi-site electrode
Minimizes trauma to the brain tissue
Newest and strongest of Plexon's Probes
Electrophysiology
Behavioral Tracking
Optogenetics

grapevine
neural interface system

Ripple's Grapevine Neural Interface System is a hardware and software platform for diverse neuroscience research applications. It is compact, portable, and optimized for real-time, high channel count, closed-loop electrophysiology studies with up to 512 recording and stimulating electrodes

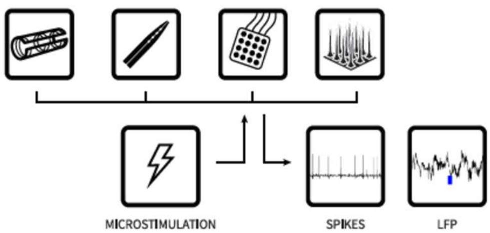

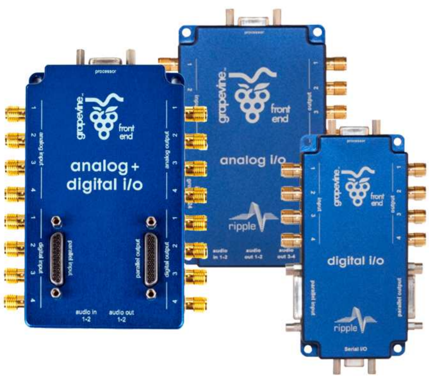

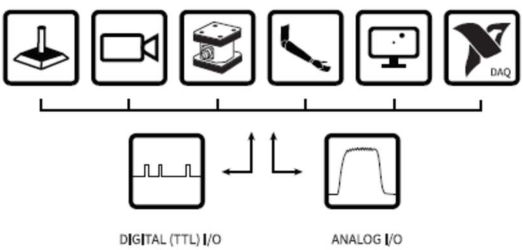

ripple

sales@rppl.com | (801) 413-0139 | rippleneuro.com


be global. be one.
oneworld® is proud to be your official airline alliance.
The **one**world network covers more than 1,000 destinations in 150 plus countries, so we can get you anywhere you need to be.
Learn more at oneworld.com
oneworld


COSYNE 2019
Cosyne goes to Europe, save the date!
MAIN MEETING Lisbon PT, Feb 28–Mar 3
WORKSHOPS Cascais PT, Mar 4–Mar 5

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 and the occasional odd bribe.

Cosyne topics include (but are not limited to): 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 2018 Leadership

Organizing Committee

General Chairs

Ilana Witten (Princeton University) and Eric Shea-Brown (University of Washington)

Program Chairs

Linda Wilbrecht (University of California, Berkeley) and Brent Doiron (University of Pittsburgh)

Workshop Chairs

Laura Busse (Ludwig-Maximilians-Universitaet Munich) and Ralf Haefner (University of Rochester)

Undergraduate Travel Chairs

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

Communications Chair

Xaq Pitkow (Rice University)

Publicity 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

Linda Wilbrecht (University of California, Berkeley), co-chair
Brent Doiron (University of Pittsburgh), co-chair
Genevera Allen (Rice University)
Bruno Averbeck (NIH)
Bing Brunton (University of Washington)
Steve Chase (Carnegie Mellon University)
Damon Clark (Yale University)
Long Ding (University of Pennsylvania)
Julijana Gjorgjieva (Max Planck Institute for Brain Research)
Benjamin Grewe (ETH Zurich)
Timothy Hanks (University of California, Davis)
Catherine Hartley (New York University)
Christopher Harvey (Harvard University)
Andrea Hasenstaub (University of California, San Francisco)
Alex Huk (University of Texas at Austin)
Santiago Jaramillo (University of Oregon)
Mehrdad Jazayeri (Massachusetts Institute of Technology)
Matthieu Louis (Centre for Genomic regulation)
Arianna Maffei (Stony Brook University)
Cris Niell (University of Oregon)
Bence Olveczky (Harvard University)
Srdjan Ostojic (Ecole Normale Superieure)
Robert Rosenbaum (University of Notre Dame)
Marshall Schuler (Johns Hopkins University)
Samuel Sober (Emory University)
Saori Tanaka (ATR Brain Information Communication Research Laboratory)
Tatjana Tchumatchenko (Max Planck Institute for Brain Research)
Srini Turaga (Janelia Farm Research Campus)
Melissa Warden (Cornell University)
Joel Zylberberg (University of Colorado)

Cosyne 2018 reviewers

Avishek Adhikari, Arash Afraz, Athena Akrami, Badr Albanna, Johnatan Aljadeff, Genevera Allen, Elizabeth Ann Amadei, Ana Amador, Bruno Averbeck, Dominik Bach, Omri Barak, Andrea Barreiro, André Bastos, Brice Bathellier, Aaron Batista, Jeff Beck, Gordon Berman, Michael Beyeler, James Bigelow, William Bishop, Aaron Bornstein, Hannah Bos, Braden Brinkman, Nicholas Cain, Ioana Carcea, Alex Cayco-Gajic, Matty Chalk, Rishidev Chaudhuri, Jerry Chen, Shizhe Chen, Eugenia Chiappe, Hannah Choi, Amelia Christensen, SueYeon Chung, Ruben Coen-Cagli, Christine Constantinople, Aurelio Cortese, Vincent Costa, Frederic Crevecoeur, Maria Dadarlat, Stephen David, Anita Disney, Jochen Ditterich, Takahiro Doi, Chunyu Duan, Eva Dyer, Hafsteinn Einarsson, Ahmed El Hady, Mark Eldridge, Ian Ellwood, Tatiana Engel, Jeffrey Erlich, Katie Ferguson, Christopher Fetsch, James Fitzgerald, Alfredo Fontanini, Birte Forstmann, Vikram Gadagkar, Andra Geana, Marc Gershow, David Gire, Matt Golub, Michael Graupner, Benjamin Grewe, Gabrielle Gutierrez, Julie Haas, Bilal Haider, Timothy Hanks, Catherine Hartley, Till Hartmamn, Christopher Harvey, Andrea Hasenstaub, Matthias Hennig, Mark Histed, Haruo Hosoya, Chengcheng Huang, Alex Huk, Kyohito Iigaya, Santiago Jaramillo, Mehrdad Jazayeri, James Jeanne, Vijay Mohan K Namboodiri, Leor Katz, Hokto Kazama, Ann Kennedy, Anisha Keshavan, Zachary Kilpatrick, Shin Kira, Takaki Komiyama, Charles Kopec, Kishore Kuchibhotla, James Kunert-Graf, Alex Kwan, Josef Ladenbauer, Guillaume Lajoie, Phillip Larimer, Kenneth Latimer, Mark Laubach, Arthur Leblois, Andrew Leifer, Talia Lerner, Anna Levina, Scott Linderman, Giuseppe Lisi, Ashok Litwin-Kumar, Matthieu Louis, Gyorgy Lur, Cheng Ly, Pedro Maia, Jeremy Manning, Jonathan Mapelli, David Margolis, Jesse Marshall, Christian David Márton, Mackenzie Mathis, Matthew McGinley, Joseph McGuire, Florian Meier, Ida Momennejad, Ilya Monosov, Ruben Moreno Bote, Masayoshi Murakami, John D. Murray, Jens-Oliver Muthmann, Kathy Nagel, Kae Nakamura, Devika Narain, Manjari Narayan, Thomas Naselaris, Matt Nassar, Richard Naud, Thierry Nieus, Gabriel Ocker, Amy Orsborn, Srdjan Ostojic, Gonzalo Otazu, Marius Pachitariu, Marino Pagan, Chethan Pandarinath, Il Memming Park, Abigail Person, Jean-Pascal Pfister, Eugenio Piasini, Xaq Pitkow, Alon Poleg-Polsky, Braden Purcell, Garvesh Raskutti, Heather Read, Tobias Reichenbach, Blake Richardss, Mattia Rigotti, Sandro Romani, Robert Rosenbaum, Peter Rudebeck, Douglas Ruff, Caroline Runyan, A Saleem, Yulia Sandamirskaya, Cristina Savin, Evan Schaffer, Geoffrey Schoenbaum, Benjamin Scholl, Nicolas Schuck, Ben Scott, Melissa Sharpe, Marion Silies, Julie Simpson, Annabelle Singer, Matt Smear, Spencer Smith, Adam Snyder, Samuel Sober, Angelika Steger, Carlos Stein, Merav Stern, Sina Tafazoli, Gaia Tavoni, Jordan Taylor, Tatjana Tchumatchenko, Dougal Tervo, John Tuthill, Takanori Uka, Floris van Breugel, Alex Vaughan, Giuseppe Vinci, Kate Wassum, Corette Wierenga, Robert Wilson, Sho Yagishita, Noriaki Yahata, Mehmet Yanik, Jacob Yates, Byron Yu, Zhengwu Zhang, Petr Znamenskiy, and Joel Zylberberg.

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

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 2018 #cosyne18

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. Four award granting programs were available for Cosyne 2018.

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
- Google

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 2018 recipients are:

Akram Bakkour, Jeremie Barral, Flora Bouchacourt, Julia Chartove, Hannah Choi, Christopher Cueva, Joshua Downer, Chunyu Duan, Lea Duncker, Dalin Guo, Caroline Haimerl, Jorge Jaramillo, Sharika K M, Anna Kutschireiter, Emily Mackevicius, Marcelo Mattar, Bayar Menzat, David Rolnick, Sofia Soares, Jake Stroud, Eszter Vertes, and Anqi Wu

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 2018 recipients are:

Denise Cai, Olga Dal Monte, Jens-Bastian Eppler, Zahara Girones, Akash Guru, Uday Jagadisan, Miaomiao Jin, Matthew Kearney, Edward Kim, Samuel Kissinger, Yunzhe Liu, Mariana Marquez Machorro, Nicolas Meirhaeghe, Mirna Mihovilovic Skanata, Bartul Mimica, Eric Mooshagian, Emily Oby, Raviv Pryluk, Katrina Quinn, Solymar Rolon-Martinez, Hana Ros, Naveen Sendhilnathan, Tristan Shuman, Yosef Singer, Daniel Takahashi, Jiannis Taxidis, Yayoi Teramoto, Sanghyun Yi, Peter Zatka-Haas, and Maarten Zwart.

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 2018 Cosyne Mentors and mentees are:

Julijana Gjorgjieva and Marina Elaine Wosniack, Michael Hausser and Francisco Sacadura, Josh McDermott and Erica Shook, Xaq Pitkow and Elizabeth Borneman, Daniel Yamins and Kevin Feigelis, and Byron Yu and Hillary Wehry.

Cosyne Undergraduate Travel Grant Program

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

The 2018 recipients are:

Noor Adra, Jonathan Aldana-Mendoza, Kathleen Esfahany, Diogo Miguel Goncalves Fortes, Ionatan Kuperwajs, Rachel Langan, Nadine Ly, Kamal Maher, Tuan Pham, Pedro Ribeiro, Jiaqi Shang, Sophie Shang, Yuzhou (Evelyn) Tong, Quinn Tran, Jacob Zavatone-Veth, and Rebecca Zhang

Program

Note: Printed copies of this document do not contain the abstracts; they can be downloaded at:
http://cosyne.org/c/index.php?title=Cosyne2018_Program
Institutions listed in the program are the primary affiliation of the first author. For the complete list, please consult the abstracts.

Thursday, 01 March

12.00n **Cosyne tutorial session sponsored by the Simons Foundation**

4.00p **Registration opens**

5.00p **Welcome reception**

5.45p **Opening remarks**

Session 1: Opening session

(Chair: Linda Wilbrecht, Brent Doiron)

8.00p **Poster Session I** *Refreshments are sponsored by the Simons Foundation*

Friday, 02 March

7.30a **Continental breakfast**

Session 2: Listen and learn

(Chair: Srdjan Ostojic)

10.00a **Coffee break**

Session 3: Reinforcement learning: listen again

(Chair: Timothy Verstynen)

Session 4: Basal ganglia and dopamine

(Chair: Santiago Jaramillo)

Session 5: Functional aspects of neural variability

(Chair: Mehrdad Jazayeri)

Saturday, 03 March

7.30a **Continental breakfast**

Session 6: Controlling the brain

(Chair: Alex Huk)

10.00a **Coffee break**

Session 7: The doors of perception

(Chair: Andrea Hasenstaub)

11.45p **Lunch break**

12.00p **Lunch—Workshop for equality and diversity in science**

Session 8: Dynamical neural codes

(Chair: Joel Zylberberg)

3.45p **Coffee break**

Session 9: Complexity in simple networks

(Chair: Robert Rosenbaum)

04.15p Emergent dynamics from network connectivity: A minimal model
Carina Curto, Pennsylvania State University **(invited)** 32

05.00p Assessing the scalability of biologically-motivated deep learning algorithms and architectures
S. Bartunov, B. Richards, A. Santoro, G. Hinton, T. Lillicrap, DeepMind 46

05.15p Propagation of spike timing and firing rate in feedforward networks reconstituted in vitro
J. Barral, X. Wang, A. Reyes, New York University 46

5.30p **Dinner break**

5.30p **Simons Foundation Theory+Experiment Match Making**

8.00p **Poster Session III** *Refreshments are sponsored by the Simons Foundation*

Sunday, 04 March

7.30a **Continental breakfast**

Session 10: Decision making and navigation

(Chair: Timothy Hanks)

08.30a The striatum and decision-making based on value
Ann M. Graybiel, Massachusetts Institute of Technology **(invited)** 32

09.15a Orbitofrontal and parietal contributions to economic decisions in rats
C. Constantinople, C. Kopec, C. Brody, Princeton University 47

09.30a Widespread cortical involvement in evidence-based navigation
L. Pinto, D. Tank, C. Brody, S. Thiberge, Princeton University 47

09.45a Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation
M. Campbell, S. Ocko, C. Mallory, S. Ganguli, L. Giocomo, Stanford University 48

10.00a **Coffee break**

Session 11: New vistas in neural integration

(Chair: Julijana Gjorgjieva)

10.30a Navigational attractor dynamics in the Drosophila brain: Going from models to mechanism
Vivek Jayaraman, HHMI Janelia Research Campus **(invited)** 33

11.15a A flexible model of working memory
F. Bouchacourt, T. Buschman, Princeton University 48

11.30a Motor preparation through rebound in an identified sensory integrator
M. Zwart, E. Yang, Z. Wei, N. Vladimirov, S. Narayan, M. Koyama, A. Abdelfattah, J. Grimm, L. Lavis, E. Schreiter, T. Kawashima, S. Higashijima, S. Druckmann, M. Ahrens, HHMI Janelia Research Campus 49

11.45p **Lunch break**

Session 12: Pathways to decisions

(Chair: Il Memming Park)

3.30p **Closing remarks**

Poster Session I — 7:30 pm Thursday 01 March

Refreshments are sponsored by the Simons Foundation

Poster Session II

7:30 pm Friday 02 March

Refreshments are sponsored by the Simons Foundation

Poster Session III

7:30 pm Saturday 03 March

Refreshments are sponsored by the Simons Foundation

The online version of this document includes abstracts for each presentation.

This PDF can be downloaded at: http://cosyne.org/cosyne18/Cosyne2018_program_book.pdf

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. Computation of instinctive escape decisions

Tiago Branco T.BRANCO@UCL.AC.UK
University College London

Escaping from imminent danger is an instinctive behaviour fundamental for survival that requires classifying sensory stimuli as harmless or threatening. The absence of threat allows animals to forage for essential resources, but as the level of threat and potential for harm increases, they have to decide whether or not to seek safety. Despite previous work on instinctive defensive behaviours in rodents, little is known about how the brain computes the threat level for initiating escape. Here we show that the probability and vigour of escape in mice scale with the intensity of innate threats, and are well described by a theoretical model that computes the distance between threat level and an escape threshold. Calcium imaging and optogenetics in the midbrain of freely behaving mice show that the activity of excitatory VGlut2+ neurons in the deep layers of the medial superior colliculus (mSC) represents the threat stimulus intensity and is predictive of escape, whereas dorsal periaqueductal gray (dPAG) VGlut2+ neurons encode exclusively the escape choice and control escape vigour. We demonstrate a feed-forward monosynaptic excitatory connection from mSC to dPAG neurons that is weak and unreliable, yet necessary for escape behaviour, which we suggest provides a synaptic threshold for dPAG activation and the initiation of escape. This threshold can be overcome by high mSC network activity because of short-term synaptic facilitation and recurrent excitation within the mSC, which amplifies and sustains synaptic drive to the dPAG. Thus, dPAG VGlut2+ neurons compute escape decisions and vigour using a synaptic mechanism to threshold threat information received from the mSC, and provide a biophysical model of how the brain performs a critical instinctive computation.

T-2. Collective sensing and decision-making in animal groups: From fish schools to primate societies

Iain D. Couzin[1,2] ICOUZIN@ORN.MPG.DE
[1]Max Planck Institute for Ornithology
[2]University of Konstanz

Understanding how social influence shapes biological processes is a central challenge in contemporary science, essential for achieving progress in a variety of fields ranging from the organization and evolution of coordinated collective action among cells, or animals, to the dynamics of information exchange in human societies. Using an integrated experimental and theoretical approach I will address how, and why, animals exhibit highly-coordinated collective behavior. I will demonstrate new imaging and virtual reality (VR) technology that allows us to reconstruct (automatically) the dynamic, time-varying sensory networks by which social influence propagates in groups. This allows us to identify, for any instant in time, the most socially-influential individuals, and to predict the magnitude of complex behavioral cascades within groups before they actually occur. By investigating the coupling between spatial and information dynamics in groups we reveal that emergent problem solving is the predominant mechanism by which mobile groups sense, and respond to complex environmental gradients. Finally I will reveal the critical role uninformed, or unbiased, individuals play in effecting fast and democratic consensus decision-making in collectives, and will test these predictions with experiments involving schooling fish and wild baboons.

T-3. Learning in neural circuits

Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK
Imperial College London

Learning is thought to be driven by changes in the connections between the neurons, a process called synaptic plasticity. In this talk, I will describe different approaches towards modelling learning in neural circuits. I will start by a top-down supervised learning approach to learn arbitrary behaviors in spiking neurons. Then, I will present a reinforcement learning set up, where we looked at the effect of acetylcholine in hippocampal plasticity and its implications for navigation towards a reward location. Finally, we will look at unsupervised learning, by modelling the maturation of auditory cortex ON and OFF responses across development.

T-4. Building models of the world for behavioural control

Timothy Behrens BEHRENS@FMRIB.OX.AC.UK
University of Oxford

I will talk about recent attempts in reinforcement learning that try to understand how models of the world might be represented and learnt in the brain to allow flexible control of behaviour. Relevant studies try to investigate neural codes and mechanisms that are used to organize this knowledge into a form that can be used efficiently and flexibly. The lecture will mostly focus on interactions between the frontal cortex and the medial temporal lobe. The neuronal codes and mechanisms discussed are often measured in both humans and model species, so there may be methodological interest in how to measure these mechanistic types of signals in humans.

T-5. What does dopamine mean?

Joshua Berke JOSHUA.BERKE@UCSF.EDU
University of California, San Francisco

Perhaps the most successful marriage between neuroscience and computer science is the theory that dopamine cell firing encodes "temporal-difference reward prediction errors" (RPEs). RPEs are learning signals, that serve to update predictions of future reward and thereby guide adaptive decision-making. Yet this theory is clearly incomplete, since dopamine is not just a learning signal but is also critically involved in motivation. Boosting dopamine immediately enhances willingness to work at behavioral tasks, and dopamine release ramps up as animals get closer to reward, in a manner that resembles motivational value rather than RPE. I will discuss our recent results and hypotheses that seek to reconcile and integrate these learning and motivational functions of dopamine

T-6. Understanding the relationship between neural variability and behavior

Marlene Cohen COHENM@PITT.EDU
University of Pittsburgh

The responses of pairs of neurons in visual cortex to repeated presentations of the same visual are typically correlated, but the importance of this correlated variability has been the subject of much debate. We took a practical approach, reasoning that if correlated variability is important for perception, it should be closely associated with visually guided behaviors. I will discuss results showing that correlated variability in visual cortex 1) has a consistent relationship to performance on perceptual tasks, regardless of whether performance is affected by cognitive factors that change quickly (e.g. attention), slowly (e.g. learning), or for reasons outside experimental control, 2) is oriented along the same dimensions in population response space that are most predictive of behavior on individual trials, and 3) is selectively communicated to downstream areas involved in decision making.

T-7. Brain-computer interfaces for basic science

Byron Yu BYRONYU@CMU.EDU
Carnegie Mellon University

Brain-computer interfaces (BCI) translate neural activity into movements of a computer cursor or robotic limb. BCIs are known for their ability to assist paralyzed patients. A lesser known, but increasingly important, use of BCIs is their ability to further our basic scientific understanding of brain function. In particular, BCIs are providing insights into the neural mechanisms underlying sensorimotor control that are currently difficult to obtain using limb movements. In this talk, I will demonstrate how a BCI can be leveraged to study how the brain learns. Specifically, I will address why learning some tasks is easier than others, as well as how populations of neurons change their activity in concert during learning.

T-8. Functional flexibility: State dependence in cortical circuits

Jessica Cardin JESS.CARDIN@YALE.EDU
Yale University

The physical components of cortical circuits, cells and synapses, are relatively stable. However, the operation of those circuits varies on a millisecond timescale with changes in behavioral state and cognition. Recent work has suggested that this extensive adaptation to changes in context and demand is supported by rapid flexibility in the ‘functional‘ circuit, the cells which are participating in circuit operations at any given moment. One potential mechanism that could promote the functional flexibility of cortical circuits is the diversity of GABAergic interneurons. Distinct populations of inhibitory interneuron have different intrinsic properties, activity profiles, and synaptic targeting, and are recruited into ongoing network activity under different conditions. To examine how different GABAergic populations regulate the evolving pattern of cortical circuit activity, we have explored the developmental and mature roles of VIP-expressing interneurons, whose activity is tightly correlated with behavioral state. We have further examined how interactions among different interneuron populations affect the degree to which soma- vs dendrite-targeting inhibition is activated under different behavioral conditions. We find a complex state-dependent relationship among VIP interneurons, somatostatin-expressing interneurons, and excitatory pyramidal neurons.

T-9. Sampling: coding, dynamics, and computation in the cortex

Mate Lengyel M.LENGYEL@ENG.CAM.AC.UK
University of Cambridge

Neural responses in the visual cortex are variable, and there is now an abundance of data characterizing how the magnitude and structure of this variability depends on stimulus attributes. These rich data sets allow us to revisit some of our basic assumptions about the computational roles and dynamical mechanisms of cortical variability. I will first argue that key aspects of cortical variability and correlations emerge naturally from a top-down theory of cortical population responses. According to this theory, population activity patterns represent statistical samples from the probability distribution arising from probabilistic inference, and thus their variability directly encodes perceptual uncertainty. Through direct comparisons to previously published data as well as original data analyses, I will show that a sampling-based probabilistic representation accounts for the structure of noise, signal, and spontaneous response variability and correlations in the primary visual cortex. I will then show that a simple model circuit, the stochastic stabilized supralinear network (SSN), featuring strong and fast non-normal amplification and non-linear interactions around a single attractor state, provides a mechanistic account of several key aspects of cortical variability, including the ubiquitously observed variability quenching following stimulus onset and the stimulus tuning-dependence of Fano factors and noise correlations in area V1 and MT. This model represents a qualitatively different regime of cortical dynamics than standard models that either rely on a slow, random walk-like exploration of a multitude of quasi-attractor states, or on chaotic spontaneous dynamics operating in the absence of any attractor states. Once again, through direct comparisons to previously published as well as original data analyses, we show that this regime accounts for the spatio-temporal patterns and temporal dynamics of variability suppression more comprehensively than previous proposals. Time permitting, I will also connect the two parts of the talk by showing that this dynamical regime may in fact be optimal for sampling-based inference. Taken together, these results allow us to take the interpretation of cortical variability beyond the well-entrenched realm of information transfer and put it centre stage in the study of cortical dynamics and computation.

T-10. Emergent dynamics from network connectivity: A minimal model

Carina Curto CPC16@PSU.EDU
Pennsylvania State University

Many networks in the brain display internally-generated patterns of activity — that is, they exhibit emergent dynamics that are shaped by intrinsic properties of the network rather than inherited from an external input. While a common feature of these networks is an abundance of inhibition, the role of network connectivity in pattern generation remains unclear. In this talk I will introduce Combinatorial Threshold-Linear Networks (CTLNs), which are simple "toy models" of recurrent networks consisting of threshold-linear neurons with effectively inhibitory interactions. The dynamics of CTLNs are controlled solely by the structure of an underlying directed graph. By varying the graph, we observe a rich variety of emergent dynamics including: multistability, neuronal sequences, and complex rhythms. These patterns are reminiscent of population activity in cortex, hippocampus, and central pattern generators for locomotion. I will present some theorems about CTLNs, and explain how they allow us to predict features of the dynamics by examining properties of the underlying graph.

T-11. The striatum and decision-making based on value

Ann M. Graybiel GRAYBIEL@MIT.EDU
Massachusetts Institute of Technology

Affective decision-making is a cornerstone of our responses in daily life and reflects crucial decision-making dependent on cost-benefit integration. Regions of the medial prefrontal cortex are known to function in organizing behavior and emotional decision-making, both key functions disturbed in a range of neurologic and neuropsychiatric disorders. In our laboratory, we are seeking to understand mechanisms underlying these functions by applying optogenetic manipulations and microstimulation to these regions and their related corticostriatal circuits. We find that we can interrupt the transition from deliberative decision-making to a habitual mode of decisions to act, and we can interrupt habitual and insistently repetitive behaviors. In other experiments, we can selectively disrupt decision-making under different contexts involving weighing the costs and benefits of the offered options. We further find that chronic exposure to stress disrupts the dynamics of corticostriatal circuits. Our models of these experimental observations suggest that stress alters the integration of cost and benefit, and that cost-benefit ratios (CBRs) derived by modeling are key to relating the behavior to neural indicators of circuit function and dysfunction. These experiments point to profoundly important functions of corticostriatal circuits, and to their exquisite specialization. These features could be critical in linking these corticostriatal circuit mechanisms to human disorders in which value-based decision-making is affected.

T-12. Navigational attractor dynamics in the Drosophila brain: Going from models to mechanism

Vivek Jayaraman VIVEK@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

Decades of neurophysiological work in mammals have revealed several evocative neural representations thought to be involved in navigation. Although numerous theoretical models have been proposed for the generation and use of these representations, the complexity of navigational circuits has made it difficult to link theory to biological implementation. My lab is studying how visual landmark and self-motion information is used to create abstract and persistent navigational representations in recurrent circuits of an organism with a much smaller brain, Drosophila melanogaster. We use two-photon calcium imaging, whole-cell patch clamp electrophysiology and optogenetics in head-fixed behaving flies to explore compass-like attractor dynamics in a brain region called the central complex, a higher-order sensorimotor processing center known to be conserved across arthropods. Combining results from physiological and behavioral experiments with electron microscopic reconstruction and theoretical modeling, we are uncovering how visual inputs are processed to obtain the animal‘s orientation relative to landmarks in its surroundings, and how the fly‘s movements update its heading representation. More broadly, I will make the case that the fly‘s powerful experimental toolkit can be exploited to uncover the mechanistic underpinnings of broadly relevant neural dynamics.

T-13. Dynamics of prefrontal computations during decision-making

Joni Wallis WALLIS@BERKELEY.EDU

University of California, Berkeley

A major challenge to understanding the neural mechanisms underlying cognitive processes is that these processes cannot be directly observed, but rather must be inferred from behavioral measures. Furthermore, there could be considerable variability in these processes from one iteration to the next. Because neuronal responses are inherently stochastic, studies of cognitive processes typically average activity across many repeated trials. However, when the dynamics of those processes vary, this approach can obscure critical mechanistic details. In the first part of my talk, I will describe recent studies in my lab which have uncovered the dynamics of decision-making in orbitofrontal cortex with single trial resolution by leveraging the power of decoding ensemble activity by recording from many orbitofrontal neurons simultaneously. During individual choices, neural representations alternate between states associated with each available option, as if the network were considering them in turn. In the second part of my talk, I will discuss our attempts to modify these dynamics using electrical microstimulation to examine whether we can alter decision-making. If successful, this would be a first step to building a brain implant that might be able to modify maladaptive behaviors underlying neuropsychiatric disorders such as addiction and obsessive-compulsive disorder.

T-14. Unsupervised discovery of neural sequences in large-scale recordings

Emily Mackevicius[1] ELM@MIT.EDU
Andrew Bahle[1] ABAHLE@MIT.EDU
Alex Williams[2] ALEX.H.WILLIA@GMAIL.COM
Shijie Gu[3,4] TECHEL@LIVE.CN
Natalia Denissenko[1] NDENISEN@MIT.EDU
Mark Goldman[5] MSGOLDMAN@UCDAVIS.EDU
Michale Fee[1] FEE@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Stanford University
[3]Shanghaitech University
[4]Harvard University
[5]University of California, Davis

A major challenge in neuroscience is to identify relevant, low-dimensional features that capture the dynamics of large-scale neural recordings. However, dynamics that include repeated temporal patterns (which we call sequences), are not captured by traditional dimensionality reduction techniques such as PCA and NMF. Recently, the importance of neural sequences has been demonstrated using visual display of trial-averaged firing rates [Fujisawa et al., 2008; Pastalkova et al., 2008; Harvey et al., 2012]. However, the field suffers from a lack of task-independent unsupervised tools for identifying sequences directly from neural data.

We propose a tool that extends a convolutional NMF technique to model the data using a minimal set of sequence factors. Our method, which we call seqNMF, provides a framework for extracting sequences in the data, the times at which they occur, and the statistical significance and variance explained for each extracted factor. In addition, seqNMF identifies each repetition of a sequence and can be easily cross-validated. We applied seqNMF to recover sequences in previously published datasets from rat hippocampus and prefrontal cortex, as well as new data from juvenile songbird vocal cortical HVC. In the rodent data, our algorithm blindly identified sequences that match published sequences calculated by reference to task epoch [Fujisawa et al., 2008; Pastalkova et al., 2008]. In the songbird data, the algorithm also discovered sequences, some of which were clearly associated with particular song syllables. Notably, other extracted sequences were “latent”–they were initially uncorrelated with song but gradually acquired tight locking to a newly learned song syllable. Thus, by identifying temporal structure directly from neural data in an unsupervised manner, seqNMF can enable dissection of complex neural circuits with noisy or changing behavioral readouts.

T-15. Supervised learning of unsupervised local learning rules

Brian Cheung[1] BCHEUNG@BERKELEY.EDU
Luke Metz[2] LMETZ@GOOGLE.COM
Jascha Sohl-Dickstein[2] JASCHASD@GOOGLE.COM

[1]University of California, Berkeley
[2]Google Brain

Supervised learning has proven extremely effective for many problems in machine learning where large amounts of labeled training data are available. However, the dependence on large labeled datasets and non-local updates make it unclear how similar algorithms might function in the brain. Conversely, biological neural networks are extremely effective at building rich, high utility, representations of sensory input with little or no labeled training data. However, unsupervised representation learning in artificial neural networks lags far behind both biological networks, and supervised artificial networks. One explanation for our failure to develop effective unsupervised learning rules is that the objective functions we propose are mismatched to the behaviorally relevant tasks for which we wish to use the learned representation. We optimize objectives such as log likelihood, sparsity, or reconstruction error, and then hope a learned representation which exposes high-level features of sensory data relevant to survival will result purely as a side effect.

Rather than proposing a hand-designed update rule, in this work we use supervised training to play the role of evolution in discovering an update rule for biological neural networks. Specifically, we perform supervised training of the unsupervised learning rule, so that it leads to representations which maximize a biologically plausible utility function. Additionally, we parameterize the learned update rule itself in a biologically plausible way. We meta-learn a local learning rule that only depends on bottom-up input from the pre-synaptic neuron and top-down feedback from the post-synaptic neuron. By re-casting unsupervised learning as meta-learning, we directly optimize an unsupervised learning rule with respect to its utility. We argue that this is a natural approach to unsupervised learning in the context of biology. Our work offers a preliminary investigation of unsupervised learning rules meta-learned using this novel perspective.

T-16. Mid-lateral cerebellar Purkinje neurons participate in visuomotor associative learning

Naveen Sendhilnathan[1] NAVEENSENDHILNATHAN@GMAIL.COM
Michael Goldberg[2] MEG2008@CUMC.COLUMBIA.EDU

[1]Columbia University
[2]Columbia university

The cerebellum has been primarily considered to have roles in motor coordination. Recent clinical, anatomical and electrophysiological evidence suggest a cognitive role for the cerebellum. We trained Rhesus monkeys to associate well-learned hand movements with arbitrary visual symbols. When the monkeys started to learn a new visuomotor association, the Purkinje neurons immediately showed a change in firing activity. In addition, during learning, Purkinje neurons reported the prior trial's outcome: simple spike activity differed between prior correct and prior wrong trials; but only in a particular epoch of the trial for each neuron. Across the population, the epochs tiled the whole trial period. The neurons reported the trial outcome independent of changes in reaction time, hand movement, visual symbol novelty and reward expectation. The neurons used this trial outcome to track the learning process. Our results argue for the role of cerebellum in higher order processing and gather evidence that cerebellum, rather than being regarded just as a motor control system, could be a generalized control system, essential in cognitive and rule learning as well as motor learning and adaptation.

T-17. Circuit mechanisms for negative valence extends temporal window of memory-linking retrospectively

Denise Cai[1] DENISECAI@GMAIL.COM
Daniel Aharoni[2] DBAHARONI@GMAIL.COM
Zhe Dong[1] ZHE.DONG@ICAHN.MSSM.EDU
Tristan Shuman[1] TRISTANSHUMAN@GMAIL.COM
Alcino Silva[2] ALCINOSILVA@GMAIL.COM

[1]Icahn School of Medicine at Mount Sinai
[2]University of California, Los Angeles

Memories are not formed in isolation. Most experiences involve the integration of multiple memories across time, with one memory affecting how others are encoded, stored and retrieved. We recently showed one way in which memories interact is through the linking of experiences encoded close in time (Cai et al., 2016). Two neutral experiences encoded within a day are functionally linked by an overlapping neural ensemble, such that recalling one memory will trigger the recall of another temporally-related memory. These findings have important implications for the linking of neutral memories, but it remains unclear how negative valence alters this process. This is particularly important for anxiety disorders such as posttraumatic stress disorder (PTSD), where unusually strong connections between aversive and neutral memories may contribute to pathological symptoms, such as experiencing anxiety in safe environments. We found that increasing negative valence during initial learning extends the memory-linking window from hours to days. This memory-linking is asymmetric, as the fear transfers retrospectively to contexts experienced up to 2 days before, but not days later. Using a new generation of Miniscopes for in vivo calcium imaging, we recorded neural activity during the encoding and retrieval phases of memory across different contexts. We found there was an increase in the neural ensemble overlap between the fearful and neutral contexts (compared to two neutral contexts) during the retrieval process that was not present during the encoding of the contexts. These findings suggest a latent circuit mechanism that retrospectively links fear from an aversive memory to past safe memories by increasing the neural overlap after initial encoding (e.g., co-reactivation of neural ensembles during sleep). The findings have profound implications for the coordinated storage and retrieval of memories and alteration of these memory processes impairs the efficient and effective recall of information, such as in PTSD.

T-18. Social behavior shapes hypothalamic neural ensemble representations of conspecific sex

Ryan Remedios[1] RYAN@CALTECH.EDU
Ann Kennedy[1] KENNEDYA@CALTECH.EDU
Moriel Zelikowsky[1] MORIEL@CALTECH.EDU
Benjamin Grewe[2] BGREWE@ETHZ.CH
Mark Schnitzer[3] MSCHNITZER@GMAIL.COM
David Anderson[1] WUWEI@CALTECH.EDU

[1]California Institute of Technology
[2]ETH Zurich
[3]Stanford University

All animals possess a repertoire of innate (or instinctive) behaviors, such as aggression, mating, and feeding, that can be performed without training. Because these behaviors are critical to survival, it is often assumed that they are mediated by anatomically distinct and/or genetically hard-wired neural pathways, although this remains a matter of debate. Here we report that hypothalamic neural ensemble representations underlying innate social behaviors are in fact shaped by social experience. Estrogen receptor 1-expressing (Esr1+) neurons in the ventrolateral subdivision of ventromedial hypothalamus (VMHvl) control mating and fighting in rodents. We used microendoscopy to image VMHvl Esr1+ neuronal activity in male mice engaged in these behaviors. The majority of the variance in Esr1+ activity was accounted for by the sex of the conspecific. This was surprising, given that optogenetic manipulations of this population can promote or inhibit attack behavior. Nevertheless, some neurons did show tuning for the animal's actions, and many cells showed mixed selectivity for multiple cues. In sexually and socially experienced adult males, distinct neural ensembles represented male vs. female conspecifics. But surprisingly, in inexperienced adult males, male and female intruders activated overlapping populations, with sex-specific ensembles gradually separating as the mice acquired social and sexual experience. In mice permitted to investigate but not mount or attack conspecifics, ensemble divergence did not occur—but 30 minutes of sexual experience with a female was sufficient to promote both male vs. female ensemble separation and attack. By tracking neurons across multiple days, we were able to characterize the ensemble separation process, and propose models of the learning involved. Our observations uncover an unexpected social experience-dependent component to the formation of hypothalamic neural assemblies controlling innate social behaviors. More generally, they reveal plasticity and dynamic coding in an evolutionarily ancient, deep subcortical structure that is

traditionally viewed as a “hard-wired“ system.

T-19. Sustained activity in midbrain dopamine neurons encodes information predicting distant reward

Akash Guru AG2249@CORNELL.EDU
Andrew Recknagel AKR33@CORNELL.EDU
Julia Schaffer JAS872@CORNELL.EDU
Durga Kullakanda DSK257@CORNELL.EDU
Changwoo Seo CS935@CORNELL.EDU
Melissa Warden MRWARDEN@CORNELL.EDU

Cornell University

Organisms are often motivated by goals that are spatially and temporally distant such as navigating toward food sources or seasonal migration sites. How does the brain enable reinforcement when goals are spatially and temporally distant from the current state? Natural behavior involves a stream of sensorimotor information that is predictive of future distant goals. Does the dopamine system use such a continuous stream of predictive information to reinforce distant goals? Dopaminergic neural activity in the ventral tegmental area (VTA) of mice was recorded using fiber photometry while they performed two different reward seeking tasks. In one, the animals spatially navigated through a maze to collect reward. We discovered that dopaminergic neural activity increased continuously as the mice approached the reward port and peaked at the reward. The ramping activity was sensitive to reward location, average reward value, and uncertainty. There was no obvious correlation between the activity and velocity or motivation. We note that spatial navigation provides a stream of sensorimotor information that is predictive of the distance to reward from the current state and could produce a continuous dopamine signal. To test this hypothesis, animals were trained in a second task in which the motor behavior is similar, but sensory information predicting the reward is minimal. This did not result in a similar continuously rising dopamine signal. Together, these findings identify a novel ramping activity pattern in the dopaminergic neurons in VTA, brings together ideas on origin of ramping dopamine concentration in striatum (1’3), and reveal that dopamine neurons utilize the external stream of predictive sensory information to aid learning.

T-20. Signatures of a wandering mind: multiple cell types across the basal ganglia reflect task engagement

Sofia Soares[1] SOFIA.SOARES@NEURO.FCHAMPALIMAUD.ORG
Bassam V. Atallah[1] BASSAM.ATALLAH@NEURO.FCHAMPALIMAUD.ORG
Asma Motiwala[2] ASMA.MOTIWALA@NEURO.FCHAMPALIMAUD.ORG
Bruno F. Cruz[2,3] BRUNO.CRUZ@NEURO.FCHAMPALIMAUD.ORG
Tiago Monteiro[1] TIAGO.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG
Thiago Gouvea[4] TSGOUVEA@CSHL.EDU
Joseph J. Paton[1] JOSEPH.PATON@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Research
[2]Champalimaud Foundation
[3]Champalimaud Neuroscience Program
[4]Cold Spring Harbor Laboratory

Decisions are a function of immediate sensory input, prior experience and internal states. While decision-making neuroscience has focused extensively on the influence of stimulus features and past history on choice, how internal states affect performance of tasks used to study decision-making is poorly understood. We trained rodents on a demanding interval discrimination task and used generalized linear models (GLMs) to identify behavioral

and task-imposed factors that predicted discrimination performance and adherence to task demands. Using fiber photometry, we recorded the activity of midbrain dopamine (DA) neurons and found that activity prior to trial initiation (pre-trial) was predictive of animals' discrimination accuracy: higher pre-trial DA activity predicted lower accuracy. Pre-trial DA activity was also higher when animals failed to complete a trial, either because they responded impulsively before stimulus termination or because they did not respond at all, suggesting that pre-trial DA activity reflected a general state of task engagement. Surprisingly, we observed similar signals (albeit of opposite sign) in pre-trial activity of both direct and indirect pathway striatal medium spiny neurons (MSNs). Critically, adding pre-trial neural activity (from DA neurons or MSNs) to the GLM improved predictions of both accuracy and impulsivity, suggesting that these neural signals exerted an additional influence on behavior beyond that accounted for by observable task-related factors. Based on previous results indicating that within-trial striatal population dynamics correlate with animals' judgements, we hypothesized that pre-trial BG network state might predispose it towards aberrant within-trial dynamics, leading to poorer performance. We quantified 'atypicalness' of within-single-trial population activity and found that it predicted animals' performance. Together, these data suggest that the state of BG networks at trial initiation modulates the fidelity with which the decision variable is encoded within the trial, providing insight into the circuit mechanisms by which engagement may influence performance of tasks requiring cognitive control.

T-21. Dopamine neurons projecting to the tail of the striatum reinforce avoidance of threatening stimuli

William Menegas WILLIAM.S.MENEGAS@GMAIL.COM
Korleki Akiti AKITI@G.HARVARD.EDU
Naoshige Uchida UCHIDA@MCB.HARVARD.EDU
Mitsuko Watabe-Uchida MITSUKO@MCB.HARVARD.EDU

Harvard University

Two key observations support a role for dopamine in reinforcement learning: 1) that dopamine neurons are activated by reward and inhibited by punishment, and 2) that dopamine release positively reinforces behaviors that result in dopamine release. The discovery that some dopamine neurons are activated by aversive stimuli challenged this view. Our previous study (Menegas et al., 2017) found that dopamine neurons projecting to the posterior tail of the striatum (TS) respond to aversive and novel stimuli. The function of these dopamine neurons, however, remained unknown. In the present study, we first demonstrate that dopamine axons in TS are excited by only a subset of negative events such as air puffs or loud tones, but not by bitter taste or the omission of an expected reward—suggesting that dopamine axons in TS specifically broadcast an external threat signal. We next demonstrate that animals avoid optogenetic activation of dopamine axons in TS during a choice task, in stark contrast to the widely-held notion that dopamine release is rewarding. Conversely, the ablation of TS-projecting dopamine neurons reduced avoidance of a threatening stimulus (air puff) but not of bitter water or reduction of reward size. Initial avoidance responses evoked by air puff were intact in lesion mice, suggesting that dopamine in TS is required for reinforcing avoidance behaviors rather than for detection of air puff. We also show that ablation of TS-projecting dopamine neurons altered animals' behavioral responses to a novel object by specifically reducing their retreats over time without affecting their approach behavior. We propose that the activation of TS-projecting dopamine neurons positively reinforces stimulus-based threat predictions or avoidance behaviors evoked by threatening stimuli. Taken together, our results indicate that there are at least two axes of reinforcement learning using dopamine: one learning from outcome values, and another learning from potential external threats.

T-22. How does activity in D1R and D2R-expressing neurons in the dorsomedial striatum influence choice in a non-lateralized task?

Kristen Delevich DELEVICH@BERKELEY.EDU
Benjamin Hoshal BDHOSHAL@BERKELEY.EDU
Yuting Zhang YUTING.ZHANG@BERKELEY.EDU
Satya Vedula SATYAVEDULA@BERKELEY.EDU
Anne Collins ANNECOLLINS@BERKELEY.EDU
Linda Wilbrecht WILBRECHT@BERKELEY.EDU

University of California, Berkeley

Optimal foraging strategy relies on an animal's ability to negotiate two opposing choices: exploit a current source of reward or explore lesser known, potentially better options. The dorsomedial striatum (DMS) plays a key role in action selection, and it is believed that action values are stored in synaptic weights onto direct (D1R+) and indirect (D2R+) neurons. Most data linking D1R+ and D2R+ neural activity and choice comes from lateralized tasks, in which DMS neurons exhibit marked contralateral bias. We aimed to understand how D1R+ and D2R+ DMS neurons contribute to value-based decision-making using an odor-based foraging task that 1) could not be solved using lateralized egocentric or spatial information 2) was rapidly acquired 3) drove robust learning and 4) was well captured by reinforcement learning (RL) models. Briefly, mice were trained to choose among 4 pots filled with scented shavings to retrieve food reward. Only one odor was rewarded and mice were trained to criterion (8/10 correct). The following day mice were tested in their recall of the learned odor discrimination. We found that chemogenetic inhibition of D1R+ neurons or chemogenetic excitation of D2R+ neurons similarly influenced behavior in recall: in both groups we observed an increase in non-rewarded choice selection after the first correct recall. Notably the rank values of the odor choices from the day before were preserved. We hypothesize these non-rewarded choices reflect a more exploratory choice policy (vs. exploitative). We are currently evaluating this hypothesis using reinforcement learning models. Our data do not support popular alternative models that suggest that (in a non lateralized context) activity in D2+ neurons in the DMS should drive simple rejection, action suppression, or "no-go" function.

T-23. Linking feature-selective attention and decision-related activity in the macaque visual cortex

Katrina Quinn[1] KATRINAROSEQUINN@GMAIL.COM
Stephane Clery[2] CLERYSTEPHANE@GMAIL.COM
Paria Pourriahi[2] PARIA1364@YAHOO.COM
Hendrikje Nienborg[2] HENDRIKJE.NIENBORG@CIN.UNI-TUEBINGEN.DE

[1]University of Tuebingen
[2]Werner Reichardt Centre for Integrative Neuroscience

During perceptual decisions the activity of task-relevant sensory neurons is typically correlated with an animal's decision. Such decision-related activity, often quantified as "choice-probability", is thought to reflect feed-forward and feedback sources [e.g.1]. For a feature-discrimination task, this typically implies that the feedback is feature-selective. A recent computational account of choice-probabilities [2] hypothesizes that the modulation of sensory neurons by feature-selective attention and the feedback source of choice-probabilities reflect the same mechanism. Here, we tested this hypothesis. To do so, we leveraged classical findings for feature-selective attention. When an animal attends to a stimulus feature, the responses of neurons selective for this feature are modulated [3]. Critically, this modulation is observed globally, i.e. even when the attended, task-relevant stimulus is in the opposite hemifield to the receptive field of the modulated neuron. Together with the above hypothesis this finding makes a strong prediction: choice-probabilities should also be observed globally, including for a task-irrelevant, ignored stimulus.

To test this prediction we trained macaques on a disparity-discrimination task on one of two random-dot stereograms, each presented in one hemifield. Only one stimulus, always validly cued, was task-relevant and informative. We blockwise switched the hemifield in which the task-relevant stimulus was presented. Once the animals successfully ignored the task-irrelevant stimulus, we performed multichannel recordings from disparity selective units in area V2. In support of the hypothesis, we find substantial choice-probabilities for the ignored stimulus (mean CP=0.59), slightly weaker than and correlated with those for the task-relevant stimulus (mean CP=0.62; r=0.41, p<10-3, n=71). Importantly, this would not be expected in a feed-forward account, in which choice-probabilities reflect the read-out of the information used by the animal, or in feedback accounts, which differ from feature-selective attention. Instead, these results provide a novel, but predicted, link between feature-selective attention and decision-related activity in sensory neurons.

T-24. Shared stochastic modulation can facilitate biologically plausible decoding

Caroline Haimerl CH2880@NYU.EDU
Eero Simoncelli EERO.SIMONCELLI@NYU.EDU

New York University

Sensory behaviors are often described in terms of encoding (from stimuli to neural responses) and decoding (from neural responses to behavior). Given full knowledge of encoding (stimulus selectivity, response variability), one can define statistically optimal decoders that provide an upper bound on performance. However, flexible and accurate decoding relies on identification of behaviorally-informative neurons (Britten et al., 1992). No current theory provides an explanation of how this can be achieved in the brain. As an example, in experiments by Cohen & Maunsell (2009) populations of V4 cells were monitored while a monkey performed an orientation discrimination task. A small subset of spatially scattered cells carried information relevant to the task in any given block of trials, and yet the animal performed the task well after viewing only a few examples. Optimizing even a linear decoder requires access to encoding details, or large amounts of labeled training data. A subsequent model-based analysis of these data (Rabinowitz et al., 2015) revealed that cells are gain-modulated by a common fluctuating low-dimensional signal, and neuron-specific strength of this modulation is correlated with task-specific informativeness. Here, we propose that this modulatory signal can serve as a key element in solving the identification step of decoding. We simulated a population of modulated Poisson neurons (Goris et al., 2014) responding to two stimuli with rate and modulation parameters matching those fit in Rabinowitz et al. (2015). We find that a linear decoder, with readout weights estimated from the strength of modulation, but without knowledge of the encoding model, achieves performance close to an optimal (ML) decoder assuming Poisson noise. Performance depends on the strength of the modulatory signal; poor at low and high levels and best for moderate levels. We conclude that fluctuating modulatory signals could be used to flexibly and rapidly identify task-specific neurons for decoding.

T-25. Learning can generate new patterns of neural population activity

Emily Oby[1] EMILYOBY@GMAIL.COM
Matthew Golub[2,3] GOLUB.MATTHEW@GMAIL.COM
Jay Hennig[2] JHENNIG@ANDREW.CMU.EDU
Alan Degenhart[1] ALAN.DEGENHART@PITT.EDU
Elizabeth Tyler-Kabara[1] ELIZABETH.TYLER-KABARA@CHP.EDU
Byron Yu[2] BYRONYU@CMU.EDU
Steve Chase[2] SCHASE@CMU.EDU
Aaron Batista[1] APB10@PITT.EDU

[1]University of Pittsburgh
[2]Carnegie Mellon University

[3] Stanford University

How does neural population activity change to facilitate learning? We study the changes in neural population activity that accompany learning using a brain-computer interface (BCI) in which Rhesus monkeys modulate primary motor cortical activity to control a computer cursor. We have previously shown that neural population activity tends to lie in a low-dimensional subspace, termed the intrinsic manifold, and that the intrinsic manifold constrains learning. Novel BCI mappings that are consistent with the intrinsic manifold are more readily learned than those that are not. Here, we show that monkeys can indeed learn new mappings that are inconsistent with the intrinsic manifold (termed outside-manifold perturbations, OMPs), but doing so takes many days. We then considered three possible neural strategies for how neural activity patterns reorganize during OMP learning. One strategy is to take advantage of the existing population activity patterns, termed the neural repertoire, by reassociating them with different intended movements. This is how mappings consistent with the intrinsic manifold are learned within a single day. While this strategy would lead to behavioral improvements, control is not optimal. Rather, the optimal strategy is to generate new population activity patterns that are outside the neural repertoire and aligned with the OMP mapping. These new patterns could be inside the intrinsic manifold or outside the intrinsic manifold. Here, we show for the first time that animals exhibit new neural population activity patterns with long-term OMP learning. We find evidence of both inside manifold and outside manifold strategies, even within an experiment. Our work demonstrates that the constraints of the intrinsic manifold can be violated, but that it takes extended practice over several days to do so.

T-26. Control of sensorimotor dynamics through adjustment of inputs and initial condition

Evan Remington EREMING@MIT.EDU
Devika Narain DNARAIN@MIT.EDU
Eghbal Hosseini EHOSEINI@MIT.EDU
Mehrdad Jazayeri MJAZ@MIT.EDU

Massachusetts Institute of Technology

Sensorimotor computations can be flexibly adjusted according to internal states and contextual inputs. The mechanisms supporting this flexibility are not understood. Here, we tested the utility of a dynamical systems perspective to approach this problem. In a dynamical system whose state is determined by interactions among neurons, computations can be rapidly and flexibly reconfigured by controlling the system's inputs and initial conditions. To investigate whether the brain employs such control strategies, we recorded spiking activity from the dorsomedial frontal cortex (DMFC) of monkeys trained to measure time intervals and subsequently produce timed motor responses according to context-specific stimulus-response rules. This timing task permitted the investigation of internally-generated neural dynamics in the absence of potentially confounding time varying sensory and reafferent inputs. We applied a novel state-space analysis to population activity recorded in DMFC to reveal that the geometry of neural trajectories were organized hierarchically according to inputs and initial conditions. Within contexts, neural trajectories associated with different produced intervals evolved from distinct initial conditions through the response epoch within an ordered structure. The position of trajectories within this structure controlled the speed of neural dynamics in the production epoch, allowing the animal to aim for the target interval. Changing the stimulus-response rule resulted in a translation of the structure in state space, as predicted by a tonic external input. In-silico activity in recurrent neural network models trained to perform the same task with or without a tonic external input corroborated this dynamical interpretation of the DMFC data. These results provide evidence that the language of dynamical systems can be used to describe the algorithms brain uses for sensorimotor coordination, and provide new insights into the mechanisms by which the brain selects activity patterns driving flexible behavior.

T-27. Motor primitives in space and time via targeted gain modulation in cortical networks

Jake Stroud[1] JAKE.STROUD@WADH.OX.AC.UK
Mason Porter[2] MASON@MATH.UCLA.EDU
Guillaume Hennequin[3] G.HENNEQUIN@ENG.CAM.AC.UK
Tim Vogels[1] TIM.VOGELS@CNCB.OX.AC.UK

[1]University of Oxford
[2]University of California Los Angeles
[3]University of Cambridge

Animals perform an extraordinary variety of movements over many different timescales. To support this diversity, the motor cortex (M1) exhibits a similarly rich repertoire of activity (Shenoy et al., 2013). Recent neuronal network models capture many qualitative aspects of M1 dynamics, but they can generate only a few distinct movements all with the same duration (Hennequin et al., 2014 and Sussillo et al., 2015). Therefore, these models can still not explain how M1 efficiently controls movements over a wide range of speeds and shapes. Here we demonstrate that simple modulation of neuronal input-output gains in recurrent network models with fixed connectivity can substantially and predictably affect downstream muscle outputs. Consistent with the observation of diffuse neuromodulatory projections to M1 (Huntley et al., 1992 and Hosp et al., 2011), our results suggest that a relatively small number of modulatory control units provide sufficient flexibility to adjust high-dimensional network activity on behaviourally relevant timescales. Such modulatory gain patterns can be obtained through a simple reward-based learning rule. Novel movements can also be assembled from previously learned primitives, thereby facilitating fast acquisition of hitherto untrained muscle outputs. Moreover, we show that it is possible to separately change movement speed while preserving movement shape, thus enabling efficient and independent movement control in space and time. Our results provide a new perspective on the role of neuromodulatory systems in controlling cortical activity and suggests that modulation of single-neuron excitability is an important aspect of learning.

T-28. Mechanisms underlying sharpening of visual response dynamics with familiarity

Sukbin Lim[1] SUKBIN.LIM@NYU.EDU
Nicolas Brunel[2] NICOLAS.BRUNEL@DUKE.EDU

[1]New York University, Shanghai
[2]Duke University

Experience-dependent modifications of synaptic connections are thought to change patterns of network activity with learning. In monkey inferotemporal cortex (ITC), changes of both time-averaged visual responses, and their dynamics, have been reported with familiarity. The average over visual stimuli of time-averaged visual responses decreases with familiarity, while the distribution of responses across visual stimuli broadens with learning [1-3]. The dynamics of visual responses was also found to change with familiarity—in particular, rapid successive presentation of familiar images, but not novel images, elicits strong periodic responses [3]. Previously, we investigated synaptic plasticity in recurrently connected circuits that reproduces changes of the distribution of time-averaged visual responses observed experimentally [4]. Here, we extended this framework to understand how the interaction between synaptic plasticity and various negative feedback mechanisms shapes response dynamics with learning. We found that a fatigue mechanism analogous to firing rate adaptation, together with Hebbian synaptic plasticity, can explain the changes of response dynamics observed experimentally. When novel stimuli are shown repeatedly, the peak response to the second stimuli is smaller than the response to the first, due to slow recovery from the adaptation current. In contrast, for serial presentation of familiar stimuli, the interaction between the positive feedback due to Hebbian-plasticity and the negative feedback due to firing rate adaptation leads to a resonance phenomenon, and to damped oscillations in response to the first familiar stimulus. Consequently, the

response to the second stimulus is less affected by the adaptation current, and the peak response can be almost as strong as the first one. Our results provide a mechanistic understanding of how interactions between synaptic plasticity and a negative feedback mechanism implementing firing rate adaptation shape network response dynamics, and account for experimental observations about the effects of visual experience on response dynamics of ITC neurons.

T-29. Predictive coding of novel versus familiar stimuli in the primary visual cortex

Jan Homann JHOMANN@PRINCETON.EDU
David Tank DWTANK@PRINCETON.EDU
Sue Ann Koay SUEANNKOAY@GMAIL.COM
Alistair M. Glidden AGLIDDEN@PRINCETON.EDU
Michael J. Berry II BERRY@PRINCETON.EDU

Princeton University

Predictive coding has long been postulated as a core set of canonical cortical computations. The idea is that only surprising signals are passed along the sensory hierarchy, whereas predictable signals are not. Unpredictable changes in the environment are what ultimately drives many behaviors and it seems therefore logical that significant computational resources should be devoted to surprising sensory signals. In our research, we searched for signatures of predictive coding in the primary visual cortex of mice. We presented mice with repeated sequences of images with novel images sparsely substituted. Under these conditions, mice could be rapidly trained to lick in response to a novel image, demonstrating a high level of performance on the first day of testing. Using 2-photon calcium imaging to record from layer 2/3 neurons in the primary visual cortex, we found that novel images evoked excess activity in the majority of neurons. When a new stimulus sequence was repeatedly presented, a majority of neurons had similarly elevated activity for the first few presentations, which then decayed to almost zero activity. The decay time of these transient responses was not fixed, but instead scaled with the length of the stimulus sequence. However, at the same time, we also found a small fraction of the neurons within the population (~2%) that continued to respond strongly and periodically to the repeated stimulus. Decoding analysis demonstrated that both the transient and sustained responses encoded information about stimulus identity. We conclude that the layer 2/3 population uses a two-channel predictive code: a dense transient code for novel stimuli and a sparse sustained code for familiar stimuli. These results extend and unify existing theories about the nature of predictive neural codes.

T-30. On the complexity of predictive strategies in noisy and changing environments

Gaia Tavoni TAVONI@SAS.UPENN.EDU
Vijay Balasubramanian VIJAY@PHYSICS.UPENN.EDU
Joshua Gold JIGOLD@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

To thrive in a noisy and dynamic world, it is important to learn from past experiences to make effective predictions about future events. It is generally assumed that in noisy environments with unsignaled change-points, predictive strategies must be adaptive to dynamically take into account only the most relevant past information. Here we challenge this assumption and provide novel insights into the role of adaptivity and memory in predictive-inference tasks. We show that only a narrow range of change-point rates (h) and signal-to-noise ratios (SNR) requires adaptive strategies. When $h > \sim 0.1$, adaptive models do not provide substantial improvements over non-adaptive ones. When SNR is very low or very high, the non-adaptive domain extends to even lower change-point

rates. Thus, simple predictive strategies are widely preferable in extreme and opposite noise conditions. Increasing the change-point rate beyond ~0.3 further reduces the demands for adaptive, history-dependent models to make effective predictions. We further characterize the non-adaptive domain, by inferring the most effective integration kernels. We show a phase transition when noise (1/SNR) is equal to the Golden Ratio. Below this value, an exponentially decaying delta-rule integration kernel with a time constant that decreases with increasing h optimizes a trade-off between model simplicity and model accuracy at low h, whereas at large h the trade-off is optimized by a model with no memory of past events. When noise is above the Golden Ratio, the optimal time constant of the delta-rule model has a stationary point that also marks a transition to a domain in which a more memory-demanding sliding-average model provides substantial benefits. Together the results provide a principled framework for understanding the need for adaptivity and memory in dynamic environments, including an “inverted-U” relationship with SNR that requires complex, adaptive models only for particular conditions with intermediate levels of SNR.

T-31. A relational odor map in piriform cortex

Stan Pashkovski STANLP86@GMAIL.COM
Guliano Iurilli GIULIANO.IURILLI@GMAIL.COM
Sandeep Datta SRDATTA@HMS.HARVARD.EDU
Harvard University

Odorous molecules trigger specific percepts. Appropriate assignment of odorants to corresponding percepts relies on the brain’s ability to both discriminate distinct odorants, as well as generalize odorants that share chemical features. Although odorants evoke correlated activity across receptor expressing neurons in the olfactory epithelium and glomeruli residing in the olfactory bulb, it is unclear how the chemical attributes of odors are encoded in cortex to support both discrimination and generalization. To address this question, we have developed a preparation that allows us to monitor odor representations in awake mice across cortical layers in the main region of the brain devoted to olfactory processing, the Piriform Cortex (PCx). Using sets of odorants rationally designed to tile different regions of chemical odor space at multiple scales of resolution, we demonstrate that responses of neural ensembles in both Layer II (LII) and Layer III (LIII) of PCx preserve information about the physicochemical relationships between odor molecules in a manner conserved across individual mice. LII and LIII differ in terms of their correlation structure, reliability, odor sensitivity, and odor preferences, suggesting parallel representational strategies optimized for odor discrimination and generalization. These findings suggest that PCx harbors an invariant relational map of odorant space. The differential expression of this map across distinct cortical layers further suggests complementary contributions to perception and odor-guided behavior.

T-32. Neuronal computations underlying orientation change detection in the mouse visual cortex

Miaomiao Jin MIAOMIAO.JIN@DUKE.EDU
Jeff Beck JEFF.BECK@DUKE.EDU
Lindsey Glickfeld GLICKFELD@NEURO.DUKE.EDU
Duke University

Sensory information is encoded by populations of cortical neurons. Yet, it is unknown how this information is used for even simple perceptual tasks such as detecting a change in stimulus orientation. To infer which of the proposed decoders are used to perform this computation, we took advantage of the robust adaptation in the mouse visual system. We manipulated the adaptation state of visual cortical responses by varying the interstimulus interval in a change detection task. Following shorter intervals, in which there is stronger adaptation, the mouse had a higher threshold for detecting changes and a lower false alarm rate. Transient optogenetic

suppression of neurons in primary visual cortex (V1) impaired the animal's ability to perform the task, and also decreased the dependence on adaptation state. Thus, we used two-photon calcium imaging in V1 of passively viewing mice to measure the effect of adaptation on the responsivity and tuning of populations of neurons and used these data to determine how different decoding models would be impacted by adaptation state in a change detection task. We found that the majority of tested models (including optimal decoders, population vectors, and logistic regressions) uniformly predicted that adaptation decreases the threshold and increases the false alarm rate. The only model that predicted the same direction of effect as our behavioral data was a linear, sum decoder model. However, while the direction of the model matched the behavior, the performance was poor, suggesting that perhaps the computation was being performed on a weighted sum. Indeed, a linear combination of single trial neuronal responses, in which the weights were optimized, could successfully recapitulate the observed behavioral data. Thus, we propose a simple, linear combination of neuronal signals can quantitatively account for behavior on an orientation change detection task.

T-33. Causal discrimination of movement generation models using patterned microstimulation

Uday Jagadisan KJ.UDAYAKIRAN@GMAIL.COM
Neeraj Gandhi NEG8@PITT.EDU

University of Pittsburgh

Many brain regions involved in sensorimotor transformations multiplex sensory, cognitive, and movement-related information. How the brain extracts movement generation signals from such multiplexed representations is an unresolved puzzle. In the gaze control system, visuomovement neurons in the superior colliculus (SC), among other regions, are activated both by the onset of a visual stimulus in (visual burst) as well as a saccade to (premotor burst) their response field. These neurons also project directly to brainstem structures responsible for generating the final movement command, raising the question - why does the high-frequency visual burst not trigger a saccade? Extant models posit threshold-based gating or low-dimensional readouts of population activity as solutions to this demuxing problem. We recently showed, using pseudo-population analyses, that the temporal patterning of the visual burst is unstable compared to the premotor burst [1], suggesting that population temporal stability may be a factor guiding movement initiation. Here, we test these alternative models using causal manipulations. We first verified that the temporal stability hypothesis holds on individual trials by using linear microelectrode arrays to record SC population activity. The visual burst was unstable while the premotor burst was more stable compared to baseline, even for mean-matched bursts. Additionally, a linear decoder operating on population activity was able to discriminate between the two bursts. We then explicitly tested these alternative mechanisms by applying patterned microstimulation simultaneously across multiple electrode sites. Stimulation patterns were designed to be either stable or unstable on different trials, with otherwise matched pulse rates and "population states". Stable patterns were more likely to evoke saccades, and at lower latencies, compared to rate-matched unstable patterns. Crucially, a linear decoding mechanism was insufficient to explain the differences in stimulation outcomes. This provides a causal demonstration that the temporal structure of instantaneous population activity is a key variable determining movement initiation.

T-34. Assessing the scalability of biologically-motivated deep learning algorithms and architectures

Sergey Bartunov[1] SBOS.NET@GMAIL.COM
Blake Richards[2] BLAKE.RICHARDS@UTORONTO.CA
Adam Santoro[1] ADAMSANTORO@GOOGLE.COM
Geoffrey Hinton[3] GEOFFHINTON@GOOGLE.COM
Timothy Lillicrap[1] COUNTZERO@GOOGLE.COM

[1]DeepMind
[2]University of Toronto
[3]Google

The backpropagation of error algorithm (BP) is often said to be impossible to implement in a real brain. The recent success of deep networks in machine learning and the similarity between deep network and neocortical representations, however, has inspired a number of proposals for understanding how the brain might learn across multiple layers, and hence, how it might implement or approximate BP. As of yet, none of these proposals have been rigorously evaluated on tasks where BP-guided deep learning has proved critical, or in architectures more structured than simple fully-connected networks. Here we present the first results on scaling up a biologically motivated model of deep learning to datasets which need deep networks with appropriate architectures to achieve good performance. For CIFAR-10 we show that our algorithm, a straightforward, weight-transport-free variant of difference target-propagation (DTP) is competitive with BP in training deep networks with locally defined receptive fields that have untied weights. For ImageNet we find that both DTP and our algorithm perform significantly worse than BP, opening questions about whether different architectures or algorithms are required to scale these approaches. Our results and implementation details help establish baselines for biologically motivated deep learning schemes going forward.

T-35. Propagation of spike timing and firing rate in feedforward networks reconstituted in vitro

Jeremie Barral JEREMIE.BARRAL@GMAIL.COM
Xiao-Jing Wang XJWANG@NYU.EDU
Alex Reyes REYES@CNS.NYU.EDU

New York University

The neural code of transformation and propagation of information across brain regions is yet unclear. Theoretical analyses of idealized feedforward networks suggest that several conditions have to be satisfied in order for activity to propagate faithfully across layers. In real networks, neurons across layers often form bidirectional connections, which could facilitate or hamper signal propagation. Verifying these concepts in biological networks has not been possible owing to the vast number of variables that must be controlled. Here, we studied experimentally how spikes propagated in modular networks. We cultured cortical neurons in a chamber with sequentially connected compartments, optogenetically stimulated individual neurons in the first layer with high spatiotemporal resolution, and monitored the subthreshold and suprathreshold potentials in subsequent layers. In sparse networks where synaptic connections are strong, spikes propagated through the network and remained synchronous when the input packet delivered to the first layer was sufficiently narrow and large, as predicted by theory. However, in dense networks where connections are weak, spikes propagated reliably but were temporally more dispersed and delayed, which contradicts theoretical predictions. The input/output firing rate relation was linear for small range suggesting that these networks could transmit firing rate information. In the first layer, a brief stimulus with different temporal precisions resulted in the modulation of the firing rate. This temporal to rate transformation was propagated to other layers as a sustained response. This novel mode of propagation occurred in the balanced excitatory-inhibitory regime and depends critically on NMDA-mediated synapses activated by recurrent activity. These results indicate that both temporal and rate information can be propagated in feedforward networks, de-

pending not only on the input pulse packet but also on network architecture and density.

T-36. Orbitofrontal and parietal contributions to economic decisions in rats

Christine Constantinople[1] CMC9@PRINCETON.EDU
Charles Kopec[1] CKOPEC@PRINCETON.EDU
Carlos Brody[2,1] BRODY@PRINCETON.EDU

[1]Princeton University
[2]Howard Hughes Medical Institute

Humans evaluate economic options based on their utility, or the subjective satisfaction those options will provide. Additionally, humans exhibit probability distortion: we tend to overestimate the probability of unlikely events and underestimate the probability of likely events. Utility and probability distortion account for many aspects of human choice behavior, but their neural basis is not understood. We have developed a high-throughput behavioral training approach for studying economic decision-making in rats, which reveals utility and probability weighting functions. Rats are presented with light flashes from the left and right side; these flashes convey the probability of receiving water reward at each port. Simultaneously presented auditory click rates convey the volume of water reward baited at each port. Rats thus make a choice between explicitly cued probabilistic gambles. Using a range of both reward probabilities and volumes allows evaluating utility and probability weighting functions. High-throughput training has yielded tens of thousands of choices per rat in 24 well-trained rats. Subjects consistently maximize rewards and are risk averse. They also exhibit probability distortion, overweighting low probabilities. Optogenetic perturbations of parietal and orbitofrontal cortices affect utility and probability distortion, biasing rats towards choosing the risky option. Electrophysiological recordings from these regions reveal encoding of aspects of rats' subjective value, including probability distortion. These experiments represent, to our knowledge, the first time that utility and probability distortion have been evaluated in rodents, enabling use of powerful tools to manipulate and monitor neural circuits underlying these phenomena.

T-37. Widespread cortical involvement in evidence-based navigation

Lucas Pinto[1] LPINTO@PRINCETON.EDU
David Tank[1] DWTANK@PRINCETON.EDU
Carlos Brody[2,1] BRODY@PRINCETON.EDU
Stephan Thiberge[1] THIBERGE@PRINCETON.EDU

[1]Princeton University
[2]Howard Hughes Medical Institute

The gradual accumulation of noisy sensory evidence for perceptual decisions is a key cognitive phenomenon, but much remains unknown regarding its underlying neural mechanisms. A particular puzzle has been the observation of similar response properties of neurons from many different brain areas [1'4]. Whether this reflects a distributed code requiring the engagement of multiple areas, or only a subset of these areas are in fact required for perceptual decision-making, remains to be determined. To address this, we trained head-fixed mice on a virtual navigation-based visual accumulation task with pulsed stimuli, which they solved by indeed accumulating evidence from throughout the maze. We cleared their skulls, allowing optical access to the entire dorsal cortical surface for systematic and comprehensive optogenetic inactivation [5]. Bilateral inactivation of 26/29 cortical patches including sensory, motor and association cortex resulted in significant decreases in performance, observed in ChR2+ mice but not in controls, albeit with some differences in the nature and timing of the behavioral deficit between areas. In stark contrast, in a control task not requiring evidence accumulation (a maze where reward location was indicated by a distal visual guide but otherwise identical), only the visual cortex and the medial portion of the premotor cortex were required. In agreement with these findings, large-scale widefield Ca2+

imaging revealed widespread activation of the entire dorsal surface of the cortex. Moreover, we observed spatially distributed weights of decoding models for task variables, and an encoding model of neural activity revealed mostly quantitative differences between different cortical regions. Together, our results suggest that evidence accumulation and decision-making recruit large regions of cortex, but that specific cortical areas contribute to distinct yet highly overlapping aspects of the computation. Thus, evidence accumulation during navigation seems to be carried out by distributed cortical circuits.

T-38. Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation

Malcolm Campbell MALCGCAMP@GMAIL.COM
Samuel Ocko SAMOCKO@GMAIL.COM
Caitlin Mallory SC8LIN@GMAIL.COM
Surya Ganguli SGANGULI@STANFORD.EDU
Lisa Giocomo GIOCOMO@GMAIL.COM

Stanford University

To guide navigation, the nervous system must integrate multisensory self-motion and landmark information. We examined how these inputs generate the representation of self-location by recording medial entorhinal grid, border and speed cells in mice navigating virtual environments. By manipulating the gain between animal locomotion and the rate of visual scene progression, we could ascertain the differential contribution of self-motion and visual information to different cell-type responses. We found that border cells responded to visual landmark cues alone while grid and speed cells responded to combinations of locomotion, optic flow, and visual landmark cues in a context-dependent manner. A network model involving a path integrator network with input from landmark cells explained these results, providing principled, quantitative regimes under which grid cells either remain coherent with or break away from the landmark reference frame depending on the degree to which self-motion and landmark cues disagree. Moreover, during path integration-based navigation, we demonstrated behaviorally that mice estimated their position according to the principles predicted by our electrophysiological recordings. Together, our combined theory and experiments provide a quantitative framework for understanding how the integration of landmark and self-motion cues during navigation generates internal spatial representations, both at the level of entorhinal response properties and behavior.

T-39. A flexible model of working memory

Flora Bouchacourt FLORAMB@PRINCETON.EDU
Tim Buschman TBUSCHMA@PRINCETON.EDU

Princeton University

Working memory is fundamental to complex cognition, providing the workspace on which thoughts are held and manipulated. A defining characteristic of working memory is its flexibility: we can hold anything in mind. However, typical models of working memory rely on tightly tuned attractors to maintain memories and so they do not allow for the flexibility observed in behavior. Here we propose a novel model of working memory that maintains representations in the recurrent interactions between two layers of neurons: one composed of selectively tuned neurons and a second, randomly connected, untuned layer. Importantly, due to the parameter-free nature of these interactions, the model is able to maintain any type of incoming information, capturing the flexibility of working memory. Furthermore, despite being completely untuned, several emergent properties of the model capture important experimental observations of working memory. First, our model has a limited capacity, a behavioral hallmark of working memory. Second, neural representations are distributed across the network, similar to the distributed selectivity seen in humans and animals. Third, neurons in the untuned network show high-dimensional,

“mixed” selectivity that has been observed in prefrontal cortex. Fourth, increasing the number of items in memory causes a divisive-normalization-like reduction in neural selectivity. Finally, although neural activity in the model is dynamic, mnemonic representations are separable within a stable subspace, consistent with previous work. In summary, we present a simple, parameter-free, network model that uniquely allows for flexible representations while still capturing key behavioral and neural characteristics of working memory.

T-40. Motor preparation through rebound in an identified sensory integrator

Maarten Zwart[1] ZWARTM@JANELIA.HHMI.ORG
En Yang[1] YANGE@JANELIA.HHMI.ORG
Ziqian Wei[1] WEIZ@JANELIA.HHMI.ORG
Nikita Vladimirov[2] NIKITA.VLADIMIROV@MDC-BERLIN.DE
Sujatha Narayan[1] NARAYANS@JANELIA.HHMI.ORG
Minoru Koyama[1] KOYAMAM@JANELIA.HHMI.ORG
Ahmed Abdelfattah[1] ABDELFATTAHA@JANELIA.HHMI.ORG
Jonathan Grimm[1] GRIMMJ@JANELIA.HHMI.ORG
Luke Lavis[1] LAVISL@JANELIA.HHMI.ORG
Eric Schreiter[1] SCHREITERE@JANELIA.HHMI.ORG
Takashi Kawashima[1] KAWASHIMAT@JANELIA.HHMI.ORG
Shin-ichi Higashijima[3] SHIGASHI@NIPS.AC.JP
Shaul Druckmann[4] SHAULD@STANFORD.EDU
Misha Ahrens[1] AHRENSM@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]Max-Delbruck-Centrum
[3]National Institute for Basic Biology
[4]Stanford University

While some actions are triggered directly by instantaneous incoming sensory information, most take into account internal information streams that represent information accumulated over longer time periods. Despite decades of study, much is still unknown at the level of neural circuits about how the brain integrates sensory information in between actions, stores this information, and transmits it to inform future actions. We made use of modern whole-brain microscopy methods and the advantages of genetically modified zebrafish to search the entire brain, cell-by-cell, during behavior, for neural integrators involved in motor preparation. Larval zebrafish respond to visual motion, but swim in discrete swim bouts. Accordingly, visual motion must be integrated in between swim bouts to influence future motor output. Our whole-brain activity screen identified a single brain region in the hindbrain that responds to visual flow, integrates and stores it in ongoing activity, and modulates future behavior. These neurons respond to stimuli that encourage the fish to swim, and integrate them so that temporally separated stimuli have an additive effect. Moreover, their activity levels are positively correlated with preceding levels of motor output. However, surprisingly, stimulating the neurons suppressed instantaneous swimming, yet advanced the onset time of future swims. Network models based on these results suggests a brainstem integrator network that, in between actions, stores visual information while actively suppressing motor output, and during actions transmits the stored information to premotor circuits through post-inhibitory rebound. Voltage imaging using novel voltage indicators revealed precise temporal dynamics of this population of neurons in relation to behavior consistent with this model. These findings suggest inhibitory integrators as a complementary alternative to integrate-to-bound models for sensory integration and motor decisions.

T-41. A cross-species analysis of accumulation of evidence under non-sensory uncertainty and its modulation by the prefrontal cortex

Pietro Vertechi[1] PIETRO.VERTECHI@GMAIL.COM
Eran Lottem[1] ERAN.LOTTEM@NEURO.FCHAMPALIMAUD.ORG
Dario Sarra[1] DARIO.SARRA@NEURO.FCHAMPALIMAUD.ORG
Beatriz Godinho[1,2] BEATRIZ.GODINHO@NEURO.FCHAMPALIMAUD.ORG
Isaac Treves[3] ITREVES@PRINCETON.EDU
Tiago Quendera[1] TIAGO.QUENDERA@NEURO.FCHAMPALIMAUD.ORG
Matthijs Oude Lohuis[1] MATTHIJS.OUDELOHUIS@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[1] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Champalimaud Neuroscience Programme
[3]Princeton University

Essential features of the world are often hidden and must be inferred by constructing internal models through indirect and uncertain evidence. To investigate this, we designed a probabilistic foraging task for parallel use in mice and humans. Subjects sought rewards (water or points, respectively), by actively probing a foraging site (nose-poking or screen-tapping, respectively, Fig1A). Each try yielded reward with probability pREW, and could cause the site to deplete with probability pDPL. This required them to travel to a second, fresh, site at some distance and bear a travel cost (cSwitch). Subjects were therefore tasked with inferring a hidden state (fresh or depleted) through a stochastic sequence of observations (rewards). Optimal agents should accrue evidence for depletion by counting consecutive misses, with the accumulation rate set by the statistics of the environment (pREW, pDPL, cSwitch). We found that both the rodents' and humans' actual switching times matched closely those predictions (Fig1). Moreover, a two-parameter Poisson counting model could parsimoniously fit the data. Importantly, this model, unlike other evidence accumulation models, captured the scalar invariance of switching time distributions over environmental statistics (Fig2). Finally, to study mechanism, we tested the causal contribution of prefrontal cortex using optogenetic stimulation of inhibitory interneurons in VGAT-ChR2 mice. We found that transient inactivation of anterior cingulate cortex (ACC) divisively modulated the rate of evidence accumulation, increasing tolerance of misses. In contrast, inactivation of the orbitofrontal cortex (OFC) disrupted the pattern of switching times, causing mice to exhibit more "model-free" behavior (Fig3C). This pattern of results indicate that while both areas are involved in this task, OFC is crucial in the inference process, whereas ACC may be a downstream node. Together, these results suggest that highly parallel cross-species analyses in mice and humans are a viable approach to study the neural bases of complex behavior.

T-42. Individual animals use distinct strategies in a changing spatial memory task

David Kastner[1] DBKASTNER@GMAIL.COM
Anna Gillespie[1] ANNA.GILLESPIE@UCSF.EDU
Peter Dayan[2] DAYAN@GATSBY.UCL.AC.UK
Loren Frank[1] LOREN@PHY.UCSF.EDU

[1]University of California, San Francisco
[2]Gatsby Computational Neuroscience Unit

Behavior is typically highly variable across individuals. To deal with this complexity, average data from many animals, grouped by genotype or experience, is analyzed for differences. This approach, while powerful, obscures the individual differences within groups. To link the workings of individual brains to their behavioral outputs, we have to identify and quantify inter-individual behavioral differences. Therefore, we designed an automated six-arm behavioral apparatus to minimize and control extraneous variables. Initially, wild type (WT) and fragile-X syndrome (FXS) model Long-Evans rats received reward on all arms, allowing us to examine their innate biases.

Subsequently, reward was restricted to just three arms, and was only delivered if the rats visited the arms in a specific sequence. Four different sets of rewarded arms were employed in sequential blocks. We then developed a generative model for the entire behavior, using a variant of a reinforcement learning (RL) model with short-term memory. Conventional RL implementations learned markedly slower than the animals. The addition of two types of arm biases that were apparent during the exploratory behavior produced a model with three free parameters–learning, forgetting, and temporal discounting rates–that could be fit to describe the behavior of each animal. A 2D subspace of the parameters identified systematic variation of parameters between individuals, such as where a group of animals had similar parameters for their learning rates, but varied in their temporal discounting rates. These analyses revealed that individual differences far outweighed genotypic differences across the full cohort of FXS and WT animals, but that when subsets of animals with similar sets of model parameters were identified, genotype-related differences emerged. These findings suggest that relatively simple behavioral models can explain individual differences and thereby reveal otherwise hidden patterns of variability across different groups of animals.

I-1. Organic learning: Spiking neural computation with topographic connection patterns

Carl Gold CARL.STEVEN.GOLD@GMAIL.COM

Independent Researcher

In this modeling study, simple patterns of local connections between neurons embedded in topographically arranged layers provide built-in object recognition capability to a spiking neural network. The strategy is inspired by recent detailed examinations of how connections between specific cell types and at specific dendritic locations implement computational functions in motion detection. For object recognition, straightforward structures create a sparse population code of topographically organized feature detectors. The main result of this study is that having such a structure allows a spiking neural network to learn to recognize objects rapidly and without assuming backpropagation of information in the network. Three dimensional structures that implement specific response patterns are illustrated, showing that for topographically organized neurons, certain patterns of connections are equivalent to computation. The resulting population code is easy to interpret, like the responses of visual neurons recorded in vivo, and the network is endowed with responsiveness to appropriate stimuli from inception rather than being initially random in its responses. Because of these characteristics, such structures may serve as means of enabling object recognition in visual systems with most of the necessary information transferred in the genome, and only limited task specific learning after an individual organism is born. This breaks with the standard of statistical learning theory but reflects the observation that most organic brains must compete to survive immediately upon birth, with little chance to learn. A secondary topic of the study is that the network design relies on the equivalence between certain classes of binary and spiking neural network models, and the biophysical limits on this equivalence is explored. This method of analyzing spiking neural networks may be applicable to a broad class of networks.

I-2. Causal roles of basal ganglia circuitry in regulating response criterion during visual selective attention

Lupeng Wang LUPENG.WANG@GMAIL.COM
Richard Krauzlis RICHARD.KRAUZLIS@NIH.GOV

National Eye Institute

Visual selective attention has profound behavioral effects on detection performance, including changes in the subjects' perceptual sensitivity and decision criterion, two distinct aspects defined in signal detection theory. What are the neural circuit mechanisms underlying attentional shifts in criterion and sensitivity? It has been shown that neuronal activity in visual area V4 is correlated with shifts in sensitivity (Luo&Maunsell, 2014) but not criterion. Here we present evidence that the Basal Ganglia (BG) are involved in shifts of criterion but not sensitivity. The BG are important for action selection. Recently, they have also been implicated in regulating decision variables, and BG pathology in humans often causes perceptual and attentional deficits. However, the contribution of BG to the control of attention is unclear. Here we tested the causal role of the BG in visual selective attention by optogenetically manipulating the activity of genetically defined neuronal populations in the dorsomedial striatum of mice during a spatially cued attention-to-orientation task. Mice displayed robust attentional effects, including spatially specific changes in sensitivity, criterion and reaction time. Brief unilateral optogenetic stimulation caused large changes in criterion and reaction time, without significant effects on sensitivity. When a spatial prior was provided (i.e., cued trials), the effects of stimulation displayed a spatial asymmetry: striatal direct pathway stimulation caused significantly larger decreases in criterion and reaction time when the visual change was expected in the contralateral visual field. In contrast, striatal indirect pathway stimulation caused larger increases in reaction time to the contralaterally cued visual change. In the absence of a spatial prior (i.e., uncued trials), these spatial asymmetries were no longer observed. Our results indicate that the BG play a causal role in regulating attention-related shifts in the criterion used for visual detection, and are part of the circuits that implement cue-based spatial priors during visual selective attention.

I-3. Functional synaptic architecture of callosal inputs in mouse primary visual cortex

Kuo-Sheng Lee LEEK2013@FAU.EDU
Kaeli Vandemark KVANDEMA@TULANE.EDU
David Mezey MEZEYDAVID@GMAIL.COM
Nicole Shultz NICOLE.SHULTZ@MPFI.ORG
David Fitzpatrick DAVID.FITZPATRICK@MPFI.ORG

Max Planck Florida Institute

Integrating the neural signals between two cerebral hemispheres is crucial for many cognitive functions in placental mammals. For example, there is growing evidence that callosal projections connecting visual cortex of both hemispheres play a significant role in binocular vision in rodents. However, the manner in which synaptic contacts are arranged within the dendritic field of individual neurons remains unknown. Here we describe a high-throughput method to map the functional organization of callosal connectivity by combining in vivo 3D random-access two-photon calcium imaging of the postsynaptic dendrites of single V1 neuron and optogenetic stimulation of the presynaptic neural population in the contralateral hemisphere. We find that a group of dendritic spines can be reliably activated by the presynaptic optogenetic stimulation, which we define as callosal recipient spines. Callosal recipient spines exhibit ipsilateral eye dominance, and stronger binocular matching of orientation preference and alignment to the somatic orientation preference, compared with the dendritic spines not driven by optogenetic stimulation (non-callosal recipient spines). Interestingly, callosal recipient spines are not homogeneously distributed in the dendritic arbor, but biased toward specific branches. Most importantly, we find that callosal recipient spines are more likely to cluster with the non-callosal recipient spines with similar orientation preference. At the end, we apply post hoc expansion microscopy to the tissue to directly visualize monosynaptic callosal connections. Our results demonstrate, for the first time, that functional synaptic clustering in a short dendritic segment could play a role in integrating two distinct neuronal circuits. This strategy could be ideal to non-linearly amplify the specific matched signals in the visual scene coming from two eyes, and thus facilitate the binocular vision.

I-4. From robustness to richness: neural code across species and regions

Raviv Pryluk[1] RAVIVPRYLUK@GMAIL.COM
Hagar Gelbard-Sagiv[2] HAGAR.SAGIV@GMAIL.COM
Yoav Kfir[1] YOAVKF@GMAIL.COM
Itzhak Fried[3] IFRIED@MEDNET.UCLA.EDU
Rony Paz[1] RONY.PAZ@WEIZMANN.AC.IL

[1]Weizmann Institute of Science
[2]Tel-Aviv University
[3]University of California, Los Angeles

Millions of evolutionary years separate humans and macaques from their last common ancestor, and although key brain regions exist in both species with large similarities, there is an evident gap in functional abilities. The mechanisms that underlie such differences are likely due to architecture and the neural code, yet this hypothesis remains untested. To address this, we recorded single-unit activity in the amygdala and the prefrontal (cingulate) cortex of behaving humans and monkeys. These two regions evolved extensively to underlie emotional learning and decision making; abnormal interactions between them underlie neuropsychiatric conditions. Using information-theoretic approaches, we compared neural representations in each region of each primate. We find that the entropy of human neurons is higher than in monkeys, taking into account changes in firing rates, and this was further the case when comparing the cortex to the amygdala within each specie. We then characterize entropy in simultaneously recorded pairs of neurons, and again find higher entropy in humans and in the cortex. Therefore, neurons enable richer code in humans and in the cortex. In contrast, the amount of significant pairwise correlation was higher in monkeys and in the amygdala, and we find more overlaps in the used vocabulary by comparing the distribution of 'words' across pairs of neurons directly with Jensen-Shannon-Divergence (JSD). Therefore, neurons enable more robust transmission of information in monkeys and in the amygdala. Our findings suggest that there is a tradeoff between robustness, i.e. using the same words; and richness, i.e. larger vocabulary of words. Changes in this balance across species and brain regions suggest fundamental differences in the neural code that can underlie the observed differences in function. We posit that such cross-species investigations are crucial for understanding basic features of the neural code and are essential for translational psychiatry.

I-5. Perirhinal cortical feedback activity modulates associative memory formation in the rat's somatosensory cortex

Guy Doron GUY.DORON@CHARITE.DE
Matthew Larkum MATTHEW.LARKUM@GMAIL.COM

Humboldt University Berlin

The hippocampus and related parahippocampal structures (entorhinal cortex [EC], perirhinal cortex [PRh], etc.) play a vital role in transforming experience into long-term memories that are then stored in the cortex, however the cellular mechanisms which designate single neurons to be part of a memory trace remain unknown. In order to examine associative memory formation we developed a behavioral paradigm using cortical microstimulation in which a rat is trained to report the direct stimulation of the somatosensory cortex (S1). Importantly, rats learn to associate the stimulus with a reward within a single session and improve over a three-day period, therefore offering an unprecedented method for locating the mechanisms of memory formation in both space and time. In order to confirm the connectivity between the PRh and S1 we expressed ChR2-EYFP via a viral vector (AAV) in the PRh and found a pronounced expression in layer 1 (L1) of S1, providing an anatomical and functional evidence of such feedback projection. Furthermore, we conducted double-blind experiments using the microstimulation paradigm and found that this learning is hippocampal and PRh dependent. Moreover, juxtacellular recordings in the PRh in awake head-restrained rats during learning showed that activity in the deep layers is increased around shortly (up to few seconds) after the microstimulation in S1 (Figure 1). Spectral analysis of the local field potential (LFP) 1 sec before and 1-4 seconds after the microstimulation revealed an increase in the theta frequency in PRh

and S1 during this time window, suggesting a mechanism for information transfer between these regions during memory formation. We are currently examining the dendritic mechanisms in L1, which allow information flow from PRh projections to L5 pyramidal neurons tuft dendrites. Overall, our data are consistent with the perirhinal cortex influencing S1 during memory formation and consolidation.

I-6. Ventral basal ganglia sends performance error signals to VTA in singing birds

Ruidong Chen RC694@CORNELL.EDU
Pavel Puzerey PAP99@CORNELL.EDU
Archana Podury AP675@CORNELL.EDU
Kamal Maher KAMALKALOID@GMAIL.COM
Jesse Goldberg JESSEHGOLDBERG@GMAIL.COM

Cornell University

Motor performance is evaluated against personal benchmarks that change with learning. Evaluating your tennis forehand relative to your past forehands is more useful than comparing it to swinging like Federer. Zebra finches learn to sing by imitating a tutor song, suggesting they have a 'target' they aspire to learn. Yet song syllables are not simply evaluated against a fixed target. Instead, recent recordings suggest that syllables are evaluated against syllable-specific performance benchmarks updated during practice. Specifically, ventral tegmental area (VTA) dopamine neurons exhibited phasic suppressions following distorted auditory feedback (DAF) during singing, consistent with a worse-than-predicted outcome. They also exhibited phasic bursts at the precise time-step of a syllable when a predicted distortion did not occur. Burst magnitude depended on distortion history, consistent with an error signal scaled by the predicted syllable quality. Upstream circuits that compute this error signal are unknown. Here we combine lesion, electrophysiology, DAF, and viral tract tracing, to identify the VTA-projecting part of the ventral basal ganglia (vBGvta) as a major hub for error processing. Juvenile birds with excitotoxic lesion to vBG failed to imitate tutor song compared to sham-lesioned siblings. Distinct subtypes of antidromically identified vBGvta neurons encode auditory error and predicted syllable quality. Other vBG cell types encode precise song timing, auditory error, and singing state. Using viral tracing we identified novel projections to vBG from: (1) the HVC-projecting part of the motor thalamus (Uva), a source of precise song timing information; (2) the VTA-projecting part of auditory cortex (AIV), a source of fast auditory error; and (3) the Area X projecting part of VTA (VTAx), a source of modulatory prediction error. We present a simple model in which vBG microcircuits integrate these three inputs to compute a syllable-specific performance benchmark, dependent on error history, against which auditory feedback is compared during singing.

I-7. The role of untuned neurons in sensory information coding

Joel Zylberberg JOEL.ZYLBERBERG@GMAIL.COM

University of Colorado School of Medicine

In the sensory systems, most neurons' firing rates are tuned to at least one aspect of the stimulus. Other neurons appear to be untuned, meaning that their firing rates do not depend on the stimulus. Previous work on information coding in neural populations has ignored untuned neurons, based on the tacit assumption that they are unimportant. Recent experimental work has questioned this assumption, showing that in some circumstances, neurons with no apparent stimulus tuning can contribute to sensory information coding. These findings are intriguing, because they suggest that – by virtue of our ignoring putatively untuned neurons – our understanding of neural population coding might be incomplete. At the same time, several key questions remain unanswered: Are the impacts of putatively untuned neurons on population coding due to weak tuning that is nevertheless below the threshold the experimenters set for calling neurons tuned (vs untuned)? And why do there appear to be

untuned neurons in the brain? Do mixed populations of tuned and untuned neurons have a functional advantage over populations containing only tuned neurons? Using theoretical calculations and analyses of in vivo neural data, I answer those questions by: a) showing how untuned neurons can enhance sensory information coding; b) demonstrating that this effect does not rely on weak tuning; and c) identifying conditions under which the neural code can be made more informative by replacing some of the tuned neurons with untuned ones. These conditions specify when there is a functional benefit to having untuned neurons in a circuit, and thus suggest a reason why the brain might contain untuned neurons. This work shows that, even in the extreme case, where some neurons have no tuning, those neurons can still contribute to sensory information coding, and thus should not be ignored.

I-8. Learning of valence-discrimination with temporal sequences in the primate amygdala

Tamar Stolero TAMAR.STOLERO@GMAIL.COM
Aryeh Taub ARYEH.TAUB@WEIZMANN.AC.IL
Rony Paz RONY.PAZ@WEIZMANN.AC.IL

Weizmann Institute of Science

Affective learning and memory formation are crucial for daily behavior and may lead to psychiatric disorders in extreme maladaptive situations. These processes are known to involve the amygdala and prefrontal-cortex. Studies in this field traditionally focus on stimulus-driven single-unit activity and its modulation. In contrast, the role of coding with temporal sequences and how they convey information during periods with no external stimulus remains unclear. Here, we focus on spike-sequences in triplets of neurons recorded while monkeys engaged in a conditioning task that required learning to discriminate pleasant from aversive associations. Importantly, we examined spike sequences during baseline period and during long periods in the inter-trial-interval (ITI), long after the stimulus evoked response. Our results suggest that triplets of neurons in the amygdala exhibit consistent sequence activity throughout time, that differs from the sequences expected from single neurons firing patterns. Additionally, amygdala sequences patterns during the ITI were valence specific, and correspondingly, decoding of stimulus valence from the ITI sequence activity was successful. Moreover, sequences based decoding reached higher accuracy than when using firing rates and inter spike interval information. Finally, behavioral discrimination performance was negatively correlated to the sequence-based decoding during ITI, indicating a teaching-signal dynamics. Compared to amygdala, prefrontal neurons were less prone to form sequences and carried less information about valence. We suggest that multi-unit temporal sequences serve as coding mechanisms in the primate amygdala that strengthen and rehearse recent associations to aid affective memory formation. Our work highlights the importance of studying learning related neural processing that is not stimulus evoked as well as the role of spike sequences across neurons in such processing.

I-9. Space representation in the goldfish brain

Ehud Vinepinsky EHPINSKY@GMAIL.COM
Ohad Ben-Shahar BEN-SHAHAR@CS.BGU.AC.IL
Opher Donchin DONCHIN@BGU.AC.IL
Ronen Segev RONENSGV@BGU.AC.IL

Ben-Gurion University of the Negev

Almost all animals engage in some form of navigation. The neural navigation system of mammals is believed to be based on place cells, grid cells, head direction cells and border cells. All of those cells types are found in the hippocampal formation. However, it remains an open question whether similar building blocks drive the navigation system of other animal classes. To address this fundamental issue, we use goldfish as a model animal. Goldfish have been shown to be able to navigate using allocentric and egocentric cues. Furthermore,

there is a defined neuroanatomical region associated with allocentric navigation, the lateral pallium. Using a novel wireless recording system, we measured the activity of single cells in the lateral pallium while fish swam in a longitudinal aquarium. We found three unique cell types: border cells, velocity cells and speed cells. Border cells are cells which are active when the fish is near the boundary of the environment. Velocity cells encode swimming direction and speed, while speed cells encode only the speed independent of direction. Our study sheds light on how spatial information is encoded in the fish brain and whether the mechanisms of the neural navigation system are preserved across evolution.

I-10. Emergence of grid-like representations by training neural networks to perform spatial localization

Christopher Cueva CCUEVA@GMAIL.COM
Xue-Xin Wei WEIXXPKU@GMAIL.COM

Columbia University

Decades of research on the neural code underlying spatial navigation have revealed a diverse set of neural response properties. The Entorhinal Cortex (EC) of the mammalian brain contains a rich set of spatial correlates, including grid cells which encode space using tessellating patterns. However, the mechanisms and functional significance of these spatial representations remain largely mysterious. As a new way to understand these neural representations, we trained recurrent neural networks (RNNs) to perform navigation tasks in 2D arenas based on velocity inputs. Surprisingly, we find that grid-like spatial response patterns emerge in trained networks, along with units that exhibit other spatial correlates, including border cells and band-like cells. All these different functional types of neurons have been observed experimentally. The order of the emergence of grid-like and border cells is also consistent with observations from developmental studies. Importantly, the network accurately performs spatial localization for path lengths that are orders of magnitude longer than those used during training, despite the expected accumulation of errors due to noise in the network. We find that the network reduces the errors through boundary interactions, in agreement with previous experimental results. The model proposed here develops the grid-like responses, boundary responses and the error-correction mechanisms all within the same neural network, thus potentially providing a unifying account of a diverse set of phenomena. Together, our results suggest that grid cells, border cells and others as observed in EC may be a natural solution for representing space efficiently given the predominant recurrent connections in the neural circuits.

I-11. A cerebellar circuit role in evidence accumulation and decision-making

Ben Deverett DEVERETT@PRINCETON.EDU
Sue Ann Koay KOAY@PRINCETON.EDU
Samuel Wang SSWANG@PRINCETON.EDU

Princeton University

Evidence-accumulation paradigms have been successfully used to probe neural correlates of decision-making in rodents. Numerous brain structures have been proposed to support this complex sensorimotor process, but the cerebellum has been largely unexplored. Aside from being a site of dense sensorimotor convergence, the cerebellum has been extensively associated with working memory (Kansal et al., 2017), a key component of the evidence-accumulation process. Moreover, crus I of the cerebellum is reciprocally connected with prefrontal and posterior parietal cortex (Strick, Dum, & Fiez, 2009), regions known to encode evidence and decision variables in primates and rodents. We set out to investigate cerebellar contributions to evidence accumulation in mice, focusing on crus I.

We designed an evidence-accumulation task in which puffs of air were delivered to the left and right whiskers, and mice licked in the direction with more puffs to retrieve a reward. Neuronal inactivations in the cerebellum significantly impaired performance, producing directional choice biases that were not explained by impairments in licking ability. Calcium imaging in Purkinje cells revealed prominent ramp-like representations of choice and evidence variables similar to those seen in neocortical regions during decision-making. In the dendrites of Purkinje cells, which report a separate class of calcium spikes known to induce plasticity in the cerebellar circuit, we observed increases in activity coincident with decision errors, a feature analogous to well-characterized cerebellar error signalling in the motor domain.

These findings provide novel evidence for longstanding hypotheses about cerebellar roles in complex cognitive processes. The similarity of signals in the cerebellum to those in complementary neocortical regions suggests a functional communication pathway through which information is routed to the cerebellum and processed by cerebellar circuitry to aid in decision-making processes. We propose that the cerebellum may modulate the trajectories of evidence accumulation and commitment to a decision, just as it modulates the accuracy of movement.

I-12. A mixture of sparse coding models for holistic and parts-based face processing in the IT cortex

Haruo Hosoya[1] HOSOYA@ATR.JP
Aapo Hyvarinen[2] A.HYVARINEN@UCL.AC.UK

[1]ATR International
[2]University College London

Experimental studies have revealed evidence of both parts-based and holistic representations of objects and faces in the primate visual system. However, it is computationally not obvious how such seemingly contradictory types of processing can coexist within a single system. Here, we propose a novel theory called mixture of sparse coding models, inspired by the formation of category-specific subregions in the inferotemporal (IT) cortex. We developed a hierarchical network that constructed a mixture of two sparse coding submodels on top of a simple Gabor analysis. The submodels were each trained with face or non-face object images, which resulted in separate representations of facial parts and object parts. Evoked neural activities were modeled by Bayesian inference, which had a top-down explaining-away effect that enabled recognition of an individual part to depend strongly on the category of the whole input. Notably, the resulting model explained, qualitatively and quantitatively, almost all response properties reported by Freiwald, Tsao, and Livingstone (2009) on the middle patch of face processing in IT. Namely, our model units exhibited (1) significant selectivity to face images over object images, (2) tuning to only a small number of facial features that were often related to geometrically large parts, (3) preference and anti-preference of extreme facial features (e.g., very large/small inter-eye distance), (4) reduction of the gain of feature tuning for partial face stimuli compared to whole face stimuli, and (5) similarity of feature tuning between inverted and normal face stimuli. Not all above properties could be reproduced with a simple sparse coding model trained with face images or a multi-layer perceptron trained to discriminate faces from objects. Thus, we hypothesize that the coding principle of facial features in the middle patch of face processing in the macaque IT cortex may be closely related to mixture of sparse coding models.

I-13. Stable memory with unstable synapses

Lee Susman LEE.SUSMAN@GMAIL.COM
Naama Brenner NBRENNER@TECHNION.AC.IL
Omri Barak OMRI.BARAK@GMAIL.COM

Technion - Israel Institute of Technology

What are the neural correlates of long term memory? The current dogma in theoretical neuroscience describes a picture in which synaptic efficacies encode for memory traces, by means of a Hebbian-like update rule that operates during memory acquisition. This view implies that in the absence of structured neural activity, synaptic

efficacies should be constant; in other words, stable memories correspond to stable connectivity patterns. However, an increasing body of experimental evidence points to significant, activity-independent dynamics of synaptic strengths [1]. These fluctuations seem to be occurring continuously, in parallel to directed plasticity, with effects on synaptic strengths of similar magnitude, and without specific reference to a learning process. Motivated by these observations, we explore an alternative hypothesis. We propose that network connectivity has global invariant features, despite the spontaneous plasticity of each individual synapse. As a concrete example, inspired by [2], we study a family of network models in which the symmetric part of the connectivity matrix is volatile, while memory items are stored in the anti-symmetric part. Stable limit cycles emerge as the neural implementation of memory traces, thus establishing a purely anti-symmetric analog of the Hopfield model [3]. Finally, we show that these representations can be learned in our framework via biologically plausible dynamic learning rules, similar in spirit to spike-timing-dependent plasticity.

References: [1] G. Mongillo, S. Rumpel, and Y. Loewenstein. Curr. Opin. Neurobiol., 2017; [2] M. O. Magnasco, O. Piro, and G. A. Cecchi. Phys. Rev. Lett., 2009; [3] J. J. Hopfield. Proc. Natl. Acad. Sci. USA, 1984

I-14. Self-organization of entorhinal grid modules through commensurate lattices

Louis Kang[1] LOUIS.KANG@BERKELEY.EDU
Vijay Balasubramanian[2] VIJAY@PHYSICS.UPENN.EDU

[1]University of California, Berkeley
[2]University of Pennsylvania

The grid system of the mammalian medial entorhinal cortex (mEC) exhibits striking modularity. Rat grid cell recordings reveal that spatial grid scales cluster around discrete values separated by constant ratios reported in the range 1.3–1.8. Although this modular organization has been shown to be a robust and efficient encoding of spatial location, its origin is unknown. We present the first proposed mechanism through which geometric sequences of grid scales arise naturally. A series of continuous attractor networks along the longitudinal mEC axis that would otherwise generate a smooth distribution of grid scales forms modules separated by discrete jumps in scale when excitatory connections are introduced. Moreover, constant scale ratios between successive modules arise through robust geometric relationships between commensurate triangular grids, whose lattice constants are separated by $\sqrt{3} \approx 1.7$ or other ratios, or between grids containing local lattice modulations called discommensurations. These relationships persist in single neuron spatial rate maps due to faithful path integration and are unaffected by perturbations to model parameters. We speculate on how excitatory connections between attractor networks can be realized by the known architecture of the mEC and suggest analyses and experiments that test our model.

I-15. The relationship between pairwise correlations and dimensionality reduction

Rudina Morina[1] RHMORINA@GMAIL.COM
Benjamin Cowley[1] BCOWLEY@CS.CMU.EDU
Akash Umakantha[1] AUMAKANT@ANDREW.CMU.EDU
Adam Snyder[1] ADAM@ADAMCSNYDER.COM
Matthew Smith[2] SMITHMA@PITT.EDU
Byron Yu[1] BYRONYU@CMU.EDU

[1]Carnegie Mellon University
[2]University of Pittsburgh

Spike count correlation (r_sc), also known as “noise correlation,” has been used extensively to characterize the interactions between pairs of neurons. With the advent of multi-electrode array recordings, dimensionality reduction techniques are also being increasingly used to study the interactions among many neurons simultaneously. Although both approaches utilize the spike count covariance matrix, the connection between the two has not been established. In this study, we explore how pairwise correlation metrics relate to the outputs of dimensionality reduction. Understanding this relationship would allow us to bridge across studies that report pairwise correlation metrics or dimensionality reduction metrics, but not both. We first used simulated data to systematically alter dimensionality reduction metrics and asked how pairwise correlation metrics were affected. For a given population of neurons, we found a systematic relationship between the distribution of spike count correlations (across all pairs of neurons) and the shared co-fluctuation patterns across the population (i.e., dimensions). This relationship depends on the number of patterns, as well as the strength of each pattern. For neurons with one dominant co-fluctuation pattern, the r_sc mean and r_sc standard deviation (std) indicate the extent to which all neurons increase and decrease their activity together. As the number of co-fluctuation patterns increases, both r_sc mean and r_sc std shrink. We used these simulations to understand how the previously-reported decrease in r_sc mean in macaque visual cortex (V4) during attention reflects a change in the population activity structure. These findings help to bridge studies that utilize the two different approaches for analyzing neural population activity.

I-16. Spatiochromatic integration by double opponent neurons in macaque V1

Abhishek De ABHISD@UW.EDU
Gregory Horwitz GHORWITZ@UW.EDU

University of Washington

The color of a light depends on surrounding lights. This effect is likely mediated, at least in part, by double opponent (DO) neurons in area V1. DO neurons have two characteristic properties: they are cone opponent and they have opposite color preferences in different parts of their spatial receptive field (RF). As a result, DO neurons respond maximally to color boundaries and weakly to full-field color stimuli. How these neurons integrate color signals across their RFs, however, is not well understood. For this reason, physiological and psychophysical spatial color processing are difficult to relate quantitatively. We identified V1 DO neurons in awake behaving monkeys using spike triggered averaging. We presented stimuli that activated non-overlapping regions of the RF individually or simultaneously. Using an adaptive closed-loop stimulus generator, we identified stimuli that drove the same neuronal response but differed in how strongly they activated two regions of the RF. We encountered two classes of DO neurons that were selective for either blue-yellow or red-green edges. Almost all blue-yellow and some red-green DO neurons responded to a weighted sum of color signals from the two non-overlapping regions of their RFs. Consequently, these neurons responded to chromatic contrast between the two regions of their RFs irrespective of the absolute chromaticities that defined the edge. For example, a blue-yellow DO neuron responded identically to a blue-yellow edge and to an edge between a saturated and an unsaturated blue (or yellow). A subset of red-green DO neurons combined color signals across their RFs nonlinearly. This nonlinearity may be due to complex interactions between cone opponent and cone non-opponent signals across space that have previously been identified with spike triggered covariance analysis.

I-17. Determining the role of VIP signaling in enhancement of adult visual cortex plasticity

Anna Lebedeva[1,2] ANNA.LEBEDEVA.17@UCL.AC.UK
Yujiao Jennifer Sun[3] J.SUNINCHINA@GMAIL.COM
Michael Stryker[3] STRYKER@PHY.UCSF.EDU

[1]University College London
[2]Sainsbury Wellcome Center
[3]University of California, San Francisco

The adult visual cortex has limited plasticity, which is reflected in decreased visual perceptual learning and incomplete rehabilitation after damage (Mitchell and Sengpiel, 2009). Therefore, it is critical to identify ways to enhance adult plasticity and elucidate its underlying neural mechanism. Recent studies demonstrated that locomotion promotes recovery of visual responses in adult mouse primary visual cortex (V1) from prolonged monocular deprivation (Kaneko and Stryker, 2014) through activation of vasoactive intestinal peptide (VIP) interneurons (Fu et al, 2015). Here, we examined how VIP interneurons drive the circuit to promote adult plasticity in visually deprived animals. By taking advantage of transgenic animal models and GCaMP6 labeling, we utilized two-photon imaging to examine two pathways that VIP neurons recruit, namely secretion of GABA and of VIP peptide. Our analysis of activity on a single cell level shows that the secretion of GABA is critical in promoting ocular dominance plasticity in adult mice.

I-18. Nonlinear filtering and learning for point emission processes

Anna Kutschireiter[1,2] ANNAK@INI.UZH.CH
Simone Carlo Surace[1,2] SURACE@INI.UZH.CH
Jean-Pascal Pfister[1,2] JPFISTER@INI.UZH.CH

[1]University of Zurich
[2]ETH Zurich

The number of neurons that can be simultaneously recorded doubles every seven years (Stevenson et al., 2011). This ever increasing number of recorded neurons opens the possibility to address new questions and extract relevant signals from high-dimensional recordings. However, existing algorithms designed to extract those features fall short when the dimensionality of the recordings is large. For example, classical particle filter algorithms that rely on importance weights are known to suffer from the curse of dimensionality (COD). Indeed, the number of particles required for a certain performance scales exponentially with the problem dimensionality. Here, we analytically and numerically investigate the reason for the COD of weighted particle filtering approaches: similarly to particle filtering with continuous-time observations, the COD with point-process observations is due to the decay of effective number of particles in the algorithm, an effect that is stronger when the number of observable dimensions, for instance the number of recorded neurons, increases. Given the success of unweighted particle filtering approaches in overcoming the COD for continuous-time observations, we propose an unweighted particle filter for point-process observations, the spiking Neural Particle Filter (sNPF), and show that it exhibits a similar favorable scaling as the number of dimensions grows. Further, we derive from a maximum likelihood approach a simple learning rule for the parameters of the sNPF, that allows both online and offline unsupervised learning of the model parameters.

I-19. Training continuous time spiking neural networks with back-propagation through spike times

David Sussillo[1] SUSSILLO@GOOGLE.COM
Chris Maddison[2] CMADDIS@GOOGLE.COM
Matthew Johnson[1] MATTJJ@GOOGLE.COM
Danny Tarlow[1] DTARLOW@GOOGLE.COM

[1]Google Brain
[2]DeepMind

A large and open problem in the field of computational neuroscience is the ability to build spiking neural network models (SNNs) that both reflect biological complexity and can be configured to perform tasks of interest. While a number of approaches exist for training spiking networks, the vast majority of these approaches do not include a gradient that is back-propagated through either time or structure, and those that do involve approximate gradients or modifying the model to be differentiable in some way. Here we propose a framework for exact simulation of SNNs that computes spikes in an event-based fashion in continuous time with spike times computed to numerical precision. We show how to compute exact derivatives of the spike times with respect to other variables, allowing errors arising from a global loss to propagate backwards through both structure and time. This enables gradient-descent based training of both network parameters such as synaptic weights, or neuronal parameters such as time constants and refractory periods. We demonstrate our method on popular tasks in the literature.

I-20. Synaptic mechanisms of interference in parametric working memory

Zachary Kilpatrick ZPKILPAT@COLORADO.EDU

University of Colorado Boulder

Information from preceding trials of cognitive tasks can bias performance in the current trial, a phenomenon referred to as interference. Subjects performing visual working memory tasks exhibit interference in their responses: the recalled target location is biased in the direction of the target presented on the previous trial.

We present a probabilistic inference model of this history-dependent bias, and demonstrate such inference can emerge from computations of a recurrent network with short-term facilitation (STF). Applying timescale separation methods, we obtain a low-dimensional description of the interference bias based on the target history. Delay-period activity is approximated by a particle in a slowly varying potential, attracting the particle in the direction of the previous stimulus. Target angles drawn from repetitive sequences are thus better retained in working memory than targets drawn from uncorrelated sequences. We also show that two timescales of memory degradation emerge in the delay-period activity, indicative of the slow timescale of STF dampening fluctuations later in the delay.

This framework can be extended to study how plasticity shapes neural architecture in networks to better execute evidence accumulation, as animals learn relevant features of the environment in other cognitive tasks. Ultimately, there is a great deal of promise in identifying links between probabilistic inference and the biophysics of neural computation.

I-21. Refinement of dynamic corticostriatal signals by GABAergic microcircuits

Kwang Lee[1] KWANGLEE@MEDNET.UCLA.EDU
Konstantin Bakhurin[2] BAKHURIN@G.UCLA.EDU
Leslie Claar[1] LDCLAAR@GMAIL.COM
Vishwa Goudar[1] VISHWA.GOUDAR@GMAIL.COM
Sotiris Masmanidis[1] SMASMANIDIS@UCLA.EDU

[1]University of California, Los Angeles
[2]Duke University

The corticostriatal pathway has a major role in learning, planning, and executing voluntary movements. Cortical input regulates medium spiny projection neuron (MSN) activity via several mechanisms including direct excitatory drive, as well as local microcircuit interactions. However, the overall contribution of these effects on movement-related striatal dynamics and information processing is not well understood. Here, we investigated how neural population activity and computation in the lateral striatum are shaped by a cortical projection from secondary motor cortex (M2). We trained mice to perform anticipatory licking movements during a Pavlovian reinforcement task. Since M2 and lateral striatum both mediate licking, we hypothesized that M2 contributes to at least a portion of the licking-related neural activity observed in the striatum. To test this, we transiently suppressed M2 projections in the striatum using optogenetics, while concurrently recording MSN population activity. These experiments provided a means to causally probe the net contribution of cortical input on striatal output. We found that suppressing M2 input attenuated the firing of MSNs, decorrelated local network activity, and disrupted population-level coding of lick initiation time on individual trials. In addition to regulating the activity of MSNs, we observed that suppressing M2 input reduced the firing of striatal fast spiking GABAergic interneurons (FSIs), suggesting a role for GABAergic microcircuits in processing cortical input. To examine the contribution of local inhibition on MSN activity, we pharmacologically blocked GABA-A receptor-mediated signaling in the striatum. This manipulation impaired animals' ability to initiate anticipatory licking, attenuated licking-evoked MSN activity, decorrelated the MSN network, and disrupted population-level coding. Furthermore, pairing optogenetics with neural recordings revealed that blocking striatal GABAergic signaling reduced the net gain of cortical drive on MSNs. Together, these results provide evidence for the refinement of movement-related signals as they propagate along the corticostriatal pathway.

I-22. Descending feedback pathways generate and optimize perception of second-order natural stimuli

Chengjie Huang CHENGJIE.HUANG@MAIL.MCGILL.CA
Michael G. Metzen MICHAEL.METZEN@GMAIL.COM
Maurice J. Chacron MAURICE.CHACRON@MCGILL.CA

McGill University

Growing evidence suggests that sensory systems are adapted to optimally encode incoming sensory stimulus information found in natural environments. However, natural stimuli often have complex stimulus statistics, in which the individual stimulus features must be separately extracted and transmitted in ascending pathways from the periphery to the brain. There is much evidence that sensory neurons one synapse away from the periphery can already deconstruct complex stimulus features and are capable of optimally transmitting this information through processes such as temporal whitening. In contrast to well-known feedforward pathways, very little is known about the role of descending feedback pathways. Furthermore, if these extensive feedback pathways do play a role in optimizing neural responses, what are the underlying mechanisms of each descending pathway? Finally, how do these optimized responses translate to the behavioral perception of the animal? To investigate these questions, we took advantage of the electrosensory system utilizing its well-characterized anatomy and physiology. Using a combination of in-vivo electrophysiology, pharmacology, and behavioral paradigms, we demonstrated that feed-

forward mechanisms were incapable of transmitting the necessary second-order stimulus information to sensory neurons in the hindbrain, and rather that it is the descending feedback pathways which are responsible for generating and optimizing the neural responses of hindbrain sensory neurons to give rise to behavioural perception of second-order natural stimuli. Our results reveal novel critical roles for descending feedback pathways and overturn the conventional hypotheses that signal demodulation occurs only at the sensory periphery and are further refined centrally. Due to the ubiquitous nature of feedback pathways in the brain, we believe that our results are generally applicable across sensory systems.

I-23. A cortico-collicular circuit for memory-guided directional licking behavior

Chunyu Duan ANNDUAN2@GMAIL.COM
Yuxin Pan YUXINPAN@ION.AC.CN
Signe Fruekilde SIGNEFRUEKILDE2@HOTMAIL.COM
Taotao Zhou CPUZTT@ION.AC.CN
Ninglong Xu XUNL@ION.AC.CN

Chinese Academy of Sciences

The fundamental ability to bridge past events with future actions has been studied using memory- guided decision tasks. Two interconnected regions critical for these tasks are the secondary motor cortex (M2) and midbrain superior colliculus (SC). In monkeys and rats, M2 and SC are linked to planning saccades or orienting responses. Recent work in head-fixed mice started to use directional licking as the choice effector, which allows powerful experimental techniques that require a head-restraint mouse model. These studies showed that M2 is important for planning licking responses. However, whether SC also plays a critical role in planning directional licking, as part of a distributed cortico-collicular circuit, remains unknown. Here, we use anatomical tracing, optogenetics, pathway-specific pharmacology and two-photon calcium imaging to characterize a cortico-collicular circuit for memory-guided licking behavior in mice. We first identified a putative “licking” region of SC and found multiple architectural features analogous to the “orienting” circuit. We then optogenetically inactivated either M2, SC, or both regions during the sensory, delay, or response period of an auditory discrimination task. Using a generalized linear model, we found this cortico-collicular circuit to be preferentially involved during the delay; and simultaneous inactivation of both regions resulted in a larger impairment than cortical inactivation alone. Simultaneous silencing also selectively affected more difficult trials, suggesting that the cooperation between cortex and SC may be crucial for maintaining the memory of harder decisions. These experiments provide a critical foundation for more detailed circuit dissection of memory-guided decisions. Currently, we are conducting pathway-specific pharmacology and two-photon calcium imaging to examine the information flow from M2 to SC. Our data provide direct functional evidence for SC’s importance in planning licking responses and expand SC’s role for general spatial decisions, irrespective of motor effectors.

I-24. Sensory codes for optimizing tradeoffs between task performance, adaptation speed, and resource use

Ann Hermundstad[1] HERMUNDSTADA@JANELIA.HHMI.ORG
Wiktor Mlynarski[2] MLYNAR@MIT.EDU

[1]HHMI Janelia Research Campus
[2]Massachusetts Institute of Technology

Organisms use sensory stimuli to support diverse tasks. To optimally perform such tasks with limited metabolic resources, the sensory systems that encode these stimuli must be adapted to the incoming stimulus distribution. In dynamic environments, adaptation requires accurate inference of changes in the stimulus distribution. The information required to infer these changes, however, can differ from information required to perform the task. Given limited metabolic resources, the system must balance inference accuracy with task performance. Importantly, if the system misallocates resources during a transient period when the stimulus distribution is changing, task performance could drop below the level required for survival. The extent of such performance deficits, and the speed with which the system can recover, are determined by the sensory encoding scheme.

Here, we study dynamic tradeoffs between energy use, task performance, and adaptation speed in a sensory system with limited metabolic resources. We model a system whose goal is to achieve high task performance (here, accurate stimulus reconstruction) in a fluctuating environment. The system has limited coding fidelity, and maps incoming stimuli onto a discrete number of response levels. To maintain an estimate of the stimulus distribution, the system uses this code to perform optimal Bayesian estimation. Because the encoding scheme impacts both inference and task performance, we consider optimal codes that interpolate between these two goals. We find that the system can adapt more quickly during transient periods at the cost of degraded performance when the environment is stable. Improving performance during stable periods requires an increase in resource use, but results in longer adaptation times and degraded performance during the transient. These results demonstrate that fast and accurate responses to environmental changes could require a compromise in either metabolic cost or behavioral performance when the environment is stable.

I-25. Robustness to real-world background noise: A physiological signature of non-primary auditory cortex

Alex Kell ALEXKELL@MIT.EDU
Erica N. Shook ERSHOOK@MIT.EDU
Josh H. McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

Despite longstanding interest, evidence for representational transformations between primary and non-primary auditory cortex remains elusive. Here we probed for one candidate transformation by measuring robustness to the presence of real-world background noise throughout auditory cortex. Background noise distorts the pattern of spikes in the auditory nerve, and thus to recognize sources of the interest in real-world conditions, the auditory system must separate or suppress concurrent noise sources. To assess the noise robustness of cortical auditory representations, we measured fMRI responses in twelve human listeners to thirty natural sounds (each: two seconds long). Sounds were presented in quiet as well as embedded in thirty everyday background noises selected to have stable statistics over time (e.g., a crowded theater, rain hitting pavement, crickets chirping). To quantify the noise robustness of neural responses, we leveraged the fact that a voxel's response typically varies across natural sounds and we simply correlated each voxel's response to the natural sounds in isolation with its response to those same natural sounds when superimposed on background noise.

Primary cortical responses were substantially affected by background noise (r2 ~ 0.40), but non- primary responses were more robust (r2 ~ 0.80). This difference between primary and non-primary areas was replicated in a second experiment in which speech and music stimuli were excluded, suggesting that this difference cannot be attributed to previously reported speech and music selectivity in non-primary areas. Lastly, in a third experiment we found that this difference between regions was only seen when the background noises exhibited the non-stationarities found in real- world background noise—both primary and non-primary responses were robust to spectrally-shaped Gaussian noise. Our results illustrate the neural basis of a core aspect of real-world hearing and suggest that robustness to real-world background noise is a physiological signature of non-primary auditory cortex.

I-26. Inferring connectivity and latent input covariance from spike train correlations

Cody Baker CBAKER9@ND.EDU
Robert Rosenbaum RROSENB1@ND.EDU

University of Notre Dame

A major goal in computational neuroscience is to obtain estimates of functional or synaptic connectivity via large scale, in vivo, extracellular recordings of neural activity. Several measures of functional connectivity have been proposed, but their relationship to synaptic connectivity is often not explored. Measuring the relationship between functional and synaptic connectivity requires knowledge of ground truth synaptic connectivity, which is typically unavailable in experiments. Some studies have used in silico simulations as benchmarks for investigating this relationship, but these approaches often use small networks or assume that synaptic inputs from outside the recorded network are uncorrelated. Inferring connectivity under a more biologically realistic assumption that neurons receive correlated input from unobserved sources, known as the “common input problem”, has only been studied in limited settings.

We combine spiking network simulations, analytical formulae, and calcium imaging data to give an in-depth analysis of when and how functional connectivity, synaptic connectivity, and latent variability can be untangled. We show numerically and theoretically that, under a large class of functional connectivity measures, pairwise synaptic connectivity cannot generally be inferred from functional connectivity for biologically realistic networks that receive correlated latent inputs. However, using a mean-field theory of correlated variability in balanced networks, some population-level statistics can be estimated from extracellular, multi-unit recordings. For example, the spatial scale of local connectivity and the structure of latent input covariance can be estimated from sampled spike count covariance under an assumption of approximate excitatory-inhibitory balance. We apply this method to in vivo calcium imaging in mouse primary visual cortex. Our analysis indicates that pairs of PV-expressing neurons receive more highly correlated latent input than other pairs, with a covariance structure indicating that PV-expressing neurons receive a private or mostly-private source of latent input.

I-27. Normalization of cortical excitatory-inhibitory balance by heterosynaptic plasticity

Rachel Field[1] RACHEL.FIELD2@NYUMC.ORG
James D’Amour[2] JIM.DAMOUR@NIH.GOV
Robin Tremblay[1] ROBIN.TREMBLAY@MED.NYU.EDU
Illya Kruglikov[3] ILLYA.KRUGLIKOV@NYUMC.ORG
Bernardo Rudy[1] BERNARDO.RUDY@NYUMC.ORG
Robert Froemke[1] ROBERT.FROEMKE@MED.NYU.EDU

[1]New York University
[2]National Institute of Health
[3]New York Stem Cell Foundation

The fine tuning of inhibition to excitation is a critical feature of neural networks that regulates spike generation and information processing and limits periods of prolonged over- or under-excitability. In sensory cortex, the balance between excitation and inhibition is established early on in postnatal development and must be maintained as experience- and activity-dependent modifications occur at excitatory inputs. In some forms of epilepsy, as excitatory-inhibitory (E:I) balance is disrupted, synaptic plasticity might be effective in repairing epileptic circuits. Here we examine how spike-timing-dependent plasticity (STDP) alters E:I balance across multiple synaptic inputs to layer 5 pyramidal cells in mouse auditory cortex and human temporal lobe tissue from epileptic patients. We simultaneously monitored multiple inputs onto cortical pyramidal neurons and induced synaptic modifications at one set of inputs. Manipulations at the paired input resulted in heterosynaptic modifications at the originally

best excitatory and inhibitory inputs and increased overall E:I balance. In mouse auditory cortex, heterosynaptic modifications were abolished by inhibiting calcium-induced calcium release, providing a mechanism for regulating heterosynaptic plasticity. Our results suggest that heterosynaptic plasticity can rapidly normalize excitation and inhibition in both neurotypical and epileptic circuits and may offer a promising approach to treating temporal lobe epilepsy.

I-28. Stable sequential activity underlying the maintenance of a precisely executed skilled behavior

Kalman Katlowitz AKATLOWITZ@GMAIL.COM
Michel Picardo MICHEL.PICARDO@INSERM.FR
Michael Long MICHAEL.LONG@NYUMC.ORG
New York University

A vast array of motor skills, like speaking or hitting a tennis ball, can be maintained throughout life via mechanisms that remain poorly understood. Does stable behavior necessitate underlying neuronal stability or do individual neurons alter their firing rates in a coordinated manner such that the behavioral effect of the neuronal ensemble remains identical? Neurons underlying simple motor behaviors have been shown to exhibit stable tuning curves over time, but the degree of flexibility required of these neurons to maintain more complex learned movements remains poorly understood. For instance, it has been proposed that premotor circuits resemble a manifold in which a desired behavioral performance can result from vastly different neuronal states. A major obstacle preventing the characterization of network stability is the confound of variability in the behavior itself, which can occur as the result of motor noise or redundancies in the muscular patterns needed to achieve a specific task. To circumvent these potential sources of bias, we consider the courtship song of the adult zebra finch, which is a learned, highly stereotyped motor act that relies precise movements of a small set of respiratory and syringeal muscles. By longitudinally tracking the activity of individual projection neurons in a premotor forebrain nucleus during singing using two-photon calcium imaging, we find that both the number and precise timing of song-related spiking events remain nearly identical over the span of several weeks. These findings demonstrate that learned, complex behaviors can be mediated by a surprising degree of network stability.

I-29. Designing an experiment without a human experimenter

ChangHwa Lee CKDGHK77@KAIST.AC.KR
SuYeon Heo SYH1743@KAIST.AC.KR
Sang Wan Lee SANGWAN@KAIST.AC.KR
Korea Advanced Institute of Science and Technology

Hypothesis testing in all experimental studies requires careful manipulation of experimental conditions, with which researchers try to associate independent variables while controlling for all of the other variables. However, task design examining complex cognitive processes is tricky because the relationships among task variables are often beyond the researchers' capacity to predict possible outcomes. This arises from not only a high-level correlation between independent and dependent variables, but also covariate effects. Consequently, the experimenter has to examine multiple competing hypotheses. Here we present a computational framework for experimental design that identifies testable hypotheses and covariates: Deep Neural Experimenter (DNE). The DNE consists of the hypothesis tester and the covariate identifier. The former module, based on a deep convolutional neural network, learns to discover all possible relationships between independent and dependent variables, and then visualizes the degree of contribution of each individual event to experimental conditions. The second module, based on a deep generative model, discovers information about latent task variables, including independent variables that experimenters intended to implement and covariates (if any) that experimenters failed to control. To demonstrate

applicability of this approach, we examined two different types of experiments that involve high temporal and spatial dependencies between events (causal inference task and two-stage Markov decision task). We found that the hypotheses and latent variables generated by the DNE predict the effects reported in the previous studies. We also showed that the DNE correctly identifies covariates in a situation in which a covariate was created by pseudo-random stimulus presentation. We further showed that the DNE replicates neural results obtained by model-based fMRI analysis in the previous study. The present study opens up the possibility of designing experiments while being none the worse for the complexity of the task.

I-30. Depressive model-based and model-free reinforcement learning

SuYeon Heo SYH1743@KAIST.AC.KR
Sang Wan Lee SANGWAN@KAIST.AC.KR
Korea Advanced Institute of Science and Technology

Two distinct types of learning are known to engage in decision making: a model-based and a model-free reinforcement learning (RL). Previous studies argued that people with depressive symptoms exhibit impaired performance in RL tasks (Chen 2015, Huys 2015). However, exact computational mechanisms of how depression influences the two RL are not well understood. To address this, we developed a computational approach allowing a depression factor to interact with model-based and model-free RL. Forty two participants performed a two-stage Markov decision task, in which changes in state transition probabilities were designed to control the preference for one learning strategy over the other (Lee 2014). We found that the accumulated reward, choice optimality, and choice consistency are inversely proportional to the depression score (CED-D) in a situation in which model-based RL is favored. We also found that the computational model of arbitration, incorporating a single parameter (exploitation sensitivity) representing the degree of severity of depression, accounts for choice behavior significantly better than the previously reported arbitration model (Lee 2014). Our model is built on the rational assumption that the degree of exploitation, as an indicator of optimality of the RL agent's policy, increases with the preference for model-based RL. The model provides computational evidence to support that depression tends to nullify this effect by lowing exploitation sensitivity. This provides a full account of how people with more severe depressive symptoms tend to make less coherent choices and that the effect is more pronounced in model-based RL (Huys 2015). In summary, our findings indicate how depression disrupts conversion from learned values to action during model-based and model-free RL.

I-31. Metacognitive exploration in a completely unknown state space

Su Jin An SUJINAN@KAIST.AC.KR
Sang Wan Lee SANGWAN@KAIST.AC.KR
Korea Advanced Institute of Science and Technology

Often humans make decisions based on a combination of past experience and predictions about future outcomes, rather than solely on the preference for currently available options. Reinforcement learning (RL) has provided a mathematical framework to describe how humans learn from experience and foresee consequences. However, understanding computational mechanisms underlying RL in the absence of such information has received scant attention. Here we aim to elucidate computational mechanisms of how humans explore a completely unknown and continuous state space to guide RL. In doing so, we combined computational modelling with behavioral data obtained from our novel experiment ("two-stage infinite-armed bandit problem"). The essence of our computational model is the capability to introspect its thought process and report its level of uncertainty in the course of learning, the ability known as metacognition[1-3]. We found the first evidence to indicate the distinct role of metacognition and RL in exploration; humans tend to learn locally when the degree of uncertainty about the reward structure (i.e., large reward prediction error) is greater than the uncertainty about the environmental structure, whereas

they tend to learn globally in the opposite case. The computational model, in which the uncertainty about the state-space is incorporated into the RL process (actor-critic), was also found to best account for subjects' choice behavior. These results allow us to characterize the exact computational steps that underlie the psychological construction of uncertainty during RL in an unknown environment. The results also open the possibility to address the fundamental question: how is it that human RL copes with a possibly infinite state space with limited learning capacity?

I-32. A geometrical description of global dynamics in trained feedback networks

Francesca Mastrogiuseppe[1,2] FRANCESCA.MASTROGIUSEPPE@ENS.FR
Srdjan Ostojic[1] SRDJAN.OSTOJIC@ENS.FR

[1]Ecole Normale Superieure
[2]GNT - LNC - DEC

A growing body of experimental evidence indicates that high-level computations emerge in cortical networks from mixed, non-stationary and heterogeneous representations. While traditional network models are characterized by highly homogeneous dynamics, non-trivial computational properties along with disordered responses can be obtained by training recurrent networks with algorithmic approaches [Barak 2017]. Apart from a very restricted number of cases, a rigorous understanding of how computations are supported by recurrent dynamics in trained networks is however still lacking.

A widely-adopted computational strategy consists of constraining the dynamics of a random reservoir network by means of external feedback loops [Jaeger 2001], which is equivalent to adding a unit-rank structured terms to the recurrent connectivity. Very recently, a mean-field theory of random networks with weak and uncorrelated low-dimensional connectivity has been developed [Mastrogiuseppe and Ostojic, 2017]. Here, we show that such framework can be used to derive a description of global dynamics in a simple feedback architecture. The theory allows to design optimal solutions, which can be compared with the ones emerging from common algorithmic approaches.

We focus on the case of a feedback architecture trained to reproduce a stationary output [Rivkind and Barak 2017]. We show that different classes of solutions, characterized by different stability properties, can be designed. Solutions differ in the geometrical arrangement of the readout vector with respect to the input vectors in the high-dimensional space defined by the reservoir population. We consider the batch least-square solution, and we show that its geometry can give rise to extensive instability regions. Furthermore, we show that correlations between the training solution and the random reservoir are not exploited by the algorithm to improve the network stability. On the other hand, we find that an online scheme, like FORCE learning [Sussillo and Abbott, 2009], succeeds in shaping the training solution into the optimal direction.

I-33. Piercing sensory readout via relationships between choice-related signal and microstimulation effect

Xuefei Yu XFYU@ION.AC.CN
Yong Gu GUYONG@ION.AC.CN

Chinese Academy of Sciences

How sensory information is readout by downstream neurons in the brain for behavioral outputs is always a fundamental yet unsolved question in Neuroscience. One proficient proposal of inferring sensory readout is to measure the covariation between neuronal activity and perceptual choice on a trial by trial basis, i.e. choice-related activity. However, such a metric is controversial because it is often confounded by other non-sensory driven factors in-

cluding the correlated noise among single neurons, and the decision or anticipatory signals sent from higher level areas. Indeed as discovered in neurophysiological experiments, choice-related activity does not always imply a necessary role of the sensory neurons as revealed by chemical inactivations. In addition, it could be opposite to the direction predicted from the tuning preference. Here we reconciled these facts with combination of a new experimental approach and a theoretical inference. We measured both choice-related activity and sensory readout, as revealed by microstimulation perturbation effect on a site by site basis in multiple sensory cortices including the medial superior temporal area (MST), the middle temporal area (MT), and the ventral intraparietal area (VIP) in macaque. The relationship between the two metrics is heterogeneous both across and within sensory cortices: in MST and MT, choice-related signals are indicative about their actual readout only for those neurons with congruent sensory and choice signals, whereas they are in the wrong direction for the sensory-choice opposite neurons. In VIP, choice-related signals cannot indicate sensory readout at all. An artificial network with a maximum likelihood estimator can reproduce these results by implementing two constraints: a reversed noise correlation structure and a heterogeneous readout weight between the sensory-choice congruent and opposite cells. Thus, our study provided a new approach to probing the sensory readout mechanism for perceptual decision making.

I-34. Glutamatergic ventral pallidal neurons modulate activity of the habenula - tegmental circuitry and constrain reward seeking

Meaghan Creed MEAGHAN.CREED@GMAIL.COM
Lauren Marconi LMARCONI@SOM.UMARYLAND.EDU
Jessica Tooley JTOOLEY@SOM.UMARYLAND.EDU

University of Maryland

Appropriately responding to aversive stimuli is critical for survival, while maladaptive processing of aversive stimuli is a cardinal feature of substance use disorders and mood disorders. The ventral pallidum (VP) has been implicated in the symptoms of each of these disorders and plays a critical role in processing both rewarding and aversive stimuli. However, the VP is a heterogeneous structure, and how VP subpopulations integrate into larger reward networks to ultimately modulate these behaviors is not known. The VP is classically considered a homogenous inhibitory nucleus with cholinergic projection neurons; here, we identify a non-canonical population of glutamatergic VP neurons that play a unique role in responding to aversive stimuli and constraining inappropriate reward seeking. Using multiplexed fluorescent in situ hybridization, patch clamp and in vivo electrophysiology along with viral genetic tracing, we show that glutamatergic VP neurons modulate activity of the lateral habenula (LHb), rostromedial tegmental area (RMTg) and ventral tegmental area (VTA) through direct synaptic and network effects. While inputs to these glutamatergic VP neurons did not differ from inputs to canonical VP cell types, stimulation of these cells induced behavioral avoidance while stimulation of canonical cell types produced robust place preference. We further probed the role of this non-canonical VP neural population virally-mediated using genetic ablation and fiber photometry in reward seeking and in mounting adaptive responses to aversive stimuli. Our results support a role for this glutamatergic VP neurons in enhancing salience of cues predicting aversive outcomes and in adaptively constraining reward seeking in response to aversive consequences. Maladaptive reward processing despite negative consequences is a hallmark feature of substance use and mood disorders; future studies will be needed to leverage the unique genetic and electrophysiological properties of this non-canonical neural population to selectively modulate their activity as a therapeutic strategy for disorders of these disorders.

I-35. Individual differences in adaptive decision-making reflect differences in inferential complexity

Alexandre Filipowicz ALSFILIP@GMAIL.COM
Kamesh Krishnamurthy KAMESHKK@GMAIL.COM
Joseph Kable KABLE@PSYCH.UPENN.EDU
Joshua Gold JIGOLD@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

To make effective decisions in noisy and uncertain environments, humans build and update mental models of environmental statistics that are used to make predictions and guide decision-making. There is a growing understanding of the general principles that underlie these models, particularly their relationship to normative theory. However, little is known about how these models are used by, and contribute to variability across, individual human subjects performing particular tasks. A primary challenge is the difficulty in assessing the particular mental model that a person is using for a given set of conditions. Here we propose and provide empirical support for the idea that a fundamental characteristic of mental models that governs their variability across individuals is their complexity. Moreover, we show that mental model complexity can be estimated directly from behavioral data, without assuming a model's particular form. Built on the principles of predictive information and the information bottleneck, this measure computes complexity as the amount of past information encoded by participants to predict future observations. We used this measure to capture the influence of complexity on participant performance of an adaptive-inference task in which the statistical structure changed unexpectedly. Consistent with the well-known tradeoff between bias and variance in statistics and machine learning, participants with more complex models were best able to adapt appropriately to real changes but were also more prone to overfit noisy, spurious events. Conversely, participants with less complex models were less adaptive but more robust to noise. These results imply that mental model complexity is both a theoretically and practically useful concept that can be estimated directly from behavior to obtain critical insights into the information-processing trade-offs made by individual subjects performing particular tasks.

I-36. Global and local excitation and inhibition shape the network dynamics for the control of movement and reward

Annalisa Scimemi[1] SCIMEMIA@GMAIL.COM
Caitlin Kennedy[2] KENNEDYC2@HAWKMAIL.NEWPALTZ.EDU
Joanna Herron[2] HERRONJ1@HAWKMAIL.NEWPALTZ.EDU
Anca Radulescu[2] RADULESA@NEWPALTZ.EDU

[1]SUNY Albany
[2]State University of New York at New Paltz

The cortico-striatal-thalamo-cortical (CSTC) pathway is a brain circuit that controls movement execution, habit formation and reward. Hyperactivity in the CSTC pathway is commonly observed in patients affected by obsessive compulsive disorder (OCD), a neuropsychiatric disorder characterized by the execution of repetitive involuntary movements. The striatum shapes the activity of the CSTC pathway through the coordinated activation of two classes of medium spiny neurons (MSNs) expressing D1 or D2 dopamine receptors. The exact mechanisms by which balanced excitation/inhibition (E/I) of these cells controls the network dynamics of the CSTC pathway remain unclear. Here we use non-linear modeling of neuronal activity and bifurcation theory to investigate how global and local changes in E/I of MSNs regulate the activity of the CSTC pathway. Our findings indicate that a global and proportionate increase in E/I pushes the system to states of generalized hyper-activity throughout the entire CSTC pathway. Certain disproportionate changes in global E/I trigger network oscillations. Local changes in the E/I of MSNs generate specific oscillatory behaviors in MSNs and in the CSTC pathway. These findings indicate that subtle changes in the relative strength of E/I of MSNs can powerfully control the network dynamics of the CSTC pathway in ways that are not easily predicted by its synaptic connections.

I-37. Prospection at 8 Hz in the hippocampus

Kenneth Kay[1,2] KAYKENNETH@GMAIL.COM
Jason Chung[1] JASON.CHUNG@UCSF.EDU
Marielena Sosa[1] MARIELENA.SOSA@UCSF.EDU
Jonathan Schor[1] JONATHAN.SCHOR@UCSF.EDU
Mattias Karlsson[1] MATTIAS@SPIKEGADGETS.COM
Margaret Larkin[1] MAGGIE.C.LARKIN@GMAIL.COM
Daniel Liu[1] DANIEL.LIU@UCSF.EDU
Loren Frank[1] LOREN@PHY.UCSF.EDU

[1]University of California, San Francisco
[2]Howard Hughes Medical Institute

How is it possible to think about the future? Prospection—the ability to represent future events—is essential to cognition yet remains poorly understood. Past work implicates a neural mechanism for prospection that (i) sustains a stable representation of the future over time, (ii) can switch between different future representations, and (iii) can operate during active behaviors, such as running. However, no known neural activity pattern supports these three capacities concurrently. To investigate this matter, we recorded neural activity in the hippocampus of rats navigating a simple maze with a left (L) vs. right (R) bifurcation, i.e. a choice between two future paths. We found that hippocampal neurons ("place" cells in hippocampal subregions CA1, CA2, and CA3) representing the L vs. R paths routinely fired in alternation, doing so continuously and at an unexpectedly fast characteristic frequency (8 Hz). Neural firing indicating 8 Hz alternation was specifically strongest when rats approached the maze bifurcation, and moreover was most prevalent in upstream hippocampal subregions (CA2/3). Lastly, we found that, when rats ran towards the maze bifurcation, populations of hippocampal neurons could represent the L and R paths in continuous alternation at 8 Hz. These findings outline a neural mechanism capable of supporting stable yet flexible prospection, even during active behavior, in the hippocampus and likely beyond.

I-38. A derivation of the ring model from the neural engineering framework

Omri Barak[1] OMRI.BARAK@GMAIL.COM
Amichai Labin[1] AMICHAI.LABIN@GMAIL.COM
Sandro Romani[2] SANDRO.ROMANI@GMAIL.COM

[1]Technion - Israel Institute of Technology
[2]HHMI Janelia Research Campus

The activity of large neuronal populations can represent various features of the external world. Often, this implies a mapping between a low dimensional feature, and a high dimensional state of the neural network. When the represented feature follows some dynamics, this induces constraints on the dynamics of the corresponding neural network. A recent framework termed 'neural engineering' [1] provides a numerical recipe for determining the connectivity of such a network. The resulting network implements the desired dynamics of the underlying feature, while its neurons have the requested tuning curves of the low dimensional feature. The properties of this recipe, however, were not explored analytically. Here, we analytically solve the neural engineering framework (NEF) for the case of a singular angular variable—the classical ring model [2]. We show that the connectivity structure of the model can be obtained analytically from this framework. The stability of the resulting neural dynamics, however, is sensitive to the details of the requested neuronal tuning curves. In particular, NEF is unable to produce stable narrow tuning curves (activity bumps). Furthermore, for a moving bump, the actual tuning curves deviate from those provided to the algorithm. By providing the first fully solvable example of the neural engineering framework, we show both its ability to recover the ring model and its limitations under this setting. Our results constitute an important step towards a rigorous characterization of recurrent neural network training algorithms. Further extensions will address training of spiking networks, as recently explored in [3,4].

[1] C. Eliasmith, C. H. Anderson (MIT Press, 2004). [2] R. Ben-Yishai, R. L. Bar-Or, H. Sompolinsky. PNAS (1995).

[3] Boerlin, M., Machens, C. K., & Deneve, S. PLoS comp bio 2013 [4] Thalmeier, D., Uhlmann, M., Kappen, H. J., & Memmesheimer, R. M. Plos comp bio 2016

I-39. Encoding of predictive information sheds light on circuit organization in both fly and mouse brain

Siwei Wang[1] SIWEI.WANG@MAIL.HUJI.AC.IL
Stephanie Palmer[2] SEPALMER@UCHICAGO.EDU
Alexander Borst[3] ABORST@NEURO.MPG.DE
Idan Segev[1] IDAN@LOBSTER.LS.HUJI.AC.IL
Oren Amsalem[1] OREN.AMSALEM1@MAIL.HUJI.AC.IL

[1]Hebrew University of Jerusalem
[2]University of Chicago
[3]Max-Planck-Institute of Neurobiology

Although efficient prediction is essential to survival, little is known about what mechanisms allow predictions to be instantiated in neural systems. It has recently been shown that optimal encoding of predictive information is at work in the early visual system of vertebrates [Palmer et al. 2015], but it is intriguing to test whether this theory generalizes to other systems. Here, we explore this both in data from the fly and in a model of mammalian cortex. We use an information bottleneck approach to compute the maximal amount of predictive information about a stimulus the fly visual discrimination system can encode, and show that this optimal encoding: 1) is present in the fly visual motion system (VS cells), and 2) is dependent on strong gap junction connections between VS cells. Gap junctions in this circuit are strikingly strong and support efficient prediction by passing information between triplets of VS cells without the delays incurred by chemical synapses. We further test the idea that efficient prediction is an organizing principle in brain by exploring how prediction error correlates with neuronal response in different layers of the blue brain column: dense computer-generated neocortical network of ~0.3-mm3 composed of ~31,000 cells and ~36 million synapses (Markram et al., 2015). Prediction error was shown to account for up to 50% of neuronal response variability in auditory cortex [J.Rubin et al., 2016]. With this model circuit, we show that prediction errors can indeed account for firing rate variability of a substantial fraction of pyramidal neurons. We further show that it most strongly correlates with layer 4 and layer 5 variability, suggesting that signals sent to subcortical areas, such as basal ganglia circuits, may preferentially represent predictive information about the stimulus. This supports the notion that cortical prediction is used to generate reward signals and guide learning.

I-40. A reservoir computing model of motor learning with parallel cortical and basal ganglia pathways

Ryan Pyle RPYLE1@ND.EDU
Robert Rosenbaum RROSENB1@ND.EDU

University of Notre Dame

Reservoir computing has been proposed as a model of how the brain learns and generates motor output. Most learning rules used for reservoir computing, including the popular First Order Reduced and Controlled Errors (FORCE) method, are fully supervised. As models of biological motor learning, such algorithms could only learn to “copy” the output generated in another area of the brain. Moreover, they can only be applied to tasks where the mapping from reservoir output to motor action and its inverse are known explicitly. How are novel motor outputs learned by biological neural networks? Biological motor learning is controlled at least in part by dopamine-modulated reinforcement learning in the basal ganglia. Reinforcement learning algorithms for reservoir computing have been proposed, but we find that they fail to converge on some relatively simple tasks. We develop a learning algorithm for reservoir computing, Supervised Transfer of Rewarded Exploration (SUPERTREX), that models

the interaction between reinforcement and supervised learning observed in mammals and songbirds. Through various learning tasks and simulations, we show that SUPERTREX performs as well as or better than existing learning algorithms for reservoir computing, is more biologically realistic, and is applicable to a larger class of motor learning tasks. Finally, we show that SUPERTREX can reproduce findings that relate Parkinson's disease and its treatments to motor learning.

I-41. Striatal action-value neurons reconsidered

Lotem Elber-Dorozko LOTEM.ELBER@MAIL.HUJI.AC.IL
Yonatan Loewenstein YONATAN.LOEWENSTEIN@MAIL.HUJI.AC.IL

Hebrew University of Jerusalem

It is generally believed that when choosing between alternative actions, striatal neurons represent the values associated with the different actions. This hypothesis is based on a large number of studies, in which the activity of striatal neurons was measured while the subject was learning to prefer the more rewarding action. The finding that for many neurons, the spike count covaries with exactly one of the two action-values was taken as an indication that these neurons represent these action-values. Here, following a systematic literature search, we maintain that all these publications are subject to at least one of two critical confounds and that most are subject to both. First, we show that even weak temporal correlations in the neuronal data may result in an erroneous identification of action-value representations. We demonstrate this by erroneously identifying action-value representations, both in simulations and in the neural activity recorded in unrelated experiments. Second, we show that neurons representing policy (probability of choice) may be erroneously identified as representing action-values, even when action-values are not calculated at all in the learning algorithm. We suggest different solutions to identifying action-value representation that are not subject to these confounds. Applying one of these solutions to previously identified action-value neurons in the basal ganglia we fail to detect action-value representations there. Thus, we conclude that the claim that striatal neurons encode action-values must await new experiments and analyses.

I-42. Cortical representation of egocentric head space and body posture in freely moving rats

Bartul Mimica[1,2] BARTUL.MIMICA@NTNU.NO
Benjamin Dunn[1,2] BENJAMIN.DUNN@NTNU.NO
Srikanth Bojja[1,2] SRIKANTH.BOJJA@NTNU.NO
Tuce Tombaz[1,2] TUCE.TOMBAZ@NTNU.NO
Jonathan Whitlock[1,2] JONATHAN.WHITLOCK@NTNU.NO

[1]Kavli Institute for Systems Neuroscience
[2]NTNU

Spatial navigation has been studied by tracking quadrupeds with head-mounted LEDs which, while easy to implement, effectively reduces animals to head-centered dots in 2D space. Despite its intrinsic connection to natural behavior, therefore, little is known about how the rodent brain represents posture and movement in its 3D form during unrestrained bodily motion. To address this, we tracked the head and back of nine freely moving rats in 3D, pioneering the measurement of whole-body behavior with the type of resolution that in the past was achieved for single effectors and typically in head-fixed animals. Simultaneously, we recorded single units from posterior parietal (PPC) and premotor (M2) cortices with chronically implanted silicon probes, as prior work suggested that these regions exhibit robust predictive tuning for impending orienting movements. We measured positions and directions of the head and back relative to arena-based (allocentric) coordinates, allowing for an independent estimate of head and body direction. In addition, we measured head rotations (azimuth, pitch and roll) relative to the body (i.e. in egocentric coordinates) together with postural states of the back alone. Our analysis of coding prop-

erties of individual neurons revealed remarkable specificity in tuning, mainly to the combination of different head angle positions relative to the body, but sometimes also to azimuthal flexion and the pitch of the back independently. Obtained results underscore the importance of the parieto-premotor network in providing the animals with a rich representation of head position relative to the body. This general coding scheme of egocentric head space possibly endows animals with a fast and reliable ability to localize surrounding stimuli and generate adequate responses to them.

I-43. Evidence accumulation and decision making on networks

Bhargav Karamched[1] BHARGAV@MATH.UH.EDU
Simon Stolarczyk[1] SPSTOLAR@MATH.UH.EDU
Zachary Kilpatrick[2] ZPKILPAT@COLORADO.EDU
Kresimir Josic[1] JOSIC@MATH.UH.EDU

[1]University of Houston
[2]University of Colorado, Boulder

A fundamental question in neuroscience is how organisms use sensory and social information to make decisions. Yet few mathematical models of decision making account for both types of information. For instance, popular evidence accumulation models describe how an ideal observer uses a sequence of sensory measurements to choose among alternatives. Such models have been applied successfully in neuroscience and psychology to explain a range of observed behaviors and neural data. However, these models describe an observer in isolation, whereas animals often make decisions in groups. It is thus natural to ask how an observer should combine private measurements with social information to make decisions. While heuristic models of this type have been proposed, few normative models exist.

We develop a normative model for collective decision making on a network of agents performing a two-alternative forced choice (TAFC) task. We assume that each agent is rational (Bayesian) and accumulates evidence privately until it makes a choice. This choice is observed by all of its neighbors on the network. Thus the flow of information is described by a directed network, and each deciding agent communicates their decision to those observing it. In this simplified setup the computations of rational agents can be intuitively explained, but can become extremely complex. We describe when and how the absence of a decision of a neighboring agent communicates information, and how an agent must marginalize over the decision states of all agents it does not observe directly. We also show how decision thresholds and network connectivity affect group evidence accumulation, and give a full description of decision-making dynamics in social cliques. Our model provides a bridge between abstractions used in the economics literature, and the evidence accumulator models used widely in neuroscience.

I-44. Distinct population codes for attention in the presence and absence of visual stimulation.

Adam Snyder[1] ADAM@ADAMCSNYDER.COM
Byron Yu[1] BYRONYU@CMU.EDU
Matthew Smith[2] MATT@SMITHLAB.NET

[1]Carnegie Mellon University
[2]University of Pittsburgh

Visual cortical neurons respond more vigorously to an attended stimulus than an unattended stimulus. Yet, attention states are not established instantaneously at stimulus onset: they are prepared prior to the appearance of an anticipated stimulus. However, the mechanisms of anticipation remain mysterious. One prominent proposal is that anticipatory states are characterized by gain-like modulations of spontaneous firing rates, similar

to gains seen for stimulus-evoked responses. However, evidence for this has been inconsistent. We hypothesized that anticipatory attention might be qualitatively different than the gain-like effects seen during stimulus processing, as this could improve the disambiguation of internal and external signals. Consider, for example, a decoder that determines the strength of an external stimulus through a weighted sum of visual cortical population activity, i.e., a stimulus-encoding dimension in the population state space. If anticipatory attention were simply a gain across all neurons in the stimulus-encoding dimension, as one might reasonably expect for attention effects during stimulus-processing, then the decoder might erroneously interpret this anticipatory change as a weak visual stimulus. As an alternative mechanism, we reasoned that modulations of population activity in dimensions orthogonal to the stimulus-encoding dimension might provide a better scheme for anticipation, since this would not affect the readout of the stimulus decoder. To test this idea, we recorded neural populations in visual cortex of monkeys performing a spatial attention task. The influence of attention on patterns of population activity was fundamentally different prior to stimulus onset than during stimulus processing. Moreover, the distinct features of these anticipatory states were predictive of the subjects' behavioral performance. These results defy an interpretation of anticipatory attention based on a time-invariant gain, and indicate the need to reconceptualize the neurophysiological mechanisms underlying the dynamic allocation of attention.

I-45. Biological simulation of scale-invariant time cells

Yue Liu LIUYUE@BU.EDU
Zoran Tiganj ZORAN.TIGANJ@GMAIL.COM
Michael Hasselmo HASSELMO@GMAIL.COM
Marc Howard MARC777@BU.EDU

Boston University

Scale-invariant timing has been observed in a wide range of behavioral experiments, from interval timing and free recall to classical conditioning, suggesting a common timing mechanism underlying a wide range of timescales. A possible neural substrate of this behavior comes from recent neurophysiological recordings in behaving animals during delay periods of a memory task. In these recordings individual cells show spiking activity at specific circumscribed time intervals. They are referred to as 'time cells'. Consistent with a scale-invariant representation of time, the firing fields of time cells scale with the elapsed time. Moreover, the distribution of peak firing times is scale free (power-law). However, it has not been addressed how a neural circuit could generate scale-invariant time cells. In particular the sequential activations in a chaining model are not scale-invariant. Here we present a biologically detailed neural network model that implements an approximation of the inverse Laplace transform to generate scale-invariant time cells. The model has a three-layer feed-forward architecture. The input layer consists of exponentially decaying persistent firing neurons maintained by the calcium-activated nonspecific (CAN) cation current. The intermediate layer consists of leaky integrate and fire neurons that relay the spiking activity to the output layer and ensure Dale's Law is satisfied. The connectivity between the intermediate and output layer neurons is given by the discretized inverse Laplace transform and has an on center, off surround symmetrical structure. The output layer neurons fire sequentially with scale-invariant firing rates. All the parameters in the model are biologically-realistic. This model bridges single cell, circuit and behavioral levels and provides the first biologically detailed neural circuit for generating scale-invariant time cells.

I-46. Three-dimensional representation of the motor space in the mouse superior colliculus

Nicolas Alexandre NICOLASA@MRC-LMB.CAM.AC.UK
Jonathan Wilson JWILSON@MRC-LMB.CAM.AC.UK
Caterina Trentin TRENTINCATERINA@GMAIL.COM
Marco Tripodi MTRIPODI@MRC-LMB.CAM.AC.UK

University of Cambridge

In order to reach objects or food, animals need to register their motor acts with their surrounding space. These goal-oriented movements can be described using an Eulerian representation of three-dimensional motion vectors. The association between spatial information and motor action would therefore be represented as a three-dimensional spatial-motor map. Yet, the mechanisms to encode such a representation in the brain still remain unclear. Previous studies in primates indicate that this three-dimensional representation requires a complex recomposition of motion vectors in multiple brain areas in which the superior colliculus decodes only two dimensions of motion. In our study, we use the mouse model system to analyse the neural response in the superior colliculus during head movements in a naturalistic foraging behaviour. We developed a head-mounted inertial sensor, inspired by aeronautic control system theory, for monitoring head movements. Mice were implanted with tetrode bundles in the superior colliculus to record neuronal activity during foraging. Isolated collicular neurons were analysed through a burst-triggered average analysis of head motion on each of the three Eulerian components. Our results show that neurons in the intermediate layers of the superior colliculus are tuned either individually or conjunctively to the three Eulerian components. These results indicate that, contrary to previous studies, the full dimensionality of head movements is represented in a spatial-motor map within the superior colliculus. This straightforward mechanism within the superior colliculus therefore eliminates the need for a downstream motion vector recomposition. These results also pave the way for the genetic dissection of the networks involved in goal-oriented movements. Moreover, we found that while many neurons maintained tuning in darkness, some lost their tuning in the absence of light. This opens an interesting question for future work on how sensory inputs impact the neural coding of targeted motion.

I-47. Interpretable model-based strategies arising from hierarchical neural networks

Necati Alp Muyesser NMUYESSE@ANDREW.CMU.EDU
Kyle Dunovan KDUNOVAN@ANDREW.CMU.EDU
Timothy Verstynen TIMOTHYV@ANDREW.CMU.EDU

Carnegie Mellon University

Advances in machine learning have allowed artificial agents to match human performance in many naturalistic tasks. While humans can learn generalized strategies in one context and apply them in the face of altered input structures or task goals, this adaptability remains a challenge for most state-of-the-art agents. Here we devised a learning agent, inspired by the hierarchical structure of frontal corticostriatal pathways where more abstract representations feed forward to more concrete action selection systems, in order to learn generalized action strategies. This hybrid neural network, called a Deep Model Network (DMN), is comprised of a Q-Network that learns state-action values and a reciprocally connected Model Network that attempts to identify lower dimensional patterns in the value landscape that can be exploited for performance gains. We evaluated the DMN in an environment that favors flexibility and model-building over accuracy, so as to explicitly evaluate the transfer of experiences across similar tasks and generate solutions that have explanatory power. To do this we trained the DMN to solve Wythoff's Game, a simple impartial game for which a single optimal strategy exists for multiple board configurations, providing a unique environment for testing strategy generalization in artificial agents. The DMN was able to learn near optimal, heuristic-like strategies for solving Wythoff's Game that generalized across changes in testing environment and to other variants of impartial games. Importantly, the neurobiologically inspired DMN outperformed state

of the art Deep Q-Networks at strategy generalization. Our results show how incorporating basic neurobiological principles into artificial agents, as well as utilizing structured training environments that favour model building over accuracy, allow for the development of explainable and generalizable strategies during learning.

I-48. Dissecting the contributions of rewards and effort on motivation

Vincent Valton VINCENT.VALTON@UCL.AC.UK
Anahit Mkrtchian A.MKRTCHIAN@UCL.AC.UK
Madeleine Payne M.PAYNE@UCL.AC.UK
Alan Gray ALAN.GRAY@ALUMNI.UCL.AC.UK
Veronika Samborska VERONIKA.SAMBORSKA.13@ALUMNI.UCL.AC.UK
Samantha Van Urk SAMANTHA.URK.14@ALUMNI.UCL.AC.UK
Peter Dayan P.DAYAN@UCL.AC.UK
Jonathan Roiser J.ROISER@UCL.AC.UK

University College London

Motivation is an essential component of neuro-economic decisions. It can be defined as the willingness to exert effort to attain a particular outcome, which is dependent on the magnitude of anticipated rewards and the effort costs required. Understanding motivation has wide ranging implications from being able to better gauge inter-individual variation in healthy decision-making, to dissecting symptoms such as anhedonia in depression and schizophrenia.

Here we present a computational analysis of motivation, as assessed by the Apple Gathering Task (AGT), in two independent datasets: a healthy sample (n=103), and a clinical sample (n=163—Low-risk, Familial Risk, Remitted, Depressed). In the AGT, participants squeezed a hand-dynamometer to win points. Effort required (force) and potential reward (money) were manipulated on a trial-by-trial basis. Subjects could either accept the challenge and exert effort, or refuse and skip the trial.

We analysed participants' choices using 70 competing hierarchical models of trial-by-trial behaviour. After identifying the winning model through model comparison we extracted parameters reflecting individual levels of reward and effort sensitivity. The same model won on both dataset. We then compared the model parameters to four latent variables measuring low-mood/anxiety, anhedonia, apathy and dysfunctional attitudes: these were extracted using factor analysis on 10 clinical questionnaires, with similar factor solutions found in both datasets.

As expected participants' willingness to accept an offer decreased significantly with increasing effort level, and increased significantly with increasing reward level; this pattern was accurately captured by our computational analysis. Most importantly, subjects with higher effort sensitivity parameters scored higher on anhedonia (r=0.22, p=0.037) and dysfunctional attitudes (r=0.25, p=0.017), while those with lower reward sensitivity parameters scored higher on low-mood/anxiety (r=0.35, p=0.0004).

This analysis suggests that we may be able to dissect the individual contributions of reward and effort on motivation.

I-49. Robustness of neural circuits with disparate components to intrinsic and synaptic perturbations

Sebastian Onasch[1] SEBASTIAN.ONASCH@BRAIN.MPG.DE
Julijana Gjorgjieva[1,2] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Technical University of Munich

Functionally equivalent neuronal circuits can generate similar activity patterns despite disparate intrinsic and

synaptic properties. Given these intrinsic and synaptic differences, we asked how circuits respond to neuromodulators or perturbations caused by environmental changes: do all circuits respond in the same way because of the same output, or do they respond differently as reflected in their different components? One example is given by the stomatogastric nervous system of crustaceans: it generates rhythmic output essential for their feeding behavior, which needs to be maintained over a wide temperature range. The operating ranges differ significantly across individual animals where some circuits appear more robust than others. We sought to uncover the interplay between circuit properties (intrinsic and synaptic) and dissect their contribution to the stability of circuit output.

We use high dimensional Hodgkin-Huxley conductance-based models to examine half-center circuits composed of two non-identical neurons coupled with inhibitory synapses. We propose a new measure of stability to classify the robustness of different circuits to changes in the maximum conductance value of specific channel types in the individual neurons. These changes are either applied separately for each channel type, or in combination, to examine whether the effects of different perturbations add up. We find that the stability of the circuits varies significantly, despite the fact that they all show similar output when unperturbed. In addition, we see exclusively positive correlations between the stability for changes in the maximum conductance value of any channel type and changes in the value of another. Our analysis suggests that circuit properties interact in non-trivial ways to yield robust output, and can be used to reveal specific combinations of intrinsic and synaptic properties of the underlying circuit from its response to neuromodulators and environmental perturbations that dynamically alter these characteristics.

I-50. Engineering functional brain connectivity patterns for therapeutic outcomes

Mehmet Yanik[1] MFYANIK@GMAIL.COM
Mostafa Ghannad-Rezaie[1,2] REZAIE@BIOMED.EE.ETHZ.CH
Peter Eimon[2] PETER.EIMON@GMAIL.COM
Yuelong Wui[2] YUELONG@MIT.EDU

[1]ETH Zurich
[2]Massachusetts Institute of Technology

There is a growing consensus that most neurological and psychiatric disorders are associated with large-scale dysfunction of brain networks and connectivity. Conventional drug screens and treatment regimens often ignore the underlying complexity of network dysfunctions, resulting in suboptimal outcomes. We sought to determine whether we could correct abnormal functional brain connectivity by combining multiple neuromodulators, each of which only normalizes connectivity in distinct subsets of the network, in order to correct the entire network and restore normal behavior. Our approach avoids the combinatorial complexity of a brute force screen in which all potential drug combinations are tested separately. First, we developed a high-speed platform capable of rapidly imaging brain activity in large numbers of zebrafish under multiple treatment conditions and accurately clustering drugs based on functional connectivity fingerprints, which can be used for discovery of functionally novel classes of drugs. Screening a panel of drugs in a zebrafish model of Dravet syndrome, an intractable genetic epilepsy in humans, we discovered that even drugs with related mechanisms can modulate functional connectivity in significantly different ways. Using these fingerprints, we were able to select small subsets of high-ranking complementary drugs and rapidly identify combinations capable of correcting the abnormal brain network and reducing seizures much more effectively than monotherapy. Even at the highest tolerated doses, monotherapies were unable to match the behavioral improvement of our polytherapy regimen at substantially lower doses with fewer side effects.

I-51. Adaptive mechanisms for the optimal discrimination of sparse olfactory signals

Nirag Kadakia NIRAG.KADAKIA@YALE.EDU
Thierry Emonet THIERRY.EMONET@YALE.EDU
Yale University

Neural pathways involved in olfactory sensing are able to identify and discriminate a vast array of odors, across differing concentrations and mixtures, using relatively few olfactory receptor neuron types (ORNs). Drosophila, for example, can discriminate up to thousands of odors with less than 100 receptor genes. Three aspects of odor environments and olfactory pathways suggest how this apparently under-determined computation is performed: i) natural odors are sparse in the space of chemical odorants, ii) many ORNs are not tuned to specific odorants or chemical families, but respond broadly to many odorants, and iii) synaptic projections diverge randomly downstream in the neural circuit. The mathematical framework of compressed sensing—which relies on this signal sparsity and measurement incoherence—could then be used to decode these high-dimensional odor signals reliably. Odors are thus identified by the unique combination of receptors from which they elicit responses—a combinatorial coding scheme.

How do combinatorial coding strategies cope with natural odor environments, exhibiting many simultaneous odor cues and large intensity fluctuations? One possibility is suggested by recent studies demonstrating that ORNs reduce their firing rate gain inversely with odor intensity, in accord with the Weber-Fechner law widely observed in vision. In this work, we investigate the role of this adaptive mechanism on combinatorial coding accuracy, by incorporating a biophysically-realistic stochastic model of odor binding and receptor neuron firing into the linear framework of compressed sensing. We find that, by reproducing observed adaptive scaling laws, the model preserves coding efficiency across many odor intensities, while non-adaptive sensing saturates the neural response and confounds odor recognition. Further, we find that the adaptive model permits robust discrimination of several distinct odors spanning widely differing concentrations and compositions. Together, these results support the viability of a combinatorial coding strategy in the discrimination of fluctuating and conflicting odor environments.

I-52. Constraining coupled neuronal networks to model stimulus-induced decorrelation in the olfactory system

Andrea Barreiro[1] ABARREIRO@SMU.EDU
Shree Hari Gautam[2] SHGAUTAM@UARK.EDU
Woodrow Shew[2] WOODROWSHEW@GMAIL.COM
Cheng Ly[3] CLY@VCU.EDU

[1]Southern Methodist University
[2]University of Arkansas
[3]Virginia Commonwealth University

Stimulus-induced decorrelation of spiking activity is a robust observation in the nervous system. Such state-dependent changes in correlation can result from multiple mechanisms, but the measurements needed to disambiguate them (such as synaptic strengths) may be challenging to obtain (Doiron et al., Nat. Neurosci., 2016). We developed a theoretical method to fill this gap, by using (easy to measure) spiking data to infer (hard to measure) characteristics of synaptic coupling in multi-region networks. First, we specify a mathematical model for a multi-region network. Some parameters are fixed by physiology; others are uncertain and thus define a high-dimensional space of admissible models. Observed statistics from array recordings are used as constraints: we then characterize the region of parameter space, for which computed statistics are consistent with the experimentally observed statistics. We applied this technique to dual microelectrode array recordings from two distinct regions of the rat olfactory system, the olfactory bulb (OB) and anterior piriform cortex (aPC), with and without olfactory stimulus. We compared spiking statistics between the two regions (OB vs. aPC) and between two

activity states (spontaneous vs. evoked) and identified several consistent trends; in particular, stimulus-induced decorrelation was observed in the downstream region (aPC), but not the afferent region (OB). We used a coupled, multi-population firing rate model as our network model, and synaptic strengths as unknown parameters. We found that less than 1% of parameter space was consistent with observed statistics; within this slim slice of parameter space, inhibition within the afferent region (OB) and inhibition within cortex (aPC) were constrained to a narrow slice of possible values. These predictions are validated in a full spiking network model. In principle this modeling framework can be adapted to other systems, other neural attributes (besides network strengths), and other state-dependent activity changes.

I-53. Γ and Θ rhythms in fast-spiking interneurons modulate oscillations of striatal projection neurons

Julia Chartove CHARTOVE@BU.EDU
Michelle McCarthy MMCCART@BU.EDU
Benjamin Pittman-Polletta BENPOLLETTA@GMAIL.COM
Nancy Kopell NK@MATH.BU.EDU

Boston University

Theta and gamma oscillations in the striatum manifest at behaviorally relevant times in movement, reward, and decision making tasks, as does their cross-frequency coupling. Fast spiking interneurons (FSIs) strongly inhibit the network of striatal projection neurons (SPNs), the output cells that make up 95% of striatum. Previous modeling and experimental work suggests that the SPN network can generate beta (15-20 Hz) oscillations, which are observed in healthy control of motor behavior as well as Parkinson's disease. The goal of the current study is to explore the interaction of rhythms produced in the FSI network with oscillations produced by the SPN network under different dopaminergic conditions. To explore the oscillatory dynamics of the striatal network, we simulated three interconnected cell populations: 100 FSIs, 100 SPNs expressing the excitatory D1 dopamine receptor, and 100 SPNs expressing the inhibitory D2 dopamine receptor. All neurons were simulated as Hodgkin-Huxley point neurons, with the addition of a D-current in FSIs and an M-current in SPNs. High dopaminergic tone in the FSI network was simulated by increased gap junction conductance, increased excitability, and decreased inhibitory conductance. High dopaminergic tone in the SPN network was simulated by exciting the D1 SPNs and inhibiting the D2 SPNs. Under high dopaminergic tone, the FSI network produced high gamma band (70 Hz) oscillations modulated by a theta (4-6 Hz) oscillation, while under low dopaminergic tone the FSI network produced low gamma band (55 Hz) oscillations alone. The D1 SPN network produced a beta rhythm in both conditions, but under high dopaminergic tone, the beta oscillation was modulated by the theta rhythm produced by the FSIs. A surrogate local field potential for the high dopaminergic state displayed packets of gamma power alternating with packets of beta power at a theta timescale. This alternating rhythmic behavior has implications for motor program coordination.

I-54. Identifying and characterizing hippocampal ripple-replay using semi-latent state-space models

Long Tao[1] LONGTAO2@BU.EDU
Loren Frank[2] LOREN@PHY.UCSF.EDU
Uri Eden[1] TZVI@BU.EDU

[1]Boston University
[2]University of California, San Francisco

A fundamental problem in neuroscience is to characterize complex, multi-scale neural phenomena. This usually requires integrating information across neural ensembles, brain regions, and both spatial and temporal scales.

For example, hippocampal replay events are often identified by detecting high frequency oscillations in the local field potential (LFP), called ripples, during which we observe replay, the reactivation of neural spike sequences corresponding to those seen during previous experience. Each replay event may reflect different features of previous experiences, and the information content of each event may play an important role for learning and memory consolidation. Currently, replay events, like many other phenomena in systems neuroscience, are identified in an ad-hoc manner, and there is a pressing need to 1) provide an explicit mathematical definition of what constitutes a hippocampal replay event in terms of the content and structure of spiking and LFP activity; 2) compute, at each instant, the probability that a replay event is occurring; and 3) decode the information content of each replay event. To address these challenges, we develop a novel semi-latent state-space model that includes one latent state variable indicating whether or not a replay event is occurring and another semi-latent state over position that can correspond to either the animal's actual position (as during active exploration) or to a non-local representation of position (as during replay events). By assuming that the place cells encode location similarly as function of this semi-latent state during exploration and during replay, we are able to decode from hippocampal data the replay state, the probability of replay and the content of replay simultaneously at each instant in time. This approach allowed us to identify previously overlooked replay events and thereby illustrates the power of developing a formal statistical framework to describe and detect specific patterns of ensemble activity.

I-55. Dorsal raphe regulation of movement depends on environmental valence

Changwoo Seo HISTUN@GMAIL.COM
Michelle Jin MJ382@CORNELL.EDU
Andrew Recknagel AKR33@CORNELL.EDU
Elias Wang EWANG921@GMAIL.COM
Christina Boada C.BOADA.M@GMAIL.COM
Nick Krupa NAK62@CORNELL.EDU
Yi Yun Ho YH557@CORNELL.EDU
Dave Bulkin DAB433@CORNELL.EDU
Melissa Warden MRWARDEN@CORNELL.EDU

Cornell University

Serotonin (5-HT) has been associated with an array of behavioral phenotypes including depression, patience, reward, and aversive processing. Although 5-HT neural activity during active/inactive behaviors is broadly consistent, modulation of 5-HT activity by the quality of the current environment is not well understood. Here, we investigated neural dynamics in the dorsal raphe nucleus (DRN), the primary source of 5-HT to the forebrain, in environments with different valences. Using fiber photometry, we recorded population activity from DRN 5-HT and GABA neurons while mice were actively behaving in rewarding and aversive environments. We first recorded neural activity while mice were subjected to the tail suspension test (bad environment), and were engaged in free running behavior on a running wheel (good environment). 5-HT and GABA activity increased during struggling in the tail suspension test and decreased during wheel running. These data suggest that environmental valence may play a role in modulating DRN activity, but leave open possibility that gross behavioral differences might underlie this effect. We therefore tested these mice in reward approach and punishment avoidance behaviors, in which they were required to cross a chamber either to obtain a reward or avoid a shock. When mice engaged in these running behaviors, DRN GABA neurons continued to be strongly modulated by environmental valence. DRN GABA activity increased during running to avoid the shock and decreased during running to obtain the reward, consistent with the dynamics observed during the tail suspension test and running wheel. DRN 5-HT neurons were systematically suppressed during running in both positive and negative environments. Optogenetic manipulation of DRN 5-HT and GABA resulted in context-dependent behavioral changes consistent with recorded neural activity in these populations. These data support a major role for environmental valence in modulating dorsal raphe neural dynamics during active and inactive behaviors.

I-56. Decision by sampling implements efficient coding of psychoeconomic functions

Rahul Bhui RBHUI@G.HARVARD.EDU
Samuel J Gershman GERSHMAN@FAS.HARVARD.EDU

Harvard University

The theory of decision by sampling (DbS) proposes that an attribute's subjective value is its rank within a sample of attribute values retrieved from memory. This can account for instances of context dependence beyond the reach of classic theories which assume stable preferences. In this project, we provide a normative justification for DbS that is based on the principle of efficient coding. The efficient representation of information in a noiseless communication channel is characterized by a uniform response distribution, which the rank transformation implements. However, cognitive limitations imply that decision samples are finite, introducing noise. Efficient coding in a noisy channel requires smoothing of the signal, a principle that leads to a new generalization of DbS. This generalization is closely connected to range-frequency theory, and helps descriptively account for a wider set of behavioral observations, such as how context sensitivity varies with the number of available response categories.

I-57. Dopaminergic changes in striatal pathway competition modify specific decision parameters

Kyle Dunovan[1] KDUNOVAN@ANDREW.CMU.EDU
Catalina Vich[2] CATALINA.VICH@UIB.ES
Matthew Clapp[3] MCLAPP@EMAIL.SC.EDU
Jonathan Rubin[4] JONRUBIN@PITT.EDU
Timothy Verstynen[1] TIMOTHYV@ANDREW.CMU.EDU

[1]Carnegie Mellon University
[2]Universitat de les Illes Balears
[3]University of South Carolina
[4]University of Pittsburgh

Mammals selecting actions in noisy contexts quickly adapt to unexpected outcomes to better resolve uncertainty in future decisions. Such feedback-based changes in behavior rely on plasticity within cortico-basal-ganglia-thalamic (CBGT) networks, driven by dopaminergic (DA) modulation of cortical inputs to the direct (D) and indirect (I) pathways of the striatum. DA error signals favor the D pathway over the I pathway for rewarding actions with the opposite tendency for aversive ones, effectively encoding the values of alternative actions. It remains unclear how changes in action value influence the mechanisms of the action selection process itself. Here we use a biologically plausible spiking model of CBGT networks to illustrate (1) how feedback-driven DA signals modify the strength of D and I pathways in accordance with a simple reinforcement learning model and (2) how asymmetries in D/I efficacy, resulting from the learning process, impact the accumulation of evidence for alternative actions. Simulations of corticostriatal synapses showed that DA feedback leads to asymmetrical weights in the D and I pathways within a given action channel and the ratio of these weights (wD/wI) effectively encodes the action's expected value (Q). We then simulated the full CBGT network in the context of a simple 2-choice value-based decision task under different weighting schemes for cortical inputs to the D and I pathways (high, medium, and low wD/wI) for one of the action channels. The simulated response times from these simulations were fit with two variants of a drift- diffusion model (DDM), leaving either the drift-rate or the boundary height free to vary with wD/wI ratio. As wD/wI increases, the speed of information accumulation in the decision process also increases, providing a direct mapping between network level properties of CBGT systems and cognitive decision processes.

I-58. Representation of sensory uncertainty in macaque visual cortex

Robbe Goris[1] ROBBE.GORIS@UTEXAS.EDU
Olivier Henaff[2] OJH221@NYU.EDU
Kristof Meding[3] K.MEDING@POSTEO.DE

[1]The University of Texas at Austin
[2]New York University
[3]Max Planck Institute for Intelligent Systems

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. Here we propose a novel framework which seeks to capture the best aspects of previously proposed theories1,2. We advocate a view of the visual cortex in which average response magnitude encodes particular stimulus features, while cross-neuron variability in response gain encodes the uncertainty of these features. In our model, visual neurons 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 variability3 ("The uncertainty receptive field"). To test our theory, we studied spiking activity of individual orientation-selective neurons in macaque V1 and V2 elicited by repeated presentations of stimuli whose uncertainty was manipulated in two distinct ways4. We found that gain variability systematically depends on stimulus uncertainty, irrespective of whether the source of this uncertainty is internal or external. These findings provide support for the coding scheme we propose, but are less compatible with other theories. Moreover, we demonstrate that our coding scheme is well equipped to support perceptual tasks. In particular, we show that a simple decoder of simulated V1 population activity faithfully recovers stimulus uncertainty on a trial-by-trial basis, and that gain fluctuations improve performance of an optimal decoder in perceptual tasks which require combining information from multiple sources.

I-59. Dynamic encoding of reward and latent task structures in human reinforcement learning

Dongjae Kim KIM10481@KAIST.AC.KR
Sang Wan Lee SANGWAN@KAIST.AC.KR

Korea Advanced Institute of Science and Technology

Model-based and model-free reinforcement learning (RL) has been known to explain behavioral flexibility and consistency, respectively, and interactions between these two RL guide coherent behavior1. One implementation of this idea is the arbitration control between two RL approaches, based on the respective estimation of the average amount of prediction error about rewards and states2. As the assumption of the arbitration control hypothesis is that everything is learnable, the prediction error signal simply reflects what the agent has not yet learned. However, this approach overlooks that there are some tasks that the agent cannot completely learn. Although there is growing evidence to support that the prediction error is sensitive to changes in the reward structure3, little is known about how information about different levels of task structures are encoded in prediction error and are distilled into each part of the arbitration control. To address this, we developed a computational model of arbitration control, in which the baseline level of reward prediction error (RPE) and state prediction error (SPE) were used to learn the dynamics of the task structure that were subsequently distilled into each part of the arbitration between MF and MB RL. These two baseline levels were hypothesized to be sensitive to different levels of task structure changes, such as an observable level (reward goal) and latent level (state-transition probability). We provide behavioral and neural evidence to suggest that the baseline level of RPE/SPE dynamically encodes a reward/latent task structure, respectively. We also demonstrated that the separate representations of these

two task structures converge at the region of ventrolateral prefrontal cortex, which was previously implicated in arbitration control. The present findings provide deeper insight into how PE signals encode different levels of task structures, and how dynamics of the task structure affect arbitration between model-based and model-free RL.

I-60. Discrete-attractor-like motion in continuous-attractor neural field models

Chi Chung Alan Fung FCCAA33@GMAIL.COM
Tomoki Fukai TFUKAI@RIKEN.JP

RIKEN Brain Science Institute

Place cells fire in sequence in the hippocampus of various mammals during navigation and these sequential activities are replayed during sleep presumably for memory consolidation. On one hand, attractor dynamics is believed to underlie the hippocampal memory processing. On the other hand, place cells in the rodent show theta phase precession during navigation. Pfeiffer and Foster recently explored functional relationships between the attractor and sequencing dynamics of hippocampal neurons and reported that the replayed activities comprise discrete attractors each of which encodes a single position. However, how the discrete attractors emerge in the hippocampal networks remain unknown. Here, we account for the phenomenon in continuous-attractor neural network models.

Continuous-attractor neural field models support a continuous family of neuronal activity profiles that represent fixed-point attractors of neuronal dynamics in the absence of external input. Continuous attractors were used to represent continuous information such as head direction and self-location in spatial navigation tasks and object orientation in visual perception tasks. To explore the reaction of neuronal networks to a continuously changing input, Fung, Wong and Wu previously analyzed the dynamics of a continuous-attractor neural field model tracking a continuously changing external input.

Here, we study the dynamics of a similar recurrent network model under the influences of an oscillatory drive as well as a continuous spatial input and reproduce experimental findings to clarify the role of oscillation in the hippocampal processing of spatial information. We derive analytically and numerically the conditions to have discontinuous tracking behavior using continuous attractors. Our model not only explains discontinuous tracking behavior but also replicates phase-locked changes of the spatial representation. Our results suggest that the discontinuous attractors observed in the experiment are consistent with the continuous attractor model of spatial navigation when it is driven by an oscillatory input.

I-61. Topological inference of the olfactory space dimension and the dimension of the Drosophila olfactory space.

Philip Egger PEE3@PSU.EDU
Vladimir Itskov VLADIMIR.ITSKOV@PSU.EDU

Penn State University

Despite many advances in understanding specific features of odor coding, the structure of the olfactory stimulus space is still poorly understood. The representation of perceivable olfactory stimuli is encoded in the responses of the olfactory receptors (ORs). While partial stimulus-response maps are available for a number of species (e.g. mice, Drosophila), the structure of the perceivable olfactory space is still poorly understood. This is because the stimulus-response maps are highly nonlinear, and the geometry of the presumably high-dimensional space of odors is unknown. Here we establish the monotone encoding principle and develop a method for estimating the dimension of the olfactory stimulus space from the stimulus-response map.

We first observe that the stimulus-response maps (obtained from the mouse OR deorphanization) are monotone with stimulus concentration, i.e. the relationship among the non-zero intensities of different OR responses are preserved across different stimulus intensities in the physiological range. This property guarantees that the intensity-invariant stimulus representation can be encoded via the rank orderings of the OR responses. Moreover, the monotone property of the OR stimulus-response maps allow to estimate the dimensionality of the olfactory space using methods of algebraic topology. We both developed the method for topological dimension inference and also estimated the dimension of the Drosophila olfactory space from the stimulus response maps. It turns out that the dimension of Drosophila olfactory space is relatively low (d less than or equal to 8). We propose that our method may be used in wider neuroscience contexts when the structure of the represented stimulus space is a priori unknown.

I-62. The spectrum of asynchronous dynamics in spiking networks: A theory for the diversity of non-rhythmic waking neocortical states

Yann Zerlaut YANN.ZERLAUT@IIT.IT
Stefano Zucca STEFANO.ZUCCA@IIT.IT
Stefano Panzeri STEFANO.PANZERI@IIT.IT
Tommaso Fellin TOMMASO.FELLIN@IIT.IT

Istituto Italiano di Tecnologia

During wakefulness, neocortical networks exhibit a variety of regimes characterized by low network synchrony. However, a clear theoretical understanding of the mechanisms underlying those diverse asynchronous regimes is lacking. Here, we demonstrate that, under the constraint of moderate recurrent interactions, spiking neural networks exhibit a spectrum of asynchronous dynamics ranging from excitatory-dominated sparse activity regimes to the classically reported states of balanced synaptic activity at higher firing levels. We recorded the spiking activity and the membrane potential fluctuations of pyramidal neurons during non-rhythmic epochs of spontaneous activity. We found that the variation of those properties across epochs could be predicted by the theoretical model, suggesting that neural dynamics transiently settle at various levels across the spectrum of asynchronous dynamics. Further, theoretical analysis suggested that, by moving along this spectrum, neural networks acquire the ability to either faithfully encode complex patterns of presynaptic activity or to faithfully encode the overall afferent input rate. Our results provide a theoretical model for the diversity of non-rhythmic waking states and their associated functional properties in the mouse neocortex.

I-63. Amplifying the redistribution of somatic and dendritic inhibition by an interneuron microcircuit

Loreen Hertaeg LOREEN.HERTAEG@TU-BERLIN.DE
Henning Sprekeler H.SPREKELER@TU-BERLIN.DE

Berlin Institute of Technology & Bernstein Center for Computational Neuroscience

Despite being outnumbered, GABAergic interneurons are undeniably an essential part in healthy brain activity [1]. Their rare occurrence is made up for by an astonishing diversity in anatomical and physiological properties [2] suggesting a functional “division of labor”. However, the computational role of these interneuron types and how they are supported by their individual characteristics is still vague.

One conspicuous difference between those interneurons is the location of their synapses onto pyramidal cells (PCs): Somatostatin-expressing interneurons (SOMs) preferably target the apical dendrites, whereas parvalbumin-expressing interneurons (PVs) mainly inhibit the perisomatic regions [2]. As SOM and PV cells are also connected, it has been hypothesized that these neurons play a key role in the redistribution of somato-dendritic

inhibition [3,4]. Here, we argue that a different sub-circuit consisting of SOMs and vasoactive intestinal peptide-expressing interneurons (VIPs) is optimized to control this redistribution by amplifying small top-down signals.

To test our hypothesis, we analyzed a computational model of an interneuron network comprising PVs, SOMs and VIPs. We show that mutual inhibition between SOMs and VIPs creates an amplifier that, in the extreme case of a winner-takes-all (WTA) condition, leads to a binary switch between somatic and dendritic inhibition. Furthermore, we interpret cell-type-specific properties including short-term facilitation between SOM and VIP, the lack of recurrent connections within VIP and SOM populations and spike-frequency adaptation in the light of this hypothesis, and show that they largely support the amplification properties of the network. Moreover, we find that favoring adaptation over recurrence as negative feedback mechanism prevents pathological conditions, enables an oscillatory WTA-regime and increases response homogeneity.

In summary, our analysis suggests that the SOM-VIP sub-circuit is optimized to serve as an amplifier that translates small top-down signals onto VIPs into large changes in the somato-dendritic distribution of inhibition onto PCs, thereby selectively gating different information streams.

I-64. Tracking the same neurons across multiple days in Ca2+ imaging data

Liron Sheintuch LIRONSHEINTUCH@GMAIL.COM
Alon Rubin ALON.RUBIN@GMAIL.COM
Noa Brande-Eilat N.BRANDE.EILAT@GMAIL.COM
Nitzan Geva NITZGEVA@GMAIL.COM
Noa Sadeh NOA.SADEH@WEIZMANN.AC.IL
Or Pinchasof OR.PINCHASOV@GMAIL.COM
Yaniv Ziv YANZIV@GMAIL.COM

Weizmann Institute of Science

Ca2+ imaging techniques permit time-lapse recordings of neuronal activity from large populations over weeks. These techniques facilitate within-subject analyses that quantify changes in neuronal activity over extended periods of time. However, without identifying the same neurons across imaging sessions (cell registration), longitudinal analysis of the neural code is restricted to population level statistics. Accurate cell registration becomes challenging with increased numbers of cells, sessions, and inter-session intervals, and depends on the stability of the preparation throughout the experiment. For example, small changes in the focal plane may affect the stability of the spatial footprints of cellular activity across sessions, precluding the reliable tracking of neurons. Current cell registration practices, whether manual or automatic, do not quantitatively evaluate registration accuracy, possibly leading to data misinterpretation in cases that accurate registration is unattainable. To address this, we developed a method (Sheintuch et al., 2017) that in addition to registering cells across multiple sessions, also estimates the probability of correct registration for each registered cell, and the overall registration error rates for any given data set. Using large-scale Ca2+ imaging data recorded over weeks from freely behaving mice, we show that our method is more accurate than previously used registration routines, yielding estimated error rates <5 %. Furthermore, registration accuracy remains high with increased numbers of sessions, demonstrating the utility of our method for longitudinal studies. We found that the method is applicable to various imaging techniques (one-photon/two-photon), cell detection algorithms (PCA-ICA/CNMF-E), and brain regions (hippocampus/cortex). We used several independent approaches to validate our registration method based on stability of hippocampal place cells, simulated data, and measures of internal consistency. Thus, our cell registration method facilitates reliable tracking of the same neurons across multiple imaging sessions, and provides objective measures that aim to help researchers quantify the suitability of their data for longitudinal analysis.

I-65. Spontaneous activity patterns in the developing mouse visual cortex in vivo guide retinotopic refinement of network connectivity

Marina Elaine Wosniack[1,2] MARINA-ELAINE.WOSNIACK@BRAIN.MPG.DE
Jan Hendrik Kirchner[3,4] JKIRCHNER@UOS.DE
Ling-Ya Chao[1,5] LINGYA_CHAO@HMS.HARVARD.EDU
Friederike Siegel[6] FRIEDERIKE.KRUPP@WIRGESTALTENEV.DE
Christian Lohmann[6] C.LOHMANN@NIN.KNAW.NL
Julijana Gjorgjieva[1,2] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Technical University of Munich
[3]Albert Ludwigs University Freiburg
[4]Bernstein Center for Computational Neuroscience
[5]Harvard University
[6]Netherlands Institute for Neuroscience

Many sensory systems generate spontaneous activity during early brain development and before the onset of sensory experience. In the visual system, spontaneous activity is first generated in the retina, and propagates to downstream areas including the visual cortex. Multiple studies suggest that this activity contains instructive cues for the topographic refinement of network connectivity between the sensory periphery and cortex, and the emergence of functional cortical responses. Calcium imaging recordings in the mouse primary visual cortex in vivo before eye opening have revealed two distinct patterns of spontaneous activity [Siegel et al. Curr Biol 2012]: (1) small-amplitude events involving 20-80% of the recorded neurons largely originating in the retina, termed L-events, and (2) large-amplitude events involving almost all recorded neurons (80-100%) largely originating in the cortex, termed H-events.

We combined a detailed analysis of in vivo recordings and a computational model to understand how these two distinct activity patterns jointly shape network connectivity and receptive fields under different plasticity rules. We found that L-events promote the emergence of cortical input selectivity and topographic organization of network connectivity. In contrast, H-events shaped receptive fields similar to adjusting various parameters in the plasticity rules. However, we found that the robustness of receptive field formation was too sensitive to the properties of learning rules. In the biologically plausible scenario when the operating plasticity rules are biophysically constrained, we predicted that H-events can play a powerful role in stabilizing network topography by dynamically adapting their amplitude to recent cortical activity. Re-examining in vivo spontaneous activity confirmed our predictions, revealing a strong positive correlation between the amplitude of preceding L-events to the amplitude of the following H-event. This demonstrates the remarkable capacity of the developing sensory cortex to adapt the statistical properties of spontaneous activity for the generation of robust connectivity maps.

I-66. Neural basis of optimal multisensory decision making

HAN HOU[1] HOUHAN@GMAIL.COM
Yuchen Zhao[1] YCZHAO@ION.AC.CN
Qihao Zheng[1] QHZHENG@ION.AC.CN
Alexandre Pouget[2] ALEXANDRE.POUGET@UNIGE.CH
Yong Gu[1] GUYONG@ION.AC.CN

[1]Chinese Academy of Sciences
[2]University of Geneva

To make effective decisions, organisms often need to integrate evidence across sensory modalities and over time, a process known as multisensory decision making. The neural basis of such decisions has yet to be elucidated despite its prevalence in real world situations. This problem is complex because the reliability of the evidence often changes across modalities and time. The optimal solution requires that sensory evidence be weighted in

proportion to their respective reliabilities. However, theoretical studies have shown that simply summing neural activity across modalities and time is sufficient to implement the optimal solution without the need for any form of reliability dependent synaptic reweighting, as long as sensory inputs are represented with a type of code known as invariant linear probabilistic population codes (iLPPC). Here we recorded from single neurons in macaque lateral intraparietal area (LIP) while the animals were optimally performing a visual-vestibular heading discrimination task. The relative reliability of the visual and vestibular cues was varied over the time course of each trial. We discovered a population of neurons whose response properties are compatible with the predictions of theoretical studies based on iLPPC. In particular, in the unimodal conditions, LIP neurons appear to simply integrate over time the sensory evidence related to velocity for the visual input and to acceleration for the vestibular input. In the cue-combined condition, these same neurons compute linear combinations of their visual and vestibular inputs. A simple extension of the original theory, along with a neural model, shows that these linear combinations of the temporally integrated inputs are indeed sufficient for optimal integration across time and modalities, without any need for reliability dependent reweighting at the synaptic level. Thus, our results provide the first neural evidence in support of the iLPPC theory of optimal multisensory decision making.

I-67. Revisiting prior biases and confidence in diffusion models

Jan Drugowitsch[1] JDRUGO@GMAIL.COM
Alexandre Pouget[2] ALEXANDRE.POUGET@UNIGE.CH

[1]Harvard University
[2]University of Geneva

Understanding how to incorporate prior beliefs in drift diffusion models (DDM) is a key question in decision making. When multiple difficulty levels are used (e.g., when coherence varies from trial to trial in the random dot motion task), the currently known optimal solution requires the accumulator to be biased toward the more likely decision by both moving the starting point and adding a time dependent urgency signal (Hanks et al., 2011). We show here that this solution in fact critically depends on how the prior distribution is being manipulated. If the prior distribution over the variable to be discriminated (e.g., signed coherence in the random dot motion task) is shifted toward the more likely choice, the prior over choices (e.g., left versus right) can be optimally implemented by simply shifting the starting point of the DDM without any urgency signal adjustment. In this model, the decision confidence is the same at either boundary for the same decision time, but the confidence for the prior-preferred decision is generally higher due to shorter reaction times. Together, this leads to a well-calibrated model, in which overall confidence equals the fraction of correct choices. Furthermore, and most importantly, this implementation considerably simplifies the learning of the prior from experience since it only involves learning the starting bias, without having to worry about elapsed time, in contrast to the case considered by Hanks et al. These results lead to a clear experimental prediction: shifting the prior sideways should only impact the LIP firing rate at the start of the integration period, but should leave the urgency signal unchanged. Our results demonstrate that there is no universal way of implementing a prior in a DDM while clarifying the conditions under which a simple shift of the starting point is all that is required.

I-68. Maximally separating and correlating model-based and model-free reinforcement learning

Sanghyun Yi SANGYI92@SNU.AC.KR
JeeHang Lee JEEHANG@KAIST.AC.KR
Sang Wan Lee SANGWAN@KAIST.AC.KR

Korea Advanced Institute of Science and Technology

Dopamine neurons are widely known to encode information about the discrepancy between a predicted and an

actual outcome, the essential component to regulate reinforcement learning (RL). Recent evidence suggests that there are two different types of RL to guide coherent behavior: model-free RL and model-based RL. Although a few studies manage to separate out model-based RL from model-free RL in certain contexts, the fact that the two RL agents produce similar behavior patterns in most situations has constituted fundamental challenges for decision making research. Here we propose a computational framework for deliberately manipulating episodes in a way that maximally separate or correlate these two types of RL. Unlike the general approach in which a computer is designed to simply maximize the satisfaction (reward) of human users, the present study exploits the potential of bidirectional interaction at the neural level - e.g., reward and state prediction error (RPE and SPE). The framework is based on the game play between (1) the human agent (i.e., participants) who uses RL strategies to obtain rewards from the environment and (2) the online task agent (i.e., environment) that uses RL to get desired RPE and SPE responses from the human agent. Simulations on a two-stage Markov decision task in 8 different scenarios show that the computational framework successfully learn an online task policy that maximizes or minimizes the estimated amount of RPE and/or SPE of the simulated human RL agent. This framework is applicable to any reinforcement learning task paradigm, and moreover, opens up the possibility for optimizing the design of reinforcement learning tasks that allow us to meticulously examine competition and interaction between model-based and model-free RL.

I-69. Leveraging low-dimensional structure in neural population activity to combine neural recordings

William Bishop[1] WBISHOP@CS.CMU.EDU
Erin Crowder[1] ERIN.C.HARE@GMAIL.COM
Amin Zandvakili[2] ZANDVAKILI@GMAIL.COM
Xiao Zhou[1] ZHOUXIAO@CMU.EDU
Steve Chase[1] SCHASE@CMU.EDU
Adam Kohn[2] ADAM.KOHN@EINSTEIN.YU.EDU
Carl Olson[1] COLSON@CNBC.CMU.EDU
Byron Yu[1] BYRONYU@CMU.EDU

[1]Carnegie Mellon University
[2]Albert Einstein College of Medicine

The ability to stably record from very large neural populations for months to years would represent a major advance for systems neuroscience. Despite rapid progress in recording technology, this goal remains elusive. To enable new scientific study with existing technology, we have developed a method of statistically combining many smaller recordings, referred to as neural stitching. Neural stitching enables the study of the activity of populations of neurons (1) in brain areas (e.g., sulci or subcortical areas) accessed through low-yield recording techniques, (2) over long periods of time (e.g., long-term learning), and (3) in response to orders of magnitude more stimuli than can be presented in a single experimental session. While previous stitching methods seek to leverage conserved dynamics in the time courses of neural activity over time, our method requires only partial overlap in recorded neural units or presented experimental conditions. It then estimates low-dimensional representations of the trial-by-trial or averaged condition-by-condition responses, respectively, of the union of all recorded units. To inform its use in experiments, we derived sufficient conditions under which neural stitching can be successfully applied. We then applied neural stitching to electrophysiological recordings from monkeys in three experimental contexts. We first demonstrate that neural stitching faithfully recovers low-dimensional representations of changes in population activity with learning over time, even when the set of neural units recorded at the beginning and end of an experiment are completely disjoint. Second, we demonstrate that neural stitching can be employed to study the co-fluctuation patterns between neurons in a brain area accessible with only a single electrode when simultaneous population recordings are made in a second brain area. Finally, we show that stitching can be used to expand the number of stimuli used in studying visual encoding.

I-70. A plasticity-coupling link in recurrent cortical networks with diverse learning rates

Yann Sweeney Y.SWEENEY@IMPERIAL.AC.UK
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK

Imperial College London

Chronic recordings of sensory cortex throughout periods of network reorganisation reveals a diverse range of stimulus response stability: some neurons retain stimulus responses that are stable over days whereas other neurons have highly plastic stimulus responses. Here, we propose that this observation suggests an underlying diversity in the synaptic plasticity of neurons within these networks. We explore this hypothesis by simulating synaptic plasticity with diverse learning rates in a rate-based recurrent network model of visual cortex.

In both a simple network model and a receptive-field based model of visual cortex we observed that neurons with fast learning rates are more coupled to population activity than neurons with slow learning rates. This relationship, which we call a plasticity-coupling link, occurs because neurons with faster synaptic plasticity recruit more non-specific recurrent synaptic input. This plasticity-coupling link predicts that neurons with high population coupling should exhibit more long-term stimulus response variability than neurons with low population coupling. We test this prediction by tracking the stimulus preference of neurons during calcium imaging experiments. The population coupling of neurons within these recordings are correlated with the plasticity of their stimulus preference, as predicted by our model. Finally, we explore the functional implications of both diverse population coupling and diverse learning rates within our network model. We find that high population coupling helps plastic neuron change its' stimulus preference to an associated stimulus, but hinders the ability of stable neurons to provide an instructive signal for learning.

Based on these findings we propose that the observed diversity of long-term stimulus response plasticity within these networks suggests a particular functional architecture: a stable 'backbone' of stimulus representation formed by neurons with slow synaptic plasticity and low population coupling, on top of which lies a flexible substrate of neurons with fast synaptic plasticity and high population coupling.

I-71. Transient population dynamics and computations in recurrent neural networks

Giulio Bondanelli GIULIO.BONDANELLI@ENS.FR
Srdjan Ostojic SRDJAN.OSTOJIC@ENS.FR

Ecole Normale Superieure

Across sensory systems, complex spatio-temporal patterns of neural activity can arise in response to even simple stimuli. A sudden stimulus presentation and removal evoke non-monotonic transient onset and offset responses that dynamically evolve in time and across neurons. These ON and OFF responses share remarkable similarities across sensory modalities [Saha et al. 2017], suggesting the possibility that they are generated by a common mechanism. At the population level, non-monotonic ON and OFF responses have been found to be particularly informative about stimulus identity [Mazor and Laurent, 2005; Roland et al., 2017] and may form the basis for computations based on transient dynamics [Rabinovich et al., 2008]. The mechanisms generating ON and OFF responses have so far not been fully elucidated. While a number of single-cell and network mechanisms can underlie ON responses, the number of candidates is more limited for non-monotonic OFF responses. Here we examine a network mechanism that generates both ON and OFF responses through recurrent interactions. Focusing on linear recurrent networks, we determine the conditions for the existence of non-monotonic ON and OFF responses and examine their computational properties. We start by identifying a general criterion that allows us to distinguish two sharply separated dynamical regimes depending on properties of the connectivity matrix: a regime in which all inputs lead to monotonic ON and OFF transients, and a regime in which specific inputs elicit non-monotonic transients. Using a geometrical analysis of the population activity patterns in the neuronal state space, we derive the optimal set of stimuli leading to non-monotonic responses. The dimensionality of this set determines the capacity of the network to discriminate stimuli based on strong ON and OFF transients. We apply this approach to a variety of network topologies, thus making the link between connectivity structure and temporal computations.

I-72. Biologically plausible online PCA without recurrent neural dynamics

Victor Minden VMINDEN@FLATIRONINSTITUTE.ORG
Dmitri Chklovskii DCHKLOVSKII@FLATIRONINSTITUTE.ORG
Cengiz Pehlevan CPEHLEVAN@FLATIRONINSTITUTE.ORG
Flatiron Institute

Networks with local learning rules performing Principal Component Analysis (PCA) and related tasks have been previously derived from the principle of similarity matching. However, the operation of such networks requires a fixed-point iteration to determine the neural response to a given stimulus before the next stimulus arrives. Such a feature is biologically implausible because it would demand unreasonably fast activity dynamics compared to the time scales over which natural stimuli vary. Other networks for the same tasks have similar issues, requiring fast fixed-point iteration or other biologically implausible features. Here, we derive a network for PCA-based dimensionality reduction that does not require fast resolution of iterative neural dynamics. Our approach is based on modifying the similarity matching objective to encourage diagonality of the lateral synaptic weight matrix. The key novelty of our approach is the observation that this near-diagonal structure may be exploited to determine the output neural firing rates explicitly in terms of the synaptic weight matrices via a truncated Neumann series approximation, avoiding fixed-point iteration. The resulting learning rule is both local and normative, and therefore both biologically implementable and interpretable. In the offline setting, we show that our derived algorithm corresponds to a nonlinear dynamical system with a single stable stationary point corresponding to the principal subspace. Online, we show through numerical tests that our algorithm converges to the principal subspace at a competitive rate with optimal computational complexity per iteration in terms of both input and output dimensionality, i.e., linear in the total number of degrees of freedom.

I-73. Human hippocampal theta oscillations reflect sequential dependencies in planning.

Raphael Kaplan[1] RAPHAELKAPLAN@GMAIL.COM
Adria Tauste[2] ADRIA.TAUSTE@GMAIL.COM
John King[1] JOHN.KING@UCL.AC.UK
Alessandro Principe[2] 60010@PARCDESALUTMAR.CAT
Raphael Koster[1] RKOSTER@GOOGLE.COM
Miguel Ley[2] 99151@PARCDESALUTMAR.CAT
Daniel Bush[1] DRDANIELBUSH@GMAIL.COM
Neil Burgess[1] N.BURGESS@UCL.AC.UK
Rodrigo Rocamora[2] RROCAMORA@PARCDESALUTMAR.CAT
Karl Friston[1] K.FRISTON@UCL.AC.UK
[1]University College London
[2]Universitat Pompeu Fabra

Human hippocampal theta oscillations have been linked to memory performance and choice certainty, but it remains unclear whether changes in theta power are associated with specific aspects of decision making. Notably, exploratory movement-induced hippocampal theta oscillations in rodents are related to sweeps of place cell activity that could be used to plan trajectories online. This raises the possibility that associated increases in human

hippocampal theta power also relate to on-the-fly planning of forward trajectories. We tested human subjects on a spatial planning paradigm, while recording from the hippocampus; either invasively, using intracranial electroencephalography (iEEG); or non-invasively, using whole- head magnetoencephalography (MEG). In each case, subjects were instructed to make a ~3s visual search for the shortest path between a starting and target location - within novel mazes that afforded multiple paths. Subjects were subsequently asked which direction they would go from one of two choice points along the shortest path. Crucially, the mazes were designed to induce forward planning in terms of a two-level tree search, where subjects needed to maintain the decisions they made at each choice point. This allowed us to dissociate the correlates of planning both initial and second/prospective choices. We used the iEEG recordings to identify a time- frequency window of interest for our MEG data, and subsequently observed a task-related increase in 2.5-6 Hz hippocampal theta power during the first half of the visual search period that correlated with faster reaction times in both iEEG and MEG datasets. Using MEG, we found that hippocampal theta power was highest when subjects viewed mazes with ambiguous prospective choices, independent of initial choice demands. Together, these results suggest that the human hippocampal theta rhythm is associated with a location/step update within a multi- step decision - to focus on bottlenecks during planning.

I-74. Does the anterior cingulate contribute to foraging decisions?

Gary Kane[1] GKANE@PRINCETON.EDU
Morgan James[2] MHJ32@CA.RUTGERS.EDU
Amitai Shenhav[3] AMITAI_SHENHAV@BROWN.EDU
Nathaniel Daw[1] NDAW@PRINCETON.EDU
Gary Aston-Jones[2] GSA35@CA.RUTGERS.EDU
Jonathan Cohen[1] JDC@PRINCETON.EDU
Robert Wilson[4] BOB@ARIZONA.EDU

[1]Princeton University
[2]Rutgers University
[3]Brown University
[4]University of Arizona

Animals must regularly choose between continuing to pursue a depleting resource or searching for a better alternative (i.e. foraging). Their choices qualitatively follow predictions of the optimal strategy: leave the depleting resource when its value depletes to the average reward rate across all possible alternative options. Studies of the neural mechanisms of foraging decisions have focused on the anterior cingulate cortex (ACC), a region often associated with cognitive control, offering two competing hypotheses: (1) ACC encodes the value of switching to an alternative course of action or (2) ACC signals the need to exert cognitive control, which is greater as reward depletes closer to the average rate, making decisions more difficult. Although these theories posit different roles for ACC in foraging decisions, they both predict that ACC activity will increase as reward depletes to the average reward rate. This common prediction has been confirmed by correlational studies, including single-unit recordings in monkeys and fMRI in humans. However, the causal role of ACC in foraging decisions has not been tested. In this work, we move closer to a causal examination of the role of ACC in foraging by developing a foraging task for rats. Using single-unit recordings, we find that, like humans and monkeys, neural activity in rat ACC increases as rewards deplete and rats are more likely to leave the depleting reward. These data indicate that ACC is playing the same role in rats as in primates. Moreover, pilot data using designer receptors (DREADDs) to selectively inhibit or excite ACC neurons shows that neither stimulation nor inhibition of rat ACC causes a substantial change in foraging decisions, suggesting that ACC may only perform a monitoring role in foraging decisions.

I-75. Cortically-controlled neuromodulation of spinal motor circuits to alleviate gait deficits of Parkinson's disease

Matthew Perich[1] MPERICH@GMAIL.COM
Tomislav Milekovic[1] TOMISLAV.MILEKOVIC@UNIGE.CH
Flavio Raschella[2] FLAVIO.RASCHELLA@EPFL.CH
Shiqi Sun[2] SHIQI.SUN@EPFL.CH
Giuseppe Schiavone[2] GIUSEPPE.SCHIAVONE@EPFL.CH
Christopher Hitz[2] CHRISTOPHER.HITZ@EPFL.CH
Yang Jianzhong[3] YANG@MOTAC.COM
Wai KD Ko[3] D.KO@MOTAC.COM
Li Qin[3] LIQIN@MOTAC.COM
Chuan Qin[4] CHUANQIN@VIP.SINA.COM
Marco Capogrosso[5] MARCO.CAPOGROSSO@UNIFR.CH
Stephanie P. Lacour[2] STEPHANIE.LACOUR@EPFL.CH
Jocelyne Bloch[6] JOCELYNE.BLOCH@CHUV.CH
Silvestro Micera[2] SILVESTRO.MICERA@EPFL.CH
Erwan Bezard[7] ERWAN.BEZARD@U-BORDEAUX.FR
Gregoire Courtine[2] GREGOIRE.COURTINE@EPFL.CH

[1]University of Geneva
[2]Ecole Polytechnique Federale de Lausanne
[3]Motac Neuroscience Ltd
[4]Chinese Academy of Medical Science
[5]University of Fribourg
[6]Lausanne University Hospital
[7]University of Bordeaux

Parkinson's Disease (PD) is a debilitating motor disorder affecting millions of people around the world. While levodopa and deep brain stimulation therapies alleviate most motor symptoms of PD, gait disorders are less responsive to these treatments. People with PD typically exhibit short and slow steps, balance deficits, and freezing of gait. These symptoms arise when cortical commands meant to activate and modulate spinal locomotor centers are rendered dysfunctional. Yet, the underlying mechanisms of PD remain poorly understood, limiting our ability to develop effective clinical interventions. Using monkeys treated with MPTP—a chemical that selectively kills dopaminergic cells in the substantia nigra—we studied the impact of PD on the interactions between the cortical and spinal circuits that control locomotion. We recorded bilateral leg motor cortical activity during attempted movements before and after the development of PD symptoms. We found that motor cortical activity could be used to accurately decode movement gait events in the presence of PD-related gait deficits. We then developed a novel neuroprosthetic intervention that uses intracortically-recorded brain activity to synchronize the delivery of epidural electrical stimulation of the spinal cord with descending cortical movement commands (Figure 1). Thus, this neuroprosthesis regulates the activity of spinal motoneuron pools, which alleviates gait deficits. We evaluated the therapeutic potential of our neuroprosthetic system in one MPTP-treated monkey that developed Parkinsonian motor symptoms. Our neuroprosthesis not only increased the stability and speed during basic walking, but also restored skilled locomotor control (Figure 2). Our experiment reveals the therapeutic potential of our neuroprosthesis to alleviate PD motor deficits, and may help to understand the impact of PD on the coordination between cortical and spinal motor circuits.

I-76. A self-organizing memory network

Callie Federer[1] CALLIE.FEDERER@UCDENVER.EDU
Joel Zylberberg[2] JOEL.ZYLBERBERG@GMAIL.COM

[1]University of Colorado School of Medicine
[2]University of Colorado

Working memory relies on us retaining information about stimuli even after they go away. Stimulus information is encoded in the activities of neurons. Those neurons' activities change over timescales of milliseconds, yet the information can be retained for tens of seconds, suggesting the question of how time-varying neural activities maintain stable representations. Prior work from Druckmann and Chklovskii shows that, if the neural dynamics are in the "null space" of the representation—so that changes to neural activity do not affect the downstream read- out of stimulus information—then information can be retained for periods much longer than the time-scale of individual-neuronal activities (called the FEVER model). According to the Gershgorin Circle Theorem, this will almost surely not happen in randomly-connected networks: finely tuned connectivity is needed. Druckmann and Chklovskii suggest a mechanism of Hebbian learning by which this fine-tuned connectivity can be learned, but this mechanism requires the read-out weights to form a 'tight-frame', which will not necessarily be true in biological circuits. We identify biologically plausible synaptic plasticity rules that organize a network to enable persistent representations of stimulus information despite time-varying neural activities with no fine-tuning requirements. We specifically address parametric working memory, where the requirement is to remember continuous values describing several different variables or dimensions such as spatial locations or sound frequencies. We performed experiments to demonstrate that networks using these plasticity rules are able to store information about multiple stimuli, work even if only 10% of the synapses are tuned, and with only 10% connectivity in the network. The networks are robust to synaptic noise with amplitude as large as the size of the update. We have also demonstrated that these networks improve over time with multiple stimuli presented, and that continuous training is not required to learn information about novel stimuli. We show that the synaptic updates can successfully form self-organizing memory networks even when reduced to local Hebbian or anti-Hebbian learning rules.

I-77. Prioritized memory access explains planning and hippocampal replay

Marcelo Mattar MARCELOMATTAR@GMAIL.COM
Nathaniel Daw NDAW@PRINCETON.EDU

Princeton University

To make decisions, animals must evaluate outcomes of candidate choices by accessing memories of relevant experiences. Recent theories suggest that phenomena of habits and compulsion can be reinterpreted computationally as selectively omitting such computations. Yet little is known about the more granular question which specific experiences are considered or ignored during deliberation, which ultimately governs decisions. We propose a normative theory to predict not just whether but which memories should be accessed at any time to enable the most rewarding future decisions. Using nonlocal "replay" of spatial locations in hippocampus as a window into memory access, we simulate a spatial navigation task where an agent accesses memories of locations sequentially in order of the expected utility of the computation: how much more reward is likely to be earned due to better choices. The theory suggests that priority for considering a particular location depends on the product of a need term (measuring the likelihood it will soon be visited) and a gain term (measuring the improvement in reward expected). Need promotes activity forward of the agent for planning; gain favors backward replay propagating prediction errors to predecessors. Our theory offers a unifying account of a large range of hitherto disconnected findings in the place cell literature such as the balance of forward and reverse replay, biases in the replayed content, and effects of experience. We suggest that the various types of non-local events during behavior and rest reflect different instances of a single choice evaluation operation, unifying seemingly disparate proposed functions of replay including planning, learning and consolidation, and whose dysfunction may explain issues like rumination and craving. Our connection between hippocampal and reinforcement learning literatures suggests many

new experiments investigating the links between experience, replay and planning.

I-78. Cell type specific spatial and temporal integration rules in retinal ganglion cell dendrites

Ziwei Huang[1] ZIWEI.HUANG@UNI-TUEBINGEN.DE
Yanli Ran[1] YANLI.RAN@CIN.UNI-TUEBINGEN.DE
Katrin Franke[1] KATRIN.FRANKE@CIN.UNI-TUEBINGEN.DE
Tom Baden[2] T.BADEN@SUSSEX.AC.UK
Philipp Berens[1] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Thomas Euler[1] THOMAS.EULER@CIN.UNI-TUEBINGEN.DE

[1]University of Tuebingen
[2]University of Sussex

Retinal ganglion cells (RGCs) are the output neurons of the eye, which send the visual information extracted by the retinal network to the brain. Recent functional and anatomical work suggests that there are more than 30 RGC types in the mouse retina, each selective for specific visual features like contrast, motion or edges(Sanes & Masland, 2015; Baden et al, 2016). How individual RGC types integrate information across their dendrites is still largely unknown. Here, we record local calcium signals to various light stimuli in dendritic segments using two-photon Calcium imaging to study how different RGC types integrate signals in their dendrites. Our preliminary data indicate that different RGC types exhibit distinct forms of spatial and temporal integration: for Off transient alpha RGCs, dendritic segments closer to the soma systematically exhibit large, overlapping receptive fields (RFs), whereas segments closer to the dendritic tips are small and spatially more independent. In contrast, Off transient “mini”-alpha RGCs, which co-stratifies with the former, exhibit no systematic change in RF size with distance from the soma. Interestingly, dendritic RF centers in Off transient “mini”-alpha RGCs are systematically spatially offset towards the soma compared to their dendritic location. Off transient alpha RGCs do not show this behavior. In agreement with these spatial findings, temporal correlations for “mini”-alpha RGCs are high throughout the entire dendritic tree of the cell, but to a lesser extent for Off transient alpha RGCs. Taken together, by relating local dendritic activity measured by calcium imaging to a cell’s morphology, we found evidence for cell type-specific spatio-temporal integration rules in mouse RGCs . We are currently establishing biophysical models and pharmacological experiments to investigate the relative contributions of cell-type intrinsic factors, such as the expressed subset of ion channels and their spatial distribution.

I-79. Trade-off between multisensory integration and cross-modal interference in rats engaged in visuotactile pattern discrimination

Alessandro Di Filippo ALESSANDRODIFILIPPO@GMAIL.COM
Alessio Ansuini ALESSIOANSUINI@GMAIL.COM
Luca Godenzini LUCA.GODENZINI@GMAIL.COM
Mathew Diamond DIAMOND@SISSA.IT
Davide Zoccolan ZOCCOLAN@SISSA.IT

SISSA

When sensing the external world, the brain needs to combine signals collected through different senses into a unified percept. Along the road to perception, suppressive interaction of cross-modal information in primary sensory cortices may coexist with optimal integration in higher brains areas. Behaviorally, finding evidence of the co-occurrence of both processes is extremely difficult, since their effects can counterbalance when measures of perceptual acuity are obtained. Here we addressed this issue by training rats in a multimodal discrimination task (with the animals sensing solid oriented grating through either the visual, tactile or visuotactile modality) and applying a classification image approach to uncover the perceptual strategies deployed by the animals under each sensory modality. This allowed expressing the multimodal strategy as a combination of the unisensory strategies, thus obtaining a measure of their integration. Crucially, this measure (being assessed at the level of perceptual strategies rather than discrimination performances) depended only on the processing at the integrative stage, being immune to the possible interference of the signals in primary areas. Concomitantly, we computed a multimodal benefit index to quantify the overall behavioral output of visuotactile perception, inclusive of both primary and integrative processing. By applying this approach we successfully isolated the impact of early sensory interference from later multisensory integration. We found instances where the multimodal benefit was either null or negative despite the visual and tactile representations being still integrated at the level of perceptual strategies. To interpret these data, we built an ideal observer model that: 1) performed an optimal integration of the sensory signals; but 2) allowed a mutual interference of the signals at the level of primary representations. The model faithfully reproduced the trends observed in the data, thus supporting the notion that the two processes (early interference and late optimal integration) do coexist during multimodal sensing.

I-80. Frequency domain structure of intrinsic infraslow dynamics of cortical microcircuits

Michael Okun[1] M.OKUN@LE.AC.UK
Kenneth Harris[2] KENNETH.HARRIS@UCL.AC.UK

[1]University of Leicester
[2]University College London

Cortical activity is organized across multiple spatial and temporal scales, yet most neurophysiological research (fMRI aside) concerns activity on timescales of 1ms-1s. Thus, little is known about neuronal dynamics on timescales of tens of seconds and minutes. We addressed this question by analyzing multi-hour recordings of ongoing activity, performed in awake head-fixed mice using chronically implanted high-density probes. For single neurons, the interspike-interval (ISI) distribution did not capture the slow changes in firing rate. We therefore developed an algorithm for generating synthetic spike trains with pre-specified ISI distribution and power spectral density (PSD), where the latter was aimed to describe the slow dynamics. These synthetic spike trains quantitatively recapitulated the original data. At the population level, PSD of the population rate was several-fold higher than the sum of PSDs of its constituent spike trains, indicating that changes in firing rate are correlated among the neurons in all slow frequencies. To investigate this correlation structure, we considered the coherency of each neuron with population rate, i.e. extending the concept of population coupling into the frequency domain. We found that the strength of coherence with population is often not maintained across frequencies: a neuron strongly coupled to population rate on fast timescales can be weakly coupled on slow timescales, and vice versa. More surprisingly, most neurons' phase with respect to population rate was significantly different from 0, particularly in the infraslow frequencies. Furthermore, three different types of observed phase structure were incompatible with an existence of a fixed time lag between individual neurons and population rate: constant phase which is not 0 or π "bumps" in phase spectrum; and phase change on logarithmic rather than linear scale. In summary, our work reveals some of the organization of neuronal infraslow dynamics at the single cell level, beyond the mesoscale organization demonstrated using fMRI.

I-81. STDP for stochastic synapses: an empirical Bayes approach

Jannes Jegminat[1,2] JANNES@INI.UZH.CH
Jean-Pascal Pfister[1,2] JPFISTER@INI.UZH.CH

[1]University of Zurich
[2]ETH Zurich

Synaptic transmission is stochastic across a wide range of species and brain areas. Each stochastic synapse represents a distribution over EPSP-amplitudes (or weights), rather than a fixed amplitude. Despite the omnipresence of stochastic synaptic transmission, its computational role is not well understood. Here we derive a Bayesian learning rule for the synaptic parameters p and q, release probability and quantal size respectively, and show that it outperforms a deterministic benchmark in a supervised learning task. Unlike dropout, our learning rule has the attractive feature of relying on natural synaptic sampling during learning and prediction. As a result of the stochastic release, it predicts that the gain of the neuronal transfer function is given by standard deviation of the membrane potential under stochastic release. Moreover, it generalizes to spiking networks and recovers STDP.

I-82. Fly motion estimates use higher-order correlations to cancel noise induced by natural scenes

Juyue Chen[1] JUYUE.CHEN@YALE.EDU
Holly Mandel[1] HOLLY.MANDEL@BERKELEY.EDU
James Fitzgerald[2] FITZGERALDJ@JANELIA.HHMI.ORG
Damon Clark[1,3] DAMON.CLARK@YALE.EDU

[1]Yale University
[2]HHMI Janelia Research Campus
[3]MCDB

Canonical motion detectors, including the motion energy model and Hassenstein-Reichardt correlator, rely exclusively on second-order spatiotemporal correlations to compute motion from visual input. Theoretical work indicates that higher-order correlations provide further information about motion in natural environments, and in particular that odd-order correlation signals can improve motion estimates when scenes are light-dark asymmetric. Here we show that flies incorporate third-order correlations to reduce noise that originates from the structure of natural images. First, we systematically characterized fly motion percepts by extracting Wiener kernels from flies' turning behavior. These comprised second-order kernels and third-order kernels, which are the lowest odd-order kernels containing directional information. The measured kernels were then used to predict flies' motion estimates to naturalistic motion, which we simulated by rigidly translating natural scenes. We separated the predicted responses into second-order motion estimates and third-order motion estimates based on which correlation types they incorporated. The motion estimates derived from the second-order kernels correlated positively with the true velocity. Interestingly, the isolated third-order motion estimates anti-correlated with true image velocities, but when the third-order motion estimates were added to the second-order motion estimates, the accuracy of motion estimation was substantially improved. We wanted to know which features of a scene were sufficient for the third-order motion estimator to improve performance. To investigate this, we generated maximum entropy models of natural scenes that preserved the pairwise correlation function of the original scenes and imposed constraints on the point statistics. Using these synthetic scenes, we found that a naturalistic skewness was sufficient for the measured third-order kernel to improve motion estimation. Since all visual systems have to cope with the light-dark asymmetry of the visual world, the fly's strategy of using third-order correlations to cancel image-induced noise might be generally relevant for dissecting visual processing in natural contexts.

I-83. Dynamics of networks of conductance-based neurons in the strong coupling limit

Alessandro Sanzeni[1] ALESSANDRO.SANZENI@NIH.GOV
Mark Histed[1,2] MARK.HISTED@NIH.GOV
Nicolas Brunel[3] NICOLAS.BRUNEL@DUKE.EDU

[1]National Institute of Mental Health
[2]National Institutes of Health

[3]Duke University

The balanced state model explains the irregular firing observed in cortical neurons with a cancellation between excitatory and inhibitory inputs to cells in the network. Such a state emerges dynamically in the strong coupling limit (large number of synaptic connections per neuron K) of networks of binary units and current-based spiking neurons. It is not known if this result generalizes to networks with more biologically accurate model neurons. In particular, a scaling argument for networks of conductance-based neurons shows that, for large K, the network should be either silent or active at maximal firing rate, raising the possibility that the balanced state idea may not be valid in these models.

We investigate the dynamics of networks of conductance-based neurons in the strong coupling limit. Using a mean field approach we show that a nontrivial activity emerges if the average synaptic efficacy J scales as K~exp(1/J)/J^(1/2), instead of the J~1/K^(1/2) scaling observed in networks of current-based neurons (Van Vreeswijk and Sompolinsky 1996). This scaling is shown to generate linear network transfer function and irregular firing (CV of order one); generalizing the balanced state model to networks of conductance-based neurons. In this model, unlike current-based models, response gain is a nonmonotonic function of recurrent inhibitory strength. Moreover, fluctuations in J exponentially modify (suppress or saturate) network response, suggesting the need for an homeostatic mechanism which regulates K or J itself to maintain the operating regime of the network. Finally, both the scaling derived here and the one observed in the current-based model agree with culture data (Barral and Reyes 2016), but they give different predictions for in-vivo connectivity.

In conclusion, our analysis shows that neuronal conductance strongly impacts how network structure affects its response and makes testable predictions on the scaling relation between J and K.

I-84. Characterization of time-variable hippocampal sequences in evidence accumulation and decision-making

Edward Nieh[1] ENIEH36@GMAIL.COM
Sue Ann Koay[1] KOAY@PRINCETON.EDU
Lucas Pinto[1] LPINTO@PRINCETON.EDU
Carlos Brody[2,1] BRODY@PRINCETON.EDU
Nicolas Freeman[1] NWF2@PRINCETON.EDU
Mark Ioffe[1] MIOFFE@PRINCETON.EDU
David Tank[1] DWTANK@PRINCETON.EDU

[1]Princeton University
[2]Howard Hughes Medical Institute

In an ever-changing environment, the ability to accumulate and evaluate evidence is crucial for optimal decision-making. The hippocampus is thought to embody a cognitive map and has been well-studied in navigation and foraging tasks. Hippocampal neurons have been shown to encode place, time, or even auditory frequency. However, it is unknown how the hippocampus behaves when evidence for decision-making is accumulated. We performed 2-photon recordings in head-fixed mice performing a virtual reality T-maze task, where noisy evidence is accumulated over time to indicate a left or right turn at the end of the maze. Surprisingly, we found a lack of reliable place cells in 3557 recorded neurons (n=5 animals) in this task, even though reliable place cells were repeatedly recorded in a related simpler task, where animals alternated between left and right turns in the same maze. The large numbers of cells that could be recorded simultaneously with our imaging technique (n=711 average simultaneously recorded cells) allowed analyzing population activity in novel ways. We searched for ordered pairs of CA1 cells that were sequentially active in many single trials, regardless of place or time, and found that these “doublets” significantly populated our dataset. These time-variable sequences appeared significantly more often than chance and were also more predictive of left or right turns than independent single cells. The doublets could also be combined to form longer sequences. These results indicate two important findings: (1) CA1 neurons do not encode place or time in an environment where evidence is variable and accumulated over time and (2) CA1 neurons form ordered yet time-variable sequences that are more predictive of animal behavior than environmental

variables such as place or time and are more predictive than independent neurons alone. Our study suggests time-variable CA1 sequences as novel and potentially critical representations in decision-making.

I-85. Adaptive spiking neural networks with efficient computation

Davide Zambrano[1] D.ZAMBRANO@CWI.NL
Roeland Nusselder[1] ROELAND.NUSSELDER@GMAIL.COM
H. Steven Scholte[2] H.S.SCHOLTE@UVA.NL
Sander Bohte[1] S.M.BOHTE@CWI.NL

[1]Centrum Wiskunde & Informatica
[2]University of Amsterdam

Recent advances in Deep Learning have proven Artificial Neural Networks (ANNs) to be highly effective, reaching and sometimes exceeding human performance on several problems. Despite such improvements current deep-ANNs are far from being a good representation of the neural coding in the brain: ANNs lack a natural notion of time, and neural units in ANNs exchange analog values in a computationally and energetically inefficient form of communication. In contrast, biological neurons communicate sparingly and efficiently using isomorphic spikes. Although Spiking Neural Networks (SNNs) can be constructed by replacing the units of an ANN with spiking neurons, current SNNs are uncompetitive with respect to deep-ANNs, and these SNNs use much higher firing rates compared to their biological counterparts, limiting their efficiency. Here we show that by employing spiking neurons with an efficient form of neural coding, SNNs can match high-performance ANNs: these SNNs exceed state-of-the-art in SNNs on important benchmarks and require much lower average firing rates. We use the firing rate limiting adaptation phenomenon observed in biological spiking neurons, that can be captured in fast adapting spiking neuron models, for which we derive the effective transfer function. Neural units in ANNs trained with this transfer function can be substituted directly with adaptive spiking neurons, and the resulting Adaptive SNNs (AdSNNs) can carry out classification in deep neural networks using up to an order of magnitude fewer spikes compared to previous SNNs. Moreover, in AdSNNs the neural coding precision can be dynamically modulated as a form of attention that selectively increases neuronal firing rates, similar to what is observed in biology. We show how a simple model of arousal in AdSNNs further halves the average required firing rate. AdSNNs thus hold promise as a novel and efficient model for neural computation that naturally fits to temporally continuous and asynchronous applications.

I-86. A dynamic connectome supports the emergence of stable computational function of neural circuits

David Kappel DAVIDK@SBOX.TUGRAZ.AT
Wolfgang Maass MAASS@IGI.TUGRAZ.AT
Robert Legenstein LEGI@IGI.TUGRAZ.AT
Stefan Habenschuss HABENSCHUSS@IGI.TUGRAZ.AT

Graz University of Technology

The connectome is dynamic: Networks of neurons in the brain rewire themselves on timescales of hours to days even in the absence of neural activity [1'3]. The recent study of [3] showed that this spontaneous component is surprisingly large, at least as large as the impact of pre- and postsynaptic neural activity. We address two questions raised by these data: How can stable network performance be achieved in spite of continuously ongoing rewiring and activity-independent synaptic plasticity? What could be a functional role of these processes? We show that spontaneous synapse-autonomous processes, in combination with reward signals such as dopamine, can explain the capability of networks of neurons in the brain to configure themselves for specific computational tasks, and to compensate automatically for later changes in the network or task. The resulting model for reward-

gated network plasticity builds on the approach from [4]. We show theoretically and through computer simulations that stable computational performance is compatible with continuously ongoing synapse-autonomous changes. After reaching good performance these synaptic dynamics cause primarily slow drifts of network architecture and neuron dynamics in task-irrelevant dimensions, as observed in motor cortex and other areas. The resulting model gives rise to a novel perspective on reinforcement learning through a reward-driven sampling process of network connectivities.

References [1] A. J. Peters, S. X. Chen, and T. Komiyama. Emergence of reproducible spatiotemporal activity during motor learning. Nature, 2014. [2] A. R. Chambers and S. Rumpel. A stable brain from unstable components: Emerging concepts, implications for neural computation. Neuroscience, 2017. [3] R. Dvorkin and N. E. Ziv. Relative Contributions of Specific Activity Histories and Spontaneous Processes to Size Remodeling of Glutamatergic Synapses. PLoS Biology, 2016. [4] D. Kappel, S. Habenschuss, R. Legenstein, and W. Maass. Network Plasticity as Bayesian Inference. PLoS Computational Biology, 2015.

I-87. Decoding motion in natural scenes from disparately tuned populations of retinal neurons

Jon Cafaro[1] CAFAROJ@GMAIL.COM
Greg Field[1] GREGORY.FIELD@DUKE.EDU
Joel Zylberberg[2] JOEL.ZYLBERBERG@GMAIL.COM

[1]Duke University
[2]University of Colorado

In the mammalian retina, four distinct populations of ON-OFF Direction-Selective (ooDS) Retinal Ganglion Cells (RGCs) respond to motion along four cardinal axes. Recent studies in rodents show these cells shape motion signals in cortex and likely contribute to visually guided behavior. Understanding how reliably ooDS RGCs can inform downstream circuits about motion requires decoding their responses to natural movies at timescales relevant for neural signal processing and behavior. However, previous studies have focused on simple stimuli (e.g moving bars) and accumulating spikes over long periods of time (~1 sec). To understand the decoding constraints imposed by natural motion and to determine decoding approaches that reliably extract motion direction, we measured spike activity from hundreds of identified RGCs types recorded on a large-scale multielectrode array while showing parameterized natural motion. We used these data to test different decoding approaches applied to actual and simulated data. We find that ooDS RGC responses are strongly influenced by the local contrast changes in natural movies, which corrupt the decoding of motion direction. Specifically, local image structure decorrelates ooDS RGC responses because nearby cells view different parts of the image. This makes linear decoding of the ooDS RGCs population subject to large errors. We find these errors can be reduced by decoders that utilize both the ooDS RGCs responses and their synchronous activity with non-DS RGCs. Specifically, an Optimal Quadratic Estimator (OQE), that has access to ooDS and non-DS RGCs provides superior performance by discounting ooDS RGC spiking as being informative of motion direction when it is accompanied by non-DS RGC spiking. These results reveal the benefits of combining divergent streams of information when decoding complex natural stimuli.

I-88. Cortical column divergence in sensory processing

He Zheng HE.ZHENG@MPFI.ORG
Hyung-bae Kwon HYUNGBAE.KWON@MPFI.ORG

Max Planck Florida Institute for Neuroscience

The cortical column is one of the most fundamental organization principles of the neocortex. In the rodent whisker pathway, a whisker deflection evokes neural activity largely confined in its corresponding column in layer 4 of the primary somatosensory cortex (S1). However, this column structure diverges in layer 2/3 (L2/3), where axons extend across columns, resulting in the whisker response spanning multiple columns. The functions of inter-column connections are unclear. How does the spatial organization in S1 affect stimulus encoding? Does it drive decision-making in higher-order cortices? Does a spatially wide-spread response to a whisker enhance its detectability, but obscures the spatial acuity? To answer these questions, we put mice in enriched environment to induce plasticity in S1. The enrichment experience enlarged the spatial span of the neural response to whisker stimulation in L2/3 neurons, likely due to enhanced inter-column connectivity. To determine how this broadened sensory response in S1 affect perceptual behavior, animals were trained on a single-whisker detection task, then a two-adjacent-whisker discrimination task. Simultaneously, we monitored activity of individual neurons using two-photon calcium imaging. In a separate group of animals, activity of L2/3 neurons in the premotor cortex (M2) were also measured to investigate how broadened sensory response in S1 is represented in a higher-order cortex. Both S1 and M2 cells were analyzed to quantify how they encode stimulus and whisker detection/discrimination decision. Preliminarily, we find that, in enriched animals, L2/3 cells in S1 decrease the acuity of whisker tuning, and information-encoding cells are more widespread spatially compared to control animals, likely due to enhanced inter-column connectivity. In M2, although there is no spatial segregation to different whisker stimulation, cells exhibit whisker tuning. By uncovering the functions of horizontal connections in S1, this study will further our understanding of the relationship between function and structure in sensory circuits.

I-89. Inferring connectivity profiles from contrast invariant network activity

Nataliya Kraynyukova NATALIYA.KRAYNYUKOVA@BRAIN.MPG.DE
Tatjana Tchumatchenko TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

Max Planck Institute for Brain Research

One of the most ubiquitous representations of neural activity observed throughout sensory modalities is a tuning curve. In the visual cortex the tuning curves have been shown to be contrast invariant. Contrast invariance means that after normalizing by the peak value tuning curves recorded at different contrast values will have a universal shape. Here, we show that if the activity of a neural rate network is contrast invariant, then network connectivity and input profiles have a specific, mathematical relation to activity tuning curves. This simple theoretical observation provides constraints on the possible network connectivities and could serve as a useful metric to infer connectomes.

I-90. The corticothalamic loop can control cortical dynamics for flexible robust motor output

Laureline Logiaco LL3041@COLUMBIA.EDU
Laurence F. Abbott LFA2103@COLUMBIA.EDU
Gary Sean Escola GSE3@COLUMBIA.EDU

Columbia University

To study the mechanisms by which motor circuits generate the activity patterns underlying sequential behavior, we employ an anatomically constrained model in which inhibitory basal ganglia (BG) output neurons project to thalamic units that are themselves bidirectionally connected to a recurrent cortical network. During movement sequences, BG neurons show sustained activity that switches at the boundaries between sequence components. Assuming that these BG signals silence their thalamic targets, the dynamics of the corticothalamic system during each sequence component is governed by the non-silenced thalamic units. Here, we consider the analytically tractable case where linear thalamic "units" are activated singly allowing them to interact with a linear cortex. This produces a piecewise linear model of a type that has been show to fit cortical dynamics. The thalamic activity switch acts as a rank-one perturbation of the corticocortical connectivity matrix that depends on corticothala-

mic connectivity. Half of the parameters describing this connectivity can be adjusted to control the eigenvalues governing the cortical dynamics, while the remaining parameters can be used to minimize the impact of noise on the motor output. Our analysis reveals that the corticothalamic architecture is well-adapted to produce long and complex motor sequences. With linear cortical dynamics, corticothalamic loops make the circuit act as a multilinear switching system without any interference between sequence components. The number of cortical units then scales linearly with the number of eigenvalues necessary to shape the motor output for each sequence component (Neigs), while the number of thalamic units scales linearly with the number of sequence components (Nmoves): a total of Neigs+Nmoves units. In contrast, a recurrent cortical network consisting of independent circuits with linearizable dynamics for each sequence component scales as Neigs* Nmoves. Thus, the corticothalamic architecture permits the learning of different sequence components independently while using units efficiently.

I-91. Dissecting the most influential features of neuronal activity on information and behavior

Ramon Nogueira[1] RNOGUEIRAMANAS@GMAIL.COM
Ruben Moreno Bote[2] RUBEN.MORENO@UPF.EDU
Nicole Peltier[3] NPELTIER@UR.ROCHESTER.EDU
Akiyuki Anzai[3] AANZAI@UR.ROCHESTER.EDU
Gregory DeAngelis[3] GDEANGELIS@UR.ROCHESTER.EDU
Julio Martinez-Trujillo[4] JULIO.MARTINEZ@ROBARTS.CA

[1]Columbia University
[2]Universitat Pompeu Fabra
[3]University of Rochester
[4]Western University

Identifying the statistical features of neuronal activity that determine information encoding is crucial for characterizing the link between brain activity and behavior1. Among all statistical features, mean pairwise correlations (MPC) and mean global activity (MGA) have been the best studied. Indeed, global activity modulations have been proposed as the mechanism by which information is modulated by attention2. On the other hand, results for the role of pairwise correlations have been mixed, sometimes deleterious3 to information coding and beneficial4 other times. Moreover, it has also been argued that the reduction of noise correlations in the visual cortex can have a positive impact on animal behavior in visual discrimination tasks5,6. However, the reported dependences between encoded information and behavior could be the result of additional unobserved variables that influence them. In the present study, we aim to uncover the hidden statistical variables of neuronal activity that directly affect the amount of encoded information and behavior. By deriving an analytical expression for the decoding performance of an optimal linear classifier, we identify the fundamental statistical features of neuronal activity that determine information coding: selectivity length (SL; distance between mean population activities) and projected precision (PP; inverse population noise projected on the tuning axis). We confirmed that SL and PP were the main features that affected encoded information in four datasets encompassing three tasks and two brain areas, by means of a conditioning bootstrap method that allows studying the effect of any single feature while controlling the rest. Surprisingly, behavioral performance was also dependent upon SL and PP, while MGA or MPC played little or no role, consistent with our predictions. Thus, we conclude that the previously observed dependences2-7 of information and behavior on MGA and MPC are likely due to their correlations with SL and PP.

I-92. Mechanisms of biased competition under balanced input: predictions from corticostriatal processing

Salva Ardid SARDID@BU.EDU
Jason Sherfey SHERFEY@BU.EDU
Michelle McCarthy MMCCART@BU.EDU
Joachim Hass JOACHIM.HASS@ZI-MANNHEIM.DE
Benjamin Pittman-Polletta BENPOLLETTA@GMAIL.COM
Nancy Kopell NK@MATH.BU.EDU

Boston University

Biased competition (BC) is a model of visual selective attention. According to BC, attention facilitates the processing of relevant information by suppressing irrelevant information. Mechanistically, BC relies on an input bias breaking the symmetry between competing neuronal ensembles of distinct selectivity, which are otherwise physiologically equal. Whether BC mechanisms can also be mediated between physiologically distinct ensembles receiving the same input remains to be explored. Corticostriatal processing fulfills two conditions to be an archetype of BC under such "balanced input". First, the striatal input-ouput transfer is accomplished by two neuronal ensembles, D1 and D2 dopamine receptor medium spiny neurons (MSNs), receiving balanced stimulation. Second, D1 and D2 MSNs differ in connectivity, and in synaptic and intrinsic dynamics. Uncovering the mechanisms of BC under balanced input is critical for understanding how corticostriatal coordination selects between the GO and the NO-GO pathways of the basal ganglia. Toward this end, we implemented a Hodgkin-Huxley circuit model of the striatum, which we constrained with the properties of D1 and D2 MSNs. We applied balanced asynchronous and rhythmic inputs to analyze the impact of input strength and frequency on the response of MSNs. This analysis identified three mechanisms of BC under balanced input: (i) rate mechanism: which MSN type shows higher firing rate depends on (low vs. high) input strength. (ii) rate vs. coherence mechanism: D2 MSNs show higher firing rate, but D1 MSNs fire more coherently, under low strength, rhythmic input. Here, two neural coding biases run in parallel, and each bias can be selectively decoded by adjusting the properties of the readout. (iii) coherence mechanism: each MSN type is more synchronized at a distinct input frequency, under high strength, rhythmic input. From these candidates, only the coherence coding mechanism is fully consistent with the current interpretation of rhythmic activity supporting rule-based decisions observed in prefrontal cortex.

I-93. LFADS: a deep learning technique to precisely estimate neural population dynamics on single trials

Chethan Pandarinath[1,2] CHETHAN.PANDARINATH@EMORY.EDU
Laurence F. Abbott[3] LFA2103@COLUMBIA.EDU
David Sussillo[4] SUSSILLO@GOOGLE.COM
Daniel O'Shea[5] DJOSHEA@STANFORD.EDU
Jasmine Collins[4] JLCOLLINS121@GMAIL.COM
Rafal Jozefowicz[4] RAFAL@OPENAI.COM
Sergey Stavisky[5] SERGEY.STAVISKY@STANFORD.EDU
Jonathan Kao[5] KAO@SEAS.UCLA.EDU
Eric Trautmann[5] ETRAUT@STANFORD.EDU
Matthew Kaufman[6] ANTIMATT@GMAIL.COM
Stephen Ryu[5] SEOULMAN@STANFORD.EDU
Leigh Hochberg[7] LHOCHBERG@MGH.HARVARD.EDU
Jaimie Henderson[5] HENDERJ@STANFORD.EDU
Krishna Shenoy[5] SHENOY@STANFORD.EDU

[1]Emory University
[2]Georgia Tech
[3]Columbia University

[4]Google Brain
[5]Stanford University
[6]Cold Spring Harbor Laboratory
[7]Massachusetts General Hospital

Neuroscience is experiencing a data revolution in which simultaneous recording of many hundreds or thousands of neurons is revealing structure in population activity that is not apparent from single-neuron responses. This structure is typically extracted from trial-averaged data. Single-trial analyses are challenging due to incomplete sampling of the neural population, trial-to-trial variability, and fluctuations in action potential timing. Here we demonstrate Latent Factor Analysis via Dynamical Systems (LFADS), a deep learning method to infer latent dynamics from single-trial neural spiking data. LFADS uses a nonlinear dynamical system (a recurrent neural network) to infer the dynamics underlying observed population activity and to extract 'de-noised' single-trial firing rates from neural spiking data. We apply LFADS to a variety of monkey and human motor cortical datasets, demonstrating its ability to predict observed behavioral variables with unprecedented accuracy, extract precise estimates of neural dynamics on single trials, infer perturbations to those dynamics that correlate with behavioral choices, and combine data from non-overlapping recording sessions (spanning months) to improve inference of underlying dynamics. In summary, LFADS leverages all observations of a neural population's activity to accurately model its dynamics on single trials, opening the door to a detailed understanding of the role of dynamics in performing computation and ultimately driving behavior. (Preprint available: Pandarinath et al., bioRxiv: http://doi.org/10.1101/152884)

I-94. A diffusive forward model for motor planning supports confidence-based hierarchical decision-making

Nicolas Meirhaeghe[1,2] NMRGHE@MIT.EDU
Mehrdad Jazayeri[1] MJAZ@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Harvard-MIT Division of Health Sciences and Technology

When Tiger Woods eventually misses after a long winning streak, he has to quickly decide whether the wind has shifted or whether his swing was off in order to make the appropriate adjustments for the upcoming stroke. Whether people can solve this type of credit assignment problem is unclear, as it would imply that they evaluate their motor variability on a trial-by-trial basis. To investigate this question, we designed a task which required subjects to produce a desired time interval in one of two covert environments. Due to unannounced transitions between environments, participants had to leverage their motor confidence to attribute failures to either timing imprecision or incorrect assumption about the environment. We find that subjects switch environment more often after making small rather than large timing errors. Switches are also more likely when the cumulative error magnitude across consecutive failures is small. These results suggest that humans monitor their trial-by-trial motor variability and use confidence across trials to update their behavior. This raises an intriguing computational question: how can the brain infer the magnitude of error, which is an analog variable, from feedback that is binary? Building on ideas developed in motor control, we hypothesized that this self-evaluation is made possible by an internal simulator (i.e., forward model). Assuming that the drift rate of a diffusion-to-bound process controls motor timing, we model the simulator as a parallel diffusion process that receives the same drift input (efference copy). We show that the output of the simulator at the time of the motor response provides a measure of trial-by-trial error and can be used to estimate confidence. Forward models are implicated in various sensorimotor functions such as state estimation, mental simulation, and sensory prediction. Our work reveals a potential new function for the forward model: estimating motor confidence.

I-95. Neuronal populations supporting vision, action, and reward across the mouse brain

Nicholas Steinmetz NICK.STEINMETZ@GMAIL.COM
Peter Zatka-Haas P.ZATKA@UCL.AC.UK
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

Behavior arises from neuronal activity patterns, but whether the relevant activity is private to a small number of brain regions, as typically studied, or instead distributed and coordinated widely across many regions, remains unknown. To answer this question, we used Neuropixels electrode arrays acutely in head-fixed mice performing a visually-guided decision task to measure the task representations of $>$20,000 neurons across $>$40 brain regions with millisecond temporal resolution.

We trained mice to perform a visually-guided perceptual decision task (Burgess et al., Cell Reports 2017). In this task, mice were trained to give one of three responses (choose left, right, or neither) depending on the relative contrast of two simultaneously presented visual stimuli. We recorded the activity of thousands of neurons during task performance using multiple acutely-inserted Neuropixels probes (Jun et al., Nature 2017), whose output was processed with Kilosort (Pachitariu et al, NIPS 2016). The arrays span ~4 mm of tissue and thus record simultaneously across diverse brain regions. By histologically reconstructing electrode locations, we established that the recorded neurons were in $>$40 brain regions, including: sensory, parietal, frontal, and motor isocortex; thalamic nuclei; hippocampus; the superior and inferior colliculi; the striatum; and midbrain structures.

To understand how this distributed neuronal activity related to multiple aspects of task performance, we fit the activity of each neuron as a sum of temporal kernels triggered on each of the task components (visual stimulus onset, auditory go cue, response initiation, reward delivery, etc). We detected prominent visual responses in superior colliculus, visual cortex, and striatum. However, representations of action execution and reward were distributed broadly, and were observed in nearly every region we recorded (for example sensory and motor cortex, superior colliculus, and multiple thalamic regions). These highly distributed representations suggest that information pertaining to actions and rewards pervades essentially the entire brain.

I-96. Heterogeneous mossy fiber activity patterns and their implications for sensorimotor encoding in the cerebellar cortex

Hana Ros H.ROS@UCL.AC.UK
Sadra Sadeh S.SADEH@UCL.AC.UK
Alex Cayco-Gajic NATASHA.GAJIC@UCL.AC.UK
R. Angus Silver A.SILVER@UCL.AC.UK

University College London

The brain gathers information about the body and the surrounding world, enabling it to build internal representations and to plan and execute movement. The cerebellum is thought to predict the sensory consequences of movements and coordinate movement by learning sensorimotor relationships. Cerebellar mossy fiber (MF) inputs convey a wide range of sensory and motor related information that is integrated by granule cells (GCs). But little is known about how populations of MFs encode sensory and motor signals locally within the input layer. To address this we used adeno-associated viruses to express the genetically encoded calcium indicator GCaMP6f in distinct precerebellar nuclei, implanted a chronic window over Crus I/II and vermis of the cerebellar cortex and performed two-photon (2P) imaging of MFs in awake behaving mice. Since MF synaptic rosettes are sparsely distributed, we used high speed 3D 2P Acousto-Optic Lens (AOL) microscopy to record their activity within a 250 x 250 x 250 μm imaging volume. We observed a wide range of activity patterns across MFs, with individual MFs exhibiting either an increase or a decrease in activity with locomotion. Surprisingly, positively and negatively modulated

MFs were often observed within the same local region (i.e. 10 - 100 μm), suggesting that individual GCs could be innervated by functionally opposed inputs that cancel out. To understand how such mixed local populations of bidirectionally modulated MFs affect population coding in this circuit, we analytically calculated the linear Fisher information in randomly connected feedforward MF-GC networks, either with bidirectionally or unidirectionally modulated inputs. We show that opposite signed MFs can counteract input noise correlations to enable better signal propagation in GCs. Our results suggest that MFs utilise a population code that conserves sensorimotor information by counteracting the deleterious effects of noise correlations.

I-97. Dopamine reward prediction errors are modulated by an internal bias during stimulus discrimination

Nestor Parga[1] NESTOR.PARGA@UAM.ES
Stefania Sarno[1] STEFANIA.SARNO@UAM.ES
Manuel Beiran[2] MANUEL.BEIRAN@ENS.FR
Jose Vergara[3] JVERGARA@EMAIL.IFC.UNAM.MX
Roman Rossi-Pool[3] ROMANR@IFC.UNAM.MX
Ranulfo Romo[3] RROMO@IFC.UNAM.MX

[1]Universidad Autonoma de Madrid
[2]Ecole Normale Superieure
[3]Universidad Nacional Autonoma de Mexico

Under uncertain stimulation conditions dopamine (DA) responses to relevant task cues reflect cortical perceptual decision-making processes, such as the certainty about stimulus detection (de Lafuente and Romo, 2011) and evidence accumulation (Nomoto et al., 2010), in a way compatible with the reward prediction error (RPE) hypothesis (Sarno et al., 2017). This suggests that the midbrain DA system receives information from cortical circuits about decision formation and transforms it into a RPE signal. If so, DA phasic activity should reflect a variety of phenomena, including internally generated biases. This is because biases influence decisions and performance and hence they are expected to modulate the error in the prediction of reward. To test this hypothesis and acquire further insight into how DA neurons behave under uncertainty we used the two-interval, two-alternative forced-choice paradigm. These tasks present a contraction bias whereby the sensory feature of the first stimulus is perceived as if its value were shifted to the center of its range. Specifically: 1.We recorded DA neurons in monkeys discriminating the frequency of two vibrotactile stimuli. 2.Although naively the response to the first stimulus should only code the predicted average reward, it was modulated (but not tuned) by the value of the frequency in the way expected from the bias. 3.Similarly, the response to the second stimulus depended on the stimulus pair in a way consistent with the bias. 4.The activity during the comparison period was modulated by the subjective difficulty of the task, defined using a Bayesian model for the choice. The model explains the contraction bias and gives a measure of the animal's choice confidence. 5.The DA activity was above baseline throughout the delay (memory) period. It was neither tuned nor modulated by the first frequency, pointing to quite different roles of delay and phasic activities.

I-98. Learning through recurrent dynamics

Florent Meyniel[1,2] FLORENT.MEYNIEL@GMAIL.COM
Alexandre Pouget[3] ALEXANDRE.POUGET@UNIGE.CH

[1]CEA
[2]Saclay
[3]University of Geneva

Learning in a world that is both changing and noisy is a difficult problem. On the one hand, since our observations

are subject to uncertainty, one must average many observations in order to learn a stable and accurate estimates of the quantities of interest (e.g. a reward rate). On the other hand, since the learned quantities may change suddenly (e.g. at change-points), one must give more weight to recent observations in order to quickly learn the new value. This can be learned with a delta rule as long as the learning rate increases transiently after change points. Previous studies have established that human behavior is indeed consistent with such a learning rule. This approach however suffers from a major limitation: it is far from clear how the confidence about the learned quantity can be accessed to by the rest of the network if it is buried in the synaptic weights. Yet, behavioral studies show that human subjects can report confidence in a way consistent with the optimal Bayesian solution to this problem. Here, we investigated whether a recurrent neural network (RNN) could learn this optimal solution approximately. We trained a RNN to predict the next observation in binary sequences whose generative probability of items undergoes change-points. Once trained, we froze the weights and tested the ability of the frozen network to learn the time varying generative probability of new series of observations. Our results show that the network behaves as if it were using a delta rule with a learning rate that increases after change points, even though the synaptic weights are frozen. In addition, we found that we could recover the confidence about the generative probability from the hidden units' activity. This demonstrates that Bayesian learning can be implemented in recurrent dynamics without any need for synaptic changes.

I-99. Transformation of population code from LGN to V1 facilitates linear decoding

Alex Cayco-Gajic[1] NATASHA.GAJIC@UCL.AC.UK
Severine Durand[2] SEVERINED@ALLENINSTITUTE.ORG
Michael Buice[2] MICHAELBU@ALLENINSTITUTE.ORG
Clay Reid[2] CLAYR@ALLENINSTITUTE.ORG
Eric Shea-Brown[3] ETSB@UW.EDU
Joel Zylberberg[4] JOEL.ZYLBERBERG@GMAIL.COM

[1]University College London
[2]Allen Institute for Brain Science
[3]University of Washington
[4]University of Colorado

How neural populations represent external stimuli, and how that representation is transformed from one brain area to another, are fundamental questions of neuroscience. In particular, dorsolateral geniculate nucleus (dLGN) and primary visual cortex (V1) represent distinct stages of early mammalian vision that are thought to use vastly different coding strategies. However, comparisons between dLGN and V1 population coding are limited by the fact that recordings are generally done in different animals, where changes in spiking activity could be due to changes in experimental variables including behavioural state or attention. To address this, we simultaneously recorded the spiking activity of mouse dLGN and V1 in vivo. V1 populations were sparser and more stimulus selective than dLGN populations, confirming that the early visual pathway transforms population activity that is more distributed across neurons to population activity that is sparser and more specific to stimulus features.

We next measured the spiking correlations of dLGN or V1 unit pairs. Correlations can have a strong impact on population coding, yet their presence or absence in cortex has been disputed. We found that V1 unit pairs were on average more correlated than dLGN pairs. Surprisingly, correlations greatly improved decoding performance in dLGN but had little effect in V1. To understand this counterintuitive result, we show that the performance of a correlated decoder (i.e., one that is sensitive to correlations) can be reduced to a quadratic separator, and the performance of an independent decoder to a linear separator. Indeed, in dLGN a quadratic separator was far more accurate than a linear separator, while in V1 they performed similarly. This agrees with recent work theorizing that the visual stream processes information to make higher-order features more linearly separable. Therefore, we find that the transformation of the correlated structure from dLGN to V1 populations facilitates downstream linear decoding.

I-100. Hebb 'n' Dale: efficient coding by time-reversible dynamics in recurrent circuits

Alberto Bernacchia[1] AB2349@CAM.AC.UK
Jozsef Fiser[2] FISERJ@CEU.EDU
Guillaume Hennequin[1] G.HENNEQUIN@ENG.CAM.AC.UK
Mate Lengyel[1] M.LENGYEL@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Central European University

The efficient coding hypothesis has accounted for many aspects of neural responses in sensory areas. However, most previous work focused on information maximization about static stimuli in networks with feedforward architectures and thus lacking internal dynamics. It remains unclear in which dynamical regime neurons should collectively operate to best represent time-varying stimuli. Here, we identify a class of network dynamics ideally suited for encoding the past history of a dynamically changing stimulus in a recurrent neural circuit. We demonstrate in simulations and prove analytically that information transmission is maximized in recurrent networks that have time-reversible dynamics when stimulus statistics are themselves time-reversible. Are such dynamics compatible with neurobiological constrains? Surprisingly, we show that recurrent circuits naturally self-organize for time-reversibility under a biological form of spike timing-dependent plasticity (STDP), and that the resulting circuit connectivity obeys the well known principles of Hebb and Dale: synaptic weights become proportional to correlations between pre- and postsynaptic activity, and neurons eventually fall into two distinct classes of excitatory and inhibitory cells. Finally, we identify signatures of adaptive time reversible circuit dynamics in experimental data. In the primary visual cortex of awake ferrets, we find that neural activity has a very small irreversible component; Furthermore, activity is initially irreversible when stimulus statistics change, but time reversibility increases with continued stimulus exposure on a timescale of a few minutes. Our work suggests a central role for time reversibility as a novel signature of efficient coding in recurrent circuits for dynamic stimuli, and provides the first joint normative account of Hebb and Dale's principles.

I-101. Blind source separation emerges in layer 2/3 from STDP

Wolfgang Maass MAASS@IGI.TUGRAZ.AT
Robert Legenstein LEGENSTEIN@IGI.TUGRAZ.AT

Graz University of Technology

This research is part of the program to elucidate the function of the complex architecture of cortical columns by analyzing the computational operations that emerge through STDP in the most important microcircuit motifs of cortical columns. We are addressing here one of the most prominent motifs: Interconnected populations of pyramidal cells and parvalbumin-positive inhibitory cells in layer 2/3. Experimental studies suggest that these inhibitory neurons impose divisive inhibition on the pyramidal cells. We show that this form of feedback inhibition, which is softer than that of the commonly considered winner-take-all (WTA) models, contributes to the emergence of an important computational function in layer 2/3 through STDP: The capability to disentangle superimposed ?ring patterns in upstream networks, and to represent their information content through a sparse assembly code (blind source separation).

I-102. Multiple timescales of adaptation in mouse primary somatosenory and visual cortices

Kenneth Latimer[1] LATIMER1@UW.EDU
Michael Sokoletsky[2] MICHAEL.SOKOLETSKY@WEIZMANN.AC.IL
Dylan Barbera[3] DYLAN.BARBERA@UTMAIL.UTEXAS.EDU
Nicholas Priebe[3] NICO@AUSTIN.UTEXAS.EDU
Ilan Lampl[2] ILAN.LAMPL@WEIZMANN.AC.IL
Adrienne Fairhall[1] FAIRHALL@UW.EDU

[1]University of Washington
[2]Weizmann Institute of Science
[3]The University of Texas at Austin

Through adaptation, neural responses can show a dependence on the statistical context of the input. We aim to uncover the rules that govern how adaptation shapes cortical responses to temporal sequences, and determine whether the resulting dynamics are shared or distinct in different cortical areas. To this end, we examined adaptation in both visual and somatosensory systems of the mouse using a common framework. We constructed stimuli consisting of sequences of 20~ms pulses, delivered as light flashes or whisker deflections, while whole cell recordings were made from cells in visual cortex (V1) or somatosensory cortex (S1). We characterized each neuron's response properties with a set of stochastic stimuli with following a homogeneous Poisson process with rates ranging from 0.5-20 pulses/s. At low pulse frequencies, neurons in both areas exhibit a biphasic response, with a second peak occurring over 300~ms after the pulse. With increased pulse frequency, the responses to a single pulse become more monophasic, and many cells show large offset responses. We show that a model of the membrane potential consisting of the sum of a small number of simple linear-nonlinear subunits accounts for the observed responses by capturing a diversity of timescales of adaptation. Additionally, we compare the adaptive dynamics observed in cortex to LGN responses to the same set of stimuli. Because our model uses a single set of linear filters across all stimulation frequencies, there is no need to invoke specialized adaptation mechanisms in the model. Instead these adaptive changes simply reflect a piece of the multidimensional, nonlinear response to sensory stimulation.

I-103. Hippocampal sequences and model-based planning in the rat

Kevin J. Miller[1] KJMILLER@PRINCETON.EDU
Sarah Jo C. Venditto[1] VENDITTO@PRINCETON.EDU
Nathaniel Daw[1] NDAW@PRINCETON.EDU
Matthew Botvinick[2,3] BOTVINICK@GOOGLE.COM
Carlos Brody[4,1] BRODY@PRINCETON.EDU

[1]Princeton University
[2]Google DeepMind
[3]University College London
[4]Howard Hughes Medical Institute

Humans and animals construct internal models of the world around them, and use these models to guide behavior. Such model-based cognition is often referred to as "planning", and its neural mechanisms remain poorly understood. Planning has been proposed to depend on hippocampal sequences, in which place cells "sweep out" trajectories through space while an animal is at rest (Foster & Knierim, 2012; Mattar & Daw, 2017). Research into the role of sequences in planning, and into the neural mechanisms of planning in general, has been hampered by a lack of tasks for animals which both demonstrably elicit planning behavior and are suitable for neural recordings. In recent work, we have lifted this limitation, adapting advances from work with humans (Daw, et al., 2011) to develop a multi-step decision task for rats, and showing that they adopt a planning strategy which depends on neural activity in the dorsal hippocampus (Miller, Botvinick, & Brody, 2017). Here, we report the results of elec-

trophysiological recordings made in dorsal hippocampus during planning behavior. We find that individual cells encode the states of the task, and that hippocampal sequences take place during sharp wave ripple events at the conclusion of most trials. The content of these sequences reflects knowledge of the structure of the task, consistent with a role in model-based planning. Ongoing work seeks to characterize the relationship between sequence content and behavior at a trial-by-trial level.

I-104. Olfactory processing by a precisely balanced network

Peter Rupprecht PETER.RUPPRECHT@FMI.CH
Rainer Friedrich RAINER.FRIEDRICH@FMI.CH

Friedrich Miescher Institute

The concept of balanced networks provides an important framework for studies of cortex and other neural circuits. In the classical balanced state, excitatory and inhibitory input currents to individual neurons are strong but uncorrelated, thus causing chaotic membrane potential fluctuations that are thought to be inefficient for stimulus encoding. More recently, theoretical studies have shown that correlating excitatory and inhibitory inputs on a short timescale (tight balance) and in a multi-dimensional coding space (detailed balance) can neutralize these drawbacks of balanced networks. We use whole-cell voltage-clamp recordings in the intact zebrafish brain to directly analyze synaptic inputs to neurons in telencephalic area Dp, the homolog of olfactory cortex. Local silencing of activity by injection of the GABAA-agonist muscimol confirmed that Dp neurons receive strong inputs from recurrent connections within Dp. During an odor response, Dp neurons exhibited the hallmarks of a balanced state: 1) strong and balanced excitatory and inhibitory inputs; and 2) a large synaptic conductance relative to the resting conductance. Using the odor-evoked 20 Hz oscillation that originate in the olfactory bulb as a reference clock, we aligned excitatory and inhibitory inputs recorded in Dp and found a tight balance, with inhibition tracking excitation by a 3 ms delay. Finally, by studying the odor-specificity of excitatory and inhibitory inputs for a set of odor stimuli, we found inhibition and excitation to be co-tuned. These findings could not be explained by a purely random network, showing that excitation and inhibition exhibit a detailed, high-dimensional balance in stimulus space. Together, our experimental results show that Dp enters a balanced state during an odor response that is precise (= tight and detailed). We propose that this network is a substrate for a pattern classification process that is fast, as in classical balanced networks, but also stable in many coding directions.

I-105. Presynaptic inhibition provides a rapid stabilization of recurrent excitation in the face of plasticity

Laura Bella Naumann[1] LAURAN@BCCN-BERLIN.DE
Henning Sprekeler[2] H.SPREKELER@TU-BERLIN.DE

[1]Technical University Berlin
[2]Berlin Institute of Technology & Bernstein Center for Computational Neuroscience

Hebbian plasticity, a mechanism believed to play a key role in learning and memory, detects and further enhances correlated neural activity. In recurrent networks this constitutes an inherently unstable positive feedback loop and therefore requires additional homeostatic control. Recent computational work indicates that slow homeostasis, as observed in experiments, is insufficient to compensate the instabilities arising from Hebbian plasticity in recurrent neural networks.

We suggest presynaptic inhibition as a compensatory process, which does not suffer from this discrepancy of timescales. Experimental studies have revealed that excess network activity can trigger inhibition of transmitter release at excitatory synapses through activation of presynaptic GABAB receptors. This effectively and reversibly attenuates recurrent interactions on timescales of 100s of milliseconds, thus serving as a candidate mechanism for the rapid compensation of elevated recurrent excitation induced by Hebbian changes.

To study the network effects of presynaptic inhibition, we analyzed a rate-based recurrent network model, in which presynaptic inhibition is mimicked by a multiplicative reduction of recurrent synaptic weights in response to increasing firing rates. Using analytical and numerical methods, we show that presynaptic inhibition ensures a gradual increase of firing rates with growing recurrent excitation, even for very strong recurrence, whereas classical subtractive postsynaptic inhibition is unable to control recurrent excitation once it has surpassed a critical strength. Moreover, we find that presynaptic inhibition stabilizes firing rates in a recurrent population subject to Hebbian plasticity, while allowing synaptic homeostasis to operate on biologically plausible timescales.

In summary, the multiplicative character of presynaptic inhibition provides a powerful compensatory mechanism to rapidly reduce effective recurrent interactions. As it conserves the underlying network connectivity, presynaptic inhibition might therefore set the stage for stable learning without interfering with plasticity at the level of single synapses.

I-106. Stabilizing the grid cell representation by coupling modules through recurrent synaptic connectivity

Noga Mosheiff NOGA.WEISS@MAIL.HUJI.AC.IL
Yoram Burak YORAM.BURAK@ELSC.HUJI.AC.IL

Hebrew University of Jerusalem

Grid cells in the entorhinal cortex encode the position of an animal in its environment with spatially periodic tuning curves of varying periodicity. Recent experiments established that these cells are functionally organized in discrete modules with uniform grid spacing. Furthermore, several experiments support a theoretical proposal, that within each module neural activity is constrained by recurrent connectivity to lie within a two-dimensional continuous attractor. Yet, not much is known about synaptic connectivity between cells from different modules. Here we argue that coupling between grid cell modules is essential in order to maintain the code stability, in the absence of sensory cues that inform the animal about its absolute position. The state of each module might be perturbed by noise that arises intrinsically within the module or by noise in the velocity inputs, leading to gradual drift in the represented position in each module. However, to avoid catastrophic readout errors and obtain a continuous joint representation of position over time, the drifts in different modules must be compatible. Here we develop a theory of coupled grid cell modules, each modeled as a continuous attractor network. We identify a simple scheme to approximately read out the drift velocity in each module. Using this scheme we propose a way to couple the drift velocities of different modules which can be implemented by plausible neural circuitry. As a result of the coupling, the activities of the different modules shift together, with a relative chosen velocity ratio that defines the grid spacing. This method eliminates the relative drift driven by the external input, and reduces the relative drift driven by the neuronal noise. Our results suggest that recurrent connectivity between grid cells belonging to different modules may help stabilize the grid cell representation of position.

I-107. Local synaptic control of global error signals permits gradient-free learning and continual circuit reconfiguration

Dhruva Raman DVR23@CAM.AC.UK
Timothy O'Leary TSO24@CAM.AC.UK

University of Cambridge

Learning can be considered as a feedback loop between behavioural performance and modifiable properties of a neural circuit. Once expert performance has been achieved, one might hypothesise a neural circuit reaches an optimal configuration, which is maintained over task-relevant timescales, and possibly indefinitely. However, this is inconsistent with recent experimental evidence from several cortical structures that show constant reconfiguration

over time during behaviour, post-learning (Driscoll et al 2017, Cell; Ziv et al 2013, Nature Neuroscience). These results imply the existence of ongoing plasticity processes unconnected with the task, and/or noise-related changes in circuit properties. How can learning be achieved and maintained in the face of these ongoing disturbances, and how might this constrain synaptic learning mechanisms? Starting from biologically plausible assumptions we consider how these recent observations constrain synaptic dynamics in a circuit. We derive a parsimonious class of learning rules that obey a so-called 'Error Control Principle' (ECP) that pairs a scalar, global reinforcement signal with local synaptic dynamics. We show analytically that plasticity rules obeying ECP allow learning to occur in the face of a wide class of disturbances that can be modelled as independent identically distributed noise processes in each weight. Extremely crude implementations of these learning rules are found to be sufficient for learning to occur, indicating compatibility with known synaptic plasticity rules that biological systems can easily implement. We demonstrate numerically that a multilayer network can learn classification tasks using several versions of the rule, which by construction does not require computation of an error gradient. Qualitatively, single unit response properties in simulations mimic the dynamic reconfiguration observed in experiments and indicate that learning may be difficult to observe in patterns of neural activation, even when performance improves dramatically.

I-108. Probabilistic inference emerges from learning in neural circuits with a cost on reliability

Laurence Aitchison LAURENCE.AITCHISON@GMAIL.COM
Guillaume Hennequin G.HENNEQUIN@ENG.CAM.AC.UK
Mate Lengyel M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

Neural responses in the cortex change over time both systematically, due to ongoing plasticity and learning, and seemingly randomly, due to various sources of noise and variability. Most previous work considered each of these processes, learning and variability, in isolation – here we study neural networks exhibiting both and show that their interaction leads to the emergence of powerful computational properties. We trained neural networks on classical unsupervised learning tasks, in which the objective was to represent their inputs in an efficient, easily decodable form, with an additional cost for neural reliability which we derived from basic biophysical considerations. Penalising reliability introduced a tradeoff between energetically cheap but inaccurate representations and energetically costly but accurate ones. Despite the learning tasks being non-probabilistic, the networks solved this tradeoff by developing a probabilistic representation: neural variability represented samples from statistically appropriate posterior distributions that would result from performing probabilistic inference over their inputs. We provide an analytical understanding of this result by revealing a connection between the cost of reliability, and the objective for a state-of-the-art Bayesian inference strategy: variational autoencoders. We show that the same cost leads to the emergence of increasingly accurate probabilistic representations as networks become more complex, from single-layer feed-forward, through multi-layer feed-forward, to recurrent architectures. Our results provide insights into why neural responses in sensory areas show signatures of sampling-based probabilistic representations, and may inform future deep learning algorithms and their implementation in stochastic low-precision computing systems.

I-109. Electrical synapses modulate the responses of principal neurons to transient inputs: a modeling study of cortical discrimination

Tuan Pham TAP218@LEHIGH.EDU
Julie Haas JUH312@LEHIGH.EDU

Lehigh University

As multimodal sensory information proceeds to the cortex, it is intercepted and processed by the nuclei of the thalamus. The main source of inhibition within thalamus is the reticular nucleus (TRN), which collects signals both from thalamocortical relay neurons and from thalamocortical feedback. Within the reticular nucleus, neurons are densely interconnected by connexin36-based gap junctions, known as electrical synapses. Electrical synapses have been shown to coordinate neuronal rhythms, including thalamocortical spindle rhythms, but their role in shaping or modulating transient activity as is propagates through the brain is less understood.

We constructed a four-cell single-compartment Hodgkin-Huxley model comprising thalamic relay and TRN neurons, and used it to investigate the impact of electrical synapses on closely timed inputs delivered to thalamic relay cells. We showed that the electrical synapses of the TRN assist in cortical discrimination of sensory inputs through effects of truncation, delay or inhibition of thalamic spike trains. We showed that increasing electrical synapse strength leads to increased thalamocortical spiking separation and independence when inputs are dissimilar and temporarily separated, through synergy of electrical synapses acting through inhibition. On the other hand, electrical synapses result in fusion of thalamocortical spiking, masking small strength and/or temporal differences, for similarly timed or sized inputs. We expect that these are principles whereby electrical synapses play similar roles in regulating the processing of transient activity in excitatory neurons across the brain. Our simulations provide specific predictions regarding the impact of electrical synapses and plasticity in thalamocortical processing, which we are currently testing in experiments in vitro.

I-110. Training recurrent networks of excitatory/inhibitory spiking or rate neurons

Alessandro Ingrosso AI2367@COLUMBIA.EDU
Laurence F. Abbott LFABBOTT@COLUMBIA.EDU

Columbia University

The construction of progressively more biologically plausible functional models of neural circuits is crucial for understanding the computational properties of the nervous system. Learning in networks composed of separate excitatory and inhibitory neurons observing Dale's law presents a number of challenges. On the one hand, there are technical difficulties involved in training sign constrained synaptic efficacies. In addition, dynamic instabilities are substantially more prominent than in the unconstrained case. We show how a simple target-based approach, when coupled with a fast online constrained optimization technique, is capable of building functional models of rate and spiking recurrent neural networks composed of excitatory and inhibitory units. The trained networks can reproduce complicated temporal patterns and solve simple input-output tasks, while retaining biologically desirable features such as Dale's law and low mean firing rates. In our framework, a Teacher network composed of two populations is driven in such a way that it generates target dynamical trajectories to be learned by a Student rate or spiking model (as in [DePasquale et al., 2017]). Suppression of chaotic dynamics in the Teacher network is required to produce neural trajectories that can be autonomously stabilized in the Student network. The suppression of chaos in excitatory/inhibitory networks shows some peculiar features with respect to temporal coherence of external drive across neurons that make learning hard. Moreover, activity correlations, reflected in the time-dependent mean firing rates of the network, can also impede learning. Finally, if unbounded activation functions (e.g. ReLu) are used, global runaway behavior has to be overcome. We show how several of these problems can be alleviated to generate targets and harness the computational capabilities of recurrent circuits for learning of stable trajectories.

I-111. Cellular noise amplified by chaotic network dynamics drives high intrinsic variability in cortex

Max Nolte MAX.NOLTE@EPFL.CH
Michael Reimann MICHAEL.REIMANN@EPFL.CH
James King JAMESGONZALO.KING@EPFL.CH
Henry Markram HENRY.MARKRAM@EPFL.CH
Eilif Muller EILIF.MUELLER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

The respective impact of cellular sources of noise and chaotic network dynamics on neural variability is not known. To unravel how chaotic network dynamics and different cellular sources of noise contribute to intrinsic cortical variability, we simulated electrical activity in a biophysically-detailed model of a neocortical microcircuit (NMC) (Markram et al., 2015). The NMC-model contains stochastic synapses and stochastic ion-channels, and also exhibits chaotic network dynamics with a naturally emerging balance of excitation and inhibition. Without any sources of noise, injected extra spikes and small sub-threshold perturbations to soma membrane voltages make network activity fully deviate from the same network state within hundreds of milliseconds. When turning on cellular noise sources such as synaptic failure, spontaneous release and stochastic ion-channels, the network deviates from the same network state within tens of milliseconds. We find that synaptic failure is the dominant source of noise, and provide evidence that other, weaker cellular sources of noise, such as ion-channel noise, are negligible for network variability as long as they are weaker than synaptic noise.

I-112. Serotonergic modulation of moving objects in the electrosensory system

Mariana Marquez Machorro MARIANA.MARQUEZMACHORRO@MAIL.MCGILL.CA
Maurice J. Chacron MAURICE.CHACRON@MCGILL.CA

McGill University

Sensory systems are able to adjust to highly dynamic changes in the external and internal environment. Although a complete understanding of the mechanisms underlying these adjusting properties is missing, feedback loops and neuromodulators are known to be crucial. For instance, serotonin has been shown to regulate sensory systems by adjusting responsiveness to stimuli. This regulatory effect has been observed in sensory systems such as auditory and visual systems [1, 2], therefore raising the question: is serotonin control ubiquitous in sensory systems? The electrosensory system in the weakly electric fish provides a well-established model system that shares similarities across vertebrates and where presence of serotonin has been extensively described[3]. This system is well known to accomplish object location (such as rocks, prey or fish) as well as conspecific communication functions. In fact, previous studies have shown that, by promoting stimulus processing through sequences of action potentials at high frequency rate of discharge (i.e. burst spikes), serotonin enhances neuronal responses to communication signals[4]. In order to investigate if serotonin exerted regulatory effect in the electrosensory system during electrolocation, we recorded pyramidal neurons within the electrosensory lateral line lobe (ELL) in the hindbrain. Briefly, these neurons receive input from primary afferents located on the fish's skin that transduce deflections of the self-generated electric field giving rise to object perception. Moreover, it has been shown that burst spiking activity of those neurons encodes changes in direction of movement[5]. We performed in vivo extracellular recordings combined with pharmacological manipulations while presenting object looming and receding motion. We found that serotonin enhances object detection and that this enhancement is encoded by burst spiking activity.

I-113. Sparse but smooth high-dimensional representations of visual stimuli

Carsen Stringer[1] CARSEN.STRINGER@GMAIL.COM
Marius Pachitariu[1] MARIUS10P@GMAIL.COM
Matteo Carandini[2] MATTEO@CORTEXLAB.NET
Kenneth Harris[2] KENNETH.HARRIS@UCL.AC.UK
Charu Reddy[2] C.REDDY@UCL.AC.UK

[1]HHMI Janelia Research Campus
[2]University College London

To investigate how large neuronal populations encode sensory stimuli, we imaged the responses of ~10,000 neurons to 2,800 images, in visual cortex of awake mice using two-photon calcium imaging. The evoked neural activity was full dimensional: population activity vectors could not be constrained to a lower dimensional subspace. Unexpectedly however, the geometry of neural activity obeyed a power-law: the variance along its n-th dimension scaled as 1/n. This power law grew more accurate as more neurons and stimuli were considered, suggesting it represents a fundamental feature of the cortical code, rather than the specific neurons or stimuli of any one experiment. The 1/n spectrum was not inherited from the 1/f spectrum of natural images, because it persisted for stimuli without this property. However, the spectrum changed when the dimensionality of the visual stimulus ensemble varied, decaying as a power law with a larger exponent when inputs were drawn from a lower dimensional distribution. Mathematically, one can show that the spectrum of a smooth (i.e. noise-resistant) stimulus representation can decay no slower than a power law of exponent 1+2/d, where d is the input dimensionality. Our results suggest that the neural activity approaches this bound, and thus that the cortical code may be as high-dimensional and sparse as possible while maintaining smooth representations of visual stimuli.

I-114. Simultaneous accessibility of multiple time scales of evidence integration for choice and confidence

Preetham Ganupuru PGANUPURU@UCDAVIS.EDU
Adam Goldring ADAMGOLDRING@GMAIL.COM
Rashed Harun RHARUN@UCDAVIS.EDU
Timothy Hanks THANKS@UCDAVIS.EDU

University of California, Davis

A sequential sampling framework has been suggested to provide a single account for both choices and one's confidence that those choices are correct during perceptual decision making tasks. To understand the neural mechanism that supports this process, it is critical to characterize how and when sensory information influences both of these components of decision making. Recent work has shown that apparent dissociations between choices and confidence in discrimination tasks involving near perfect integration of sensory information can be modeled as a common process that takes into account temporal delays in information processing for each. 1 Here, we extend this line of inquiry to a change detection task where optimal behavior does not involve perfect integration of sensory information. In brief, human subjects are trained to respond when a stochastic click stimulus increases its generative rate, or to withhold a response for the full trial if there is no change. 2 In both cases, subjects then report confidence of being correct. We find that larger changes in click rate are associated with both improved detection rates and higher confidence reports. Furthermore, using psychophysical reverse correlation, we show that the time period over which sensory information has leverage differs based on what's being reported (i.e. choice vs. confidence) as well as on trial outcome (i.e. subject reports a change vs. subject does not report a change). Strikingly, confidence judgments can be influenced by evidence both preceding and following the time period of influence for detection reports. These findings suggest multiple time scales of evidence integration are simultaneously accessible for use in different components of decision making. Our next step will be to develop computational models of this and test potential neural implementations.

I-115. Fast learning without forgetting by synaptic consolidation

Pascal Leimer LEIMER@PYL.UNIBE.CH
Walter Senn SENN@PYL.UNIBE.CH

University of Bern

For some synaptic plasticity protocols the induced change in the synaptic strength decays back to baseline within minutes or it decays partially before it reaches a long-term value. These transients are referred to as the early phase of long-term plasticity and can be observed both for long-term potentiation (LTP) and for long-term depression (LTD).

We hypothesize that the distinction between decaying, partially decaying and non-decaying weights is beneficial for learning in an environment where one context follows another. In different contexts, similar sensory stimuli may require different responses. Consolidating all synaptic weight changes acquired in one context may therefore cause learning interferences in a subsequent context and lead to catastrophic forgetting of previously learned stimulus-response associations.

We suggest that early LTP and its consolidation process is one of nature's choices to deal with the plasticity-stability dilemma. The network should only consolidate a minimum number of synapses, but still be able to change enough synapses if required by fast learning. We show that synaptic plasticity with early LTP/LTD and an additional synaptic tag mitigates this conflict. In our model, plasticity induction is derived from a gradient descent procedure of a cost function. A synapse-specific tag is set when strong plasticity is induced. The induced plasticity at a tagged synapse is consolidated if the low-pass filtered activity of the post-synaptic neuron exceeds a threshold. A decay of the synaptic strength only arises for non-tagged synapses. The synaptic forgetting appears to be beneficial as it reduces learning inferences caused by context switches.

I-116. Feature detectors drive dimorphic behaviors in Drosophila

David Deutsch[1] DDEUTSCH@PRINCETON.EDU
Jan Clemens[2] CLEMENSJAN@GMAIL.COM
Mala Murthy[1] MMURTHY@PRINCETON.EDU

[1]Princeton University
[2]European Neuroscience Institute (Goettingen)

During acoustic communication animals need to detect relevant features from a complex and dynamic signal and make appropriate decisions. During courtship, D. melanogaster males sing a song that contains two primary modes: pulse and sine. The song evolves over multiple timescales and is dynamically sculpted by feedback from the female. Males and females respond to song with sexually dimorphic behaviors: males court/sing while females slow and eventually allow or reject copulation. However, we do not know 1) which song features are relevant in males and females, 2) how this auditory information is used to drive the sex-specific behaviors and 3) how are these features represented in the fly brain. Using a new high throughput behavioral assay, we mapped the behavioral response (changes in locomotion) of males and females to courtship song, testing $>$50 stimuli. The behavioral tuning in both sexes matches conspecific song features. Whereas the responses of males and females to sine song are similar, the responses to pulse song are sexually dimorphic. We therefore looked for neurons in the fly brain that are tuned to conspecific song features in both genders, but drive sex-specific behaviors. Focusing on sexually dimorphic, doublesex-expressing cells, we found a set of neurons that has a strong response to pulse but not to sine song. These cells, called 'pC2 neurons', are tuned to multiple conspecific pulse-song features, and their auditory responses match behavioral responses. Inactivating pC2 neurons in females delays time to copulation and affects female response to pulse song during courtship. Optogenetic activation of pC2 cells induces slowing in solitary females, while driving courtship song in males. pC2 cells are therefore feature detectors that are tuned to conspecific song features and drive different behaviors in males and females. This suggests that the sexual dimorphism in the acoustomotor pathway lies downstream of the feature-detectors.

I-117. A simple model for low variability in neural spike trains

Ulisse Ferrari[1] ULISSE.FERRARI@GMAIL.COM
Stephane Deny[2] STEPHANE.DENY.PRO@GMAIL.COM
Thierry Mora[3] THIERRY.MORA@GMAIL.COM
Olivier Marre[1,4] OLIVIER.MARRE@GMAIL.COM

[1]Universite de Pierre et Marie Curie
[2]Stanford University
[3]Ecole Normale Superieure
[4]INSERM

Neural noise sets a limit to information transmission in sensory systems. In several areas, the spiking response (to a repeated stimulus) has shown a higher degree of regularity than predicted by a Poisson process (Kara et al, 2000). However, a simple model to explain this low variability is still lacking. Here we introduce a new model, with a correction on Poisson statistics, which can accurately predict the regularity of neural spike trains in response to a repeated stimulus. The model has only two parameters, but can reproduce the relation between mean firing rate and variance in retinal recordings in various conditions. We show analytically why this approximation can work. In a model of the spike emitting process where a refractory period is assumed, we derive that our simple correction can well approximate the spike train statistics over a broad range of firing rates. Our model can be easily plugged to LN or LNLN models to replace the Poisson spike train hypothesis that is commonly assumed. We show that it estimates the amount of information transmitted much more accurately than Poisson models in retinal recordings. Thanks to its simplicity this model has the potential of explaining low variability in other areas (Gur et al 1997; DeWeese and Zador, 2002; Maimon and Assad, 2009; Baudot et al, 2013).

I-118. Local diff?usion geometry for automated cellular structure extraction in calcium imaging data

Gal Mishne[1] GAL.MISHNE@YALE.EDU
Ronald Coifman[1] COIFMAN-RONALD@YALE.EDU
Maria Lavzin[2] MARIA.LAVZIN@GMAIL.COM
Jackie Schiller[2] JACKIE@TECHNION.AC.IL

[1]Yale University
[2]Technion - Israel Institute of Technology

Calcium imaging enables measuring in-vivo activity of large populations of neurons at cellular level resolution. The analysis of this complex data relies first and foremost on extracting neurons as high-resolution regions of interest, while demixing overlapping spatial components and denoising the time-series of each neuron. We propose a novel data-driven solution to these challenges, based on representing the spatio-temporal volume as a graph in the image plane. Aggregating the spectral embeddings of this graph from multi-trial data, we develop a new greedy selective spectral clustering method, capable of handling overlapping clusters and disregarding clutter. Based on the spectral embedding, we also present a new non-linear mapping for visualization of the structural map of the neurons and dendrites, and perform global video denoising. We demonstrate our approach on in-vivo calcium imaging of neurons and apical dendrites, automatically extracting complex-shaped neuronal and dendritic structures in the image domain, while separating overlapping ROIs, and demixing and denoising their time-series.

I-119. Flexible, optimal motor control in a thalamo-cortical circuit model

Mahdieh Sadabadi MSS69@CAM.AC.UK
Ta-Chu Kao TCK29@CAM.AC.UK
Guillaume Hennequin G.HENNEQUIN@ENG.CAM.AC.UK
University of Cambridge

How does the brain control movement? Experiments suggest that the (pre-)motor cortex behaves like an "engine" whose dynamics drive movement, and whose activity must first be initialised into movement-specific states (the "optimal subspace hypothesis", Shenoy et al, 2013). Both the computational and mechanistic underpinnings of this preparatory process remain poorly understood. Here, we propose a realistic circuit model for movement preparation and execution. We formalise movement preparation as an optimal control problem, under an internal forward model that predicts (future) patterns of muscle activity from momentary cortical preparatory states. We compute the preparatory input to cortex that drives fastest convergence to preparatory states predicted to yield the correct motor outputs. We also request that the control input keeps preparatory activity in an appropriate "output-null" subspace to prevent unwanted, premature motor outputs (Kaufman et al, 2014).

Critically, we show that optimal control inputs can be realised via feedback in realistic neural circuit architectures. Specifically, we model cortex as an inhibition-stabilised network, whose dynamics resembles those of monkey M1 during reaching (Hennequin et al, 2014). Optimal movement preparation is accomplished by a thalamo-cortical loop, gated by the basal ganglia. The loop is open by default, closed to drive movement preparation, and eventually re-opened to initiate movement. Importantly, we find that control loops can be flexibly combined to generate movements assembled from a few movement primitives.

The model produces naturalistic patterns of preparatory activity, including complex transients early during preparation, and displaying substantial variability in output-null dimensions. Consistent with data, across-trial variability is suppressed during preparation. Moreover, preparation may be as short as 200ms without loss of motor accuracy, also consistent with recent experimental observations. Our work brings together several threads of experimental research on both cortical and subcortical areas, and offers a new computational, normative perspective on the dynamics of motor circuits.

I-120. Co-occurrence statistics of natural sound features predict perceptual grouping

Wiktor Mlynarski MLYNAR@MIT.EDU
Josh H McDermott JHM@MIT.EDU
Massachusetts Institute of Technology

Events and objects in the world must be inferred from sensory signals to support behavior. Because sensory signals are transduced with measurements that are temporally and spatially local, the estimation of a signal component arising from a particular object or event can be viewed as the result of grouping these local measurements into representations of their common causes. In the auditory system, perceptual grouping is believed to exploit acoustic regularities of natural sounds, such as the tendency of frequencies to be harmonically related or to share a common onset. However, acoustic grouping cues have traditionally been identified using intuitions and informal observation, and investigated using simple, artificial stimuli. As a result, the relevance of known grouping cues to real-world auditory scene analysis remains unclear, and additional or alternative cues remain a possibility. We attempted to link auditory grouping principles to the structure of natural sounds by measuring feature co-occurrences in natural signals and exploring their relation to perception. We first derived a set of primitive auditory patterns by learning a dictionary of spectrotemporal features from a corpus of natural sounds using sparse convolutional coding. We then identified sets of features that are frequently co-active in natural sounds, as well as sets of features that do not naturally co-occur. In a psychophysical experiment, we found that stimuli created by superimposing naturally co-occurring features were more likely to be recognized as a single source than pairs of features that do not commonly co-occur. Clustering these stimuli revealed traditional grouping cues such as harmonicity and common onset. Other clusters suggested novel grouping principles such as the tendency of noise-like, tonal, and click-like sounds to segregate from each other. Our results suggest that auditory grouping cues are adapted to natural stimulus statistics, and that considering these statistics can reveal previously unappreciated aspects of grouping.

I-121. Inferring mesoscopic population models from population spike trains

Alexandre Rene[1] ALEXANDRE.RENE@CAESAR.DE
Andre Longtin[2] ALONGTIN@UOTTAWA.CA
Jakob Macke[1] JAKOB.MACKE@CAESAR.DE

[1]Research Center Caesar
[2]University of Ottawa

How does the interplay of single-neuron dynamics and neural connectivity give rise to the rich dynamical properties of neural populations? To tackle this question, it is desirable to have models which exhibit a wide range of population dynamics but remain interpretable in terms of connectivity and single-neuron dynamics. However, many commonly-used statistical models of neural population dynamics are based on generic models of dynamics (e.g. in Macke et al. 2011). Conversely, it has been challenging to link mechanistic spiking network models to empirical population data. To close this gap, we propose to model such data using mechanistic, but low-dimensional and hence statistically tractable models. We approximate neural populations as being composed of multiple homogeneous 'pools' of neurons, and model the dynamics of the aggregate population activity within each pool. We derive the likelihood of parameters (both single-neuron parameters and inter-pool connectivity) given this activity, which can then be used to either optimize parameters by gradient ascent on the log-likelihood, or to perform Bayesian inference using Markov Chain Monte Carlo (MCMC) sampling. We illustrate this approach on a model based on generalized integrate-and-fire neurons (Schwalger et al., 2017). Using micro- and mesoscopic simulations of multiple neuron pools, we demonstrate that both single-neuron properties (membrane and adaptation constants) and connectivity-parameters (excitatory vs inhibitory connections and connection strengths) can be recovered on simulated data. Moving beyond point estimates, we compute the Bayesian posterior for combinations of parameters using MCMC sampling. Finally, we investigate how the approximations inherent to a mesoscopic population model impact the accuracy of the inferred single-neuron parameters. Ultimately, our method ensures compatibility between experimental multi-population data and mesoscopic dynamical models, by providing methods for statistical inference of low-dimensional mesoscopic models.

I-122. Effective learning is accompanied by high dimensional & efficient representations of neural activity

Evelyn Tang[1] EVELYNT@SEAS.UPENN.EDU
Marcelo Mattar[2] MARCELOMATTAR@GMAIL.COM
Chad Giusti[3] CGIUSTI@UDEL.EDU
Sharon Thompson-Schill[1] SSCHILL@PSYCH.UPENN.EDU
Danielle Bassett[1] DSB@SEAS.UPENN.EDU

[1]University of Pennsylvania
[2]Princeton University
[3]University of Delaware

A fundamental cognitive process is the ability to map objects into meaningful categories depending on the task at hand. How such mental constructs emerge and what kind of space best embeds this mapping are poorly understood. Here we develop tools to quantify the space and organization of such a mapping, using a geometric

perspective adapted from machine learning. This approach represents neural responses from spatially distributed brain regions as points in a multidimensional space, enabling the identification of the low dimensional subspaces despite noisy measurements. While such methods have been effective in the analysis of neuron-level data, we extend these tools to study human neural responses as indirectly estimated in functional MRI, and we introduce new metrics to quantify a notion of efficient cognitive coding. Applying these tools to data acquired as adults learned the values of novel stimuli, we show that quick learners have a higher dimensional geometric representation than slow learners, and hence more easily distinguishable whole-brain responses to objects of different value. Furthermore, we find that quick learners display a more compact dimension of their neural responses without the task-relevant labels, which reflects the underlying embedding of the neural activity used. The ratio of the first task-relevant dimension to the second embedding dimension is hence much larger for fast learners, suggesting the construction of a more informative set of task-relevant features using a smaller amount of effective resources. This pattern of results is consistent with a greater efficiency of cognitive coding. Our results demonstrate a spatial organization of neural responses characteristic of the successful optimization of reward, and offer geometric measures applicable to the study of efficient coding in higher-order cognitive processes more broadly.

II-1. Bayesian inference of neural activity and connectivity from all-optical interrogation of a neural circuit

Laurence Aitchison[1] LAURENCE.AITCHISON@GMAIL.COM
Lloyd Russell[2] LLERUSSELL@GMAIL.COM
Adam Packer[2] ADAMPACKER@GMAIL.COM
Jinyao Yan[3] YANJ11@JANELIA.HHMI.ORG
Phillipe Castonguay[3] PH.CASTONGUAY@GMAIL.COM
Michael Hausser[2] M.HAUSSER@UCL.AC.UK
Srini Turaga[3] TURAGAS@JANELIA.HHMI.ORG

[1]University of Cambridge
[2]University College London
[3]HHMI Janelia Research Campus

Population activity measurement by calcium imaging can be combined with cellular resolution optogenetic activity perturbations to enable the mapping of neural connectivity in vivo. This requires accurate inference of perturbed and unperturbed neural activity from calcium imaging measurements, which are noisy and indirect, and can also be contaminated by photostimulation artifacts. We have developed a new fully Bayesian approach to jointly inferring spiking activity and neural connectivity from in vivo all-optical perturbation experiments. In contrast to standard approaches that perform spike inference and analysis in two separate maximum-likelihood phases, our joint model is able to propagate uncertainty in spike inference to the inference of connectivity and vice versa. We use the framework of variational autoencoders to model spiking activity using discrete latent variables, low-dimensional latent common input, and sparse spike-and-slab generalized linear coupling between neurons. Additionally, we model two properties of the optogenetic perturbation: off-target photostimulation and photostimulation transients. Using this model, we were able to fit models on 30 minutes of data in just 10 minutes. We performed an all-optical circuit mapping experiment in primary visual cortex of the awake mouse, and use our approach to predict neural connectivity between excitatory neurons in layer 2/3. Predicted connectivity is sparse and consistent with known correlations with stimulus tuning, spontaneous correlation and distance.

II-2. Myopic control: A new control objective for neural population dynamics

David Hocker DLHOCKER@GMAIL.COM
Il Memming Park MEMMING.PARK@STONYBROOK.EDU
Stony Brook University

Manipulating the dynamics of neural systems through targeted stimulation is key to causal investigation in systems neuroscience and clinical neuroscience; however, the modern control objectives in engineering are mismatched for the unique needs of manipulating neural dynamics. An appropriate control method should respect the variability in neural systems, incorporating moment to moment "input" to the neural dynamics and behaving based on the current neural state, irrespective of the past trajectory. We propose such a controller under a nonlinear state-space feedback framework that steers one dynamical system to function as though it were another dynamical system entirely. This "myopic" control objective manipulates dynamics instantaneously, omitting the need to pre-compute a rigid and computationally costly feedback control solution. To demonstrate the breadth of this control's utility, two simulated examples with distinctly different applications in neuroscience are studied. First an unhealthy motor-like system containing an unwanted beta-oscillation spiral attractor is controlled to function as a healthy motor system, a relevant clinical example for neurological disorders. Second, we show myopic control's utility to probe the causal dynamics of cognitive processes by transforming a winner-take-all decision-making system [Wong and Wang, J. Neurosci. 2006] to operate as a robust neural integrator of evidence [Koulakov et al. Nat. Neurosci. 2002].

II-3. Homeostatic plasticity in neural networks induces diverse dynamic states

Johannes Zierenberg JOHANNES.ZIERENBERG@DS.MPG.DE
Jens Wilting JWILTING@NLD.DS.MPG.DE
Viola Priesemann VIOLA@NLD.DS.MPG.DE
Max Planck Institute for Dynamics and Self-organization

Surprisingly, dynamics of spiking neural networks exhibit clear differences between in vivo and in vitro experiments, despite similar neural and synaptic properties. In vivo networks show continuous, fluctuating dynamics, whereas in vitro networks typically develop strong bursts separated by periods of silence. We propose that the different dynamics result from an interplay between (1) network input, which is much stronger in vivo than in vitro, and (2) homeostatic plasticity, a slow negative feedback mechanism adapting the neural spike rate. Using mean-field calculations and simulations, we find consistently across different network topologies that homeostatic plasticity clearly adapts the network dynamics to the input strength. Thereby our results account for the irregular or reverberating network dynamics typically observed in vivo, could explain close-to-critical dynamics for layer 2/3 cortex if one assumes a dominance of recurrent interactions in this layer, as well as the unavoidable bursts characteristic for in vitro networks. Moreover, assuming increasingly weaker input for higher brain areas, the hierarchy of network timescales observed across cortical areas naturally arises owing to the action of homeostasis. Thereby we propose a simple mechanism that can tune the integration times for information processing. Last but not least, our results predict that homeostasis can be harnessed to abolish the ubiquitous bursts observed in vitro: The framework predicts that already weak network stimulation can render in vitro-dynamics in vivo-like instead. This would allow to establish a physiological ground state in vitro—a key prerequisite to study the effects and the treatment of neurological and psychiatric disorders on the network level in an cost-efficient in vitro assay.

II-4. Mechanisms underlying place field development and remapping. A computational study.

Victor Pedrosa VICTORPEDROSABC@GMAIL.COM
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK
Imperial College London

The hippocampus plays a key role in spatial representation. A subset of hippocampal neurons, called place cells, are modulated by the animal's location within each environment. Even before the first exploration, future place cells and future spatially untuned silent cells can be distinguished by their intrinsic properties. Manipulating these properties can switch a silent cell into a place cell. Additionally, silencing all place cells active in a familiar environment has been shown to unveil a complete new subset of previously quiet cells, enabling the emergence of a new place map. Although vastly explored experimentally, the mechanisms underlying place field development and remapping are not fully understood. Using a computational model, we show how a simple inhibition-mediated network of neurons can lead to the emergence of two subsets of cells: active place cells and spatially untuned silent cells. We study how the suppression of place cells leads to a rapid and transient emergence of a new place map. Using a Hebbian-like plasticity model, we illustrate how this new place map can be stabilized with the repeated suppression of the initially active cells. Moreover, we propose an input-dependent inhibitory plasticity rule. When combined with heterogeneous, sparse connections, this rule leads to a rapid—but not instantaneous—place map shift while maintaining global neuronal activity. Our study provides a possible mechanism for place map plasticity in the hippocampal CA1 network.

II-5. Extracting nonlinear manifolds from spike train data with Gaussian process latent variables

Anqi Wu ANQIWU.ANGELA@GMAIL.COM
Nicholas Roy NROY@PRINCETON.EDU
Stephen Keeley STEPHENLKEELEY@GMAIL.COM
Jonathan Pillow PILLOW@PRINCETON.EDU
Princeton University

A large body of recent work in computational neuroscience focuses on methods for extracting low-dimensional latent structure from multi-neuron spike train data. Most of these methods have used linear dimensionality reduction of firing rates, linear models of latent dynamics, or both. Here we propose a nonlinear latent variable model that can identify low-dimensional, highly-nonlinear structure underlying high-dimensional spike train data, which we call the Poisson Gaussian-Process Latent Variable Model (P-GPLVM). This model specifies the joint distribution over multi-neuron spike trains in terms of conditionally Poisson spiking with firing rates driven by the composition of two underlying Gaussian processes (Gps)—one governing the trajectory of a low-dimensional temporal latent variable, and another governing a set of tuning curves that map the latent variable to high-dimensional firing rates. The use of nonlinear tuning curves allows the model to discover low-dimensional latent structure even when spike responses are themselves high-dimensional (e.g., as in hippocampal place cell or entorhinal grid cell codes). To infer the model parameters and GP hyperparameters from data, we introduce the decoupled Laplace approximation, a fast approximate inference method that allows us to efficiently optimize the latent path while marginalizing over latent tuning curves. We show that this method outperforms previous Laplace based inference methods in both the speed and accuracy. We apply the model to spike trains recorded from hippocampal place cells and show that it outperforms a variety of previous methods for latent structure discovery, including variational auto-encoder (VAE) based methods that parametrize the nonlinear mapping from latent space to spike rates with a deep neural network.

II-6. Action-outcome signals on multiple timescales in medial prefrontal cortex drive flexible decisions

Bilal Bari BBARI1@JHMI.EDU
Cooper Grossman CGROSSM4@JHMI.EDU
Jeremiah Cohen JEREMIAH.COHEN@JHMI.EDU

Johns Hopkins University

In an unstable world, the brain relies on recent experience to generate adaptive decisions. The medial prefrontal cortex (mPFC) is among the areas involved in generating flexible behavior, but how it does so is unclear. To determine how mPFC encodes outcomes and actions dynamically, we developed a foraging task in thirsty, head-restrained mice, adapted from one in monkeys (Sugrue et al., 2004; Lau & Glimcher, 2005). On each trial, an olfactory "go" cue signaled to the mouse to make a choice. Mice chose freely between two "lick ports," one on the left of the tongue, one on the right. Each lick port delivered water reward probabilistically and these probabilities changed over time. To maximize reward, mice used outcomes from previous trials to choose actions flexibly. An action-value-based reinforcement-learning model predicted behavior well, suggesting a simple generative algorithm for choice behavior.

To determine how mPFC contributed to behavior, we reversibly inactivated it, using the GABAA receptor agonist muscimol. This prevented mice from updating actions adaptively. Experiments using other behavioral tasks revealed that this was not due to a global deficit in value updating (i.e., memory for recent rewards) nor was it due to a deficit in mapping decisions onto motor outputs (i.e., licking appropriately in either direction). We next recorded action potentials from 3,419 neurons in six mice performing the foraging task. Most neurons encoded a selective action-outcome feedback signal (65%). Nearly a third of neurons (29%) contributed to a persistent representation of action-specific reward history, bridging the time between trials. Two smaller populations of neurons encoded the upcoming action of the mouse (7%) and the overall reward history (6%). Together, these findings demonstrate that activity in mPFC over multiple timescales dynamically encodes behavioral variables necessary for flexible decision making.

II-7. Accessing neural states in real time: recursive variational Bayesian dual estimation

Yuan Zhao YUAN.ZHAO@STONYBROOK.EDU
Il Memming Park MEMMING.PARK@STONYBROOK.EDU

Stony Brook University

How neural systems compute is often studied in the framework of continuous dynamical systems theory both qualitatively and quantitatively. State space model provides an interpretable view of complex time series by combining an intuitive dynamical system model with a probabilistic observation model. Learning both the latent state trajectory and the latent (nonlinear) state space model is known as the dual estimation problem. Offline batch analyses such as expectation maximization (EM) based methods and more recently variational autoencoder methods have been proposed and widely used in practice. However we are more interested in real-time signal processing and state-space control setting where we need online algorithms that can recursively solve the dual estimation problem on streaming neural observations. We developed a flexible online black-box inference framework for latent nonlinear state dynamics and filtered latent states that is applicable to a wide range of nonlinear state space models. Our method utilizes the stochastic gradient variational Bayes method to jointly optimize the parameters of the nonlinear dynamics, observation model, and the recognition model. Unlike previous approaches, our framework can incorporate non-trivial observation noise models and infer in real-time meaning the computational demand of the algorithm is constant per time step. We test our method on point process observations driven by continuous attractor dynamics, demonstrating its ability to recover the phase portrait, filtered trajectory, and produce long-term predictions for neuroscience applications. Our primary target application is neural spike train data analysis, where inferring state space models underlying neural activity can provide insights into neural dynamics, neural computation, and development of neural prosthetics and treatment through feedback control.

II-8. A local grid score for individual spikes of grid cells

Simon Weber[1] WEBER@TU-BERLIN.DE
Henning Sprekeler[1,2] H.SPREKELER@TU-BERLIN.DE

[1]Berlin Institute of Technology
[2]Bernstein Center for Computational Neuroscience

The location-specific firing of cells in the entorhinal cortex is subject to extensive experimental and theoretical research. When classifying the tuning properties of entorhinal cells, researchers distinguish between cells that fire on a hexagonal grid of locations (grid cells), cells that fire periodically but without hexagonal symmetry and cells without periodic firing patterns. This classification requires a measure for the symmetry of spatially modulated firing patterns—a grid score. The most established grid score is computed in multiple stages. First, spike locations are transformed into a rate map. Subsequently, an autocorrelogram of the rate map is cropped, rotated and correlated with its unrotated copy. The final grid score is obtained from the resulting correlation-vs-angle function at selected angles. This procedure results in a global grid score for the firing pattern. Here we suggest a new approach that computes a local grid score—and the local grid orientation—for each individual spike, directly from spike locations. Averaging over spikes, we obtain a global grid score and show that it is as reliable as existing grid scores in quantifying the global hexagonal symmetry of a firing pattern. The new score enables the plotting of spike locations, color-coded with the local grid score or the local orientation of the grid and could thus simplify the visualization of experimental data. More specifically, it could be used to quantify and highlight recent experimental findings on local properties of grid patterns, like boundary effects in asymmetric enclosures and drifts in grid orientation along the arena. The grid score is applicable to any n-fold symmetry. We provide a public Python package (using SciPy and NumPy) and a website that efficiently determines the grid score directly from spike locations.

II-9. Multiplicative interaction between perceptual biases and sensory input in motion estimation

Jacob Yates[1] JYATES7@UR.ROCHESTER.EDU
Shaun Cloherty[2] SHAUN.CLOHERTY@MONASH.EDU
Gregory DeAngelis[1] GDEANGELIS@UR.ROCHESTER.EDU
Jude Mitchell[1] JMITCHELL@MAIL.BCS.ROCHESTER.EDU

[1]University of Rochester
[2]Monash University

The perception of low-level visual features has repeatedly been shown to be biased. This bias is often explained by considering perception to result from Bayesian inference wherein intrinsic knowledge of environmental statistics is combined with noisy sensory input [2]. A key prediction of such statistical inference is that intrinsic biases (reflecting environmental statistics) and sensory input should combine multiplicatively. Using the visual motion perception as a model, we tested this prediction in both humans and marmosets performing a continuous direction estimation task. Using probabilistic models of the subjects' trial-by-trial choices, we learned each subject's internal bias distribution while characterizing their sensory noise. These bias distributions were idiosyncratic for each subject, but tended to have peaks on the cardinal axes. We tested whether single trial responses were better described by an additive mixture of sensory noise and biased guesses, or as a multiplicative interaction between the internal bias distribution and the incoming sensory input. Our data from both humans and marmosets are best described by a multiplicative model. This model approach and our behavioral paradigm provides a platform

for studying the integration of sensory signals with intrinsic and learned priors at the level of neural populations in the behaving marmoset.

II-10. Efficient tracking of dynamic psychophysical behavior during learning

Nicholas Roy[1] NROY@PRINCETON.EDU
Ji Hyun Bak[2] JHBAK@KIAS.RE.KR
Athena Akrami[1] AAKRAMI@PRINCETON.EDU
Carlos Brody[3,1] BRODY@PRINCETON.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Korea Institute for Advanced Study
[3]Howard Hughes Medical Institute

Animal learning is a critical aspect of modern neuroscientific research. Having an untrained animal learn to perform a task adequately well is often prerequisite to systems level research, whereas the acquisition of a new task is itself the focus of many cognitive studies. Unfortunately, standard methods for assessing the behavior of an animal over the course of learning have progressed slowly from static measures that approximate behavior over a fixed window of trials, e.g. psychometric curves. Smith & Brown (2004) introduced a method to quantify learning dynamics, approximating the trajectory of an unobservable process corresponding to an animal's probability of a correct response. We propose a new method for 2-alternative forced choice tasks that instead recovers trajectories of psychophysical weights corresponding to specific behavioral phenomena, such as choice bias, reward history dependence, and stimulus weighting. Our model requires only that these psychophysical weights evolve smoothly in time, a constraint imposed by a prior over the trial-to-trial changes in these weights [Bak et al., 2016]. While previous methods used local approximations to find these trajectories, we leverage sparsity in our model to efficiently find the global MAP estimate of our weights, increasing accuracy as well as allowing for application to large behavioral datasets (N = 100,000+ trials) and complex models (K = 10+ weights). Additionally, we introduce the decoupled Laplace method which permits efficient estimation of the hyperparameters parameterizing the prior over weight changes, optionally allowing smoothness to vary between distinct weights and across time. Here we demonstrate the accuracy of our method in recovering psychophysical weights & hyperparameters, as well as propose practical applications in current training regimens.

II-11. How fast is neural winner-take-all when deciding between many options?

Rishidev Chaudhuri[1] RCHAUDHURI@AUSTIN.UTEXAS.EDU
Birgit Kriener[2] BIRGIT.KRIENER@MEDISIN.UIO.NO
Ila Fiete[1] FIETE@MAIL.CLM.UTEXAS.EDU

[1]The University of Texas at Austin
[2]University of Oslo

Finding the largest (or best) of N options (max, argmax) is a ubiquitous computation, involved in tasks like inference, optimization, decision-making, action selection, consensus, and foraging. A serial strategy, as would be implemented on a computer, must examine each option, taking a time that scales with N if the options are noise-less. If the options fluctuate around their means, then each must be integrated; we find the time-complexity of this procedure is then Nlog(N).

The computational power of the brain is attributed partly to its parallelism, yet it is theoretically unclear whether, even on an elemental task like max, it is possible for leaky and noisy neurons with constrained firing rates to perform distributed computations in a time that achieves a full parallel speedup over serial strategies. The benchmark

for parallel computation on max is a factor-N speedup: a constant decision time for noiseless inputs and a time of log(N) for noisy inputs.

We show that neural winner-takes-all (WTA) circuits, in which neurons compete to choose the best option and silence the others, fail asymptotically (large N) to achieve the parallelism benchmark and worse, in the presence of noise altogether fail to produce a winner. If, however, neurons are equipped with a second nonlinearity so that weakly active neurons cannot contribute inhibition to the circuit, the network can complete the computation N times faster than serial strategies at comparable accuracy, partially self-adjusting integration time for task difficulty and number of options, and saturating the parallelism benchmark without parameter fine-tuning. Finally, in the regime of few choices (small N), the circuit predicts Hick's law of decision-making. Thus, Hick's law behavior is a symptom of efficient parallel computation.

Our work shows that distributed computation saturating the parallelism benchmark is possible in networks of noisy and finite-memory neurons.

II-12. Navigation of an abstract discrete environment by rhesus macaques

James Butler[1] JAMES.BUTLER@UCL.AC.UK
Shirley Mark[1] MARKSHIR@GMAIL.COM
Tim Behrens[1,2] BEHRENS@FMRIB.OX.AC.UK
Steve Kennerley[1] S.KENNERLEY@UCL.AC.UK

[1]University College London
[2]University of Oxford

Since the discovery of place cells and the 'cognitive map' by O'Keefe & Dostrovsky in 1971 much research has revealed the mechanisms underlying navigation of spatial environments. However, how mammals navigate more abstract conceptual spaces, such as those that underlie complex cognitive processes, remains unclear. We therefore taught two rhesus macaques (Macaca mulatta) several abstract discrete environments composed of networks of associated stimuli (e.g. a 5x5 grid with 4 edges per node). After just seven days of training subjects could successfully navigate to rewarded target locations several stimuli away with high precision. When learning a new environment, due to the high number of unique pathways ($>$500), some routes would be encountered for the first time near the beginning of training whereas others would be encountered later in training. Subjects performed better in their first encounter of a unique pathway if it occurred later in training rather than earlier, therefore demonstrating a global representation of the map.

When presented with two options equidistant to the target, it is advantageous to choose the one with more possible subsequent routes to the goal. This is however computationally-expensive to calculate in an environment with such a large state space. A successor representation-like strategy which assigns a value based on the potential reward of all possible future states would perform well at this task, whereas a tree search-like strategy that individually searches all possible routes would take exponentially longer to solve this task. We found both subjects preferentially chose options with multiple possible routes to the target over those with fewer possible routes, and this was consistent across multiple environments tested. We have therefore shown that non-human primates can navigate complex abstract environments, can learn knowledge about the structure of these environments, and can navigate in an optimal manner which indicates a computationally-efficient process.

II-13. Chloride dynamics alter the input-output properties of neurons

Christopher Currin[1] CHRIS.CURRIN@GMAIL.COM
Andrew Trevelyan[2] ANDYTREV@GMAIL.COM
Joseph Raimondo[1] JOSEPH.RAIMONDO@UCT.AC.ZA

[1]University of Cape Town
[2]Newcastle University

Different populations of inhibitory interneurons change the input-output (IO) function of their targets in different ways [1]–[5]. A factor which has received less consideration is changes in the reversal potential of GABA ($\Delta EGABA$), arising from changes in neuronal chloride ion concentration ([Cl-]i). Chloride-loading of neurons is known to be a factor in various neurological conditions [6]. Here we investigate its impact on the input-output function of neurons. We explored how local dendritic parameters affect the dynamics of [Cl-]i in response to synaptic input, and how this might affect computation depending on the location of synaptic input. By introducing a novel measure of the functional influence of Cl- on the neuron's output, the 'chloride index', we quantify the difference between a model with and without Cl- dynamics. Our results highlight the importance of accounting for Cl- dynamics when designing computational models that include dendritic inhibition. Inhibition can quickly change the computational properties of a neuron over time as Cl- accumulates inside the dendrites. Neurons with finely balanced excitatory and inhibitory synaptic input are particularly susceptible to transient changes in Cl- with shifts in firing rate being more pronounced over time, as well as over increasing frequency of input. Distal GABAergic dendritic input is more susceptible to [Cl-]i accumulation than proximal input, which impacts the amount of control inhibition has on the output of the neuron and maps to physiological findings of GABA synapse distribution in the cortex. These results are highly relevant for assessing computational models which incorporate static inhibition and generally apply to models exhibiting inhibitory plasticity. We highlight the important computational changes neurons can undergo over short periods of time due to inhibitory input, which applies to models of dendritic processing, the behaviour of individual neurons, and also to networks of neurons.

II-14. Nonlinear synaptic interaction as a computational resource in the Neural Engineering Framework

Andreas Stoeckel ASTOECKE@UWATERLOO.CA
Aaron Voelker ARVOELKE@UWATERLOO.CA
Chris Eliasmith CELIASMITH@UWATERLOO.CA

University of Waterloo

Nonlinear interaction in the dendritic tree is known to be an important computational resource in biological neurons. Yet, high-level neural compilers—such as the Neural Engineering Framework (NEF), or the predictive coding method published by Deneve et al. in 2013—tend not to include conductance-based nonlinear synaptic interactions in their models, and so do not exploit these interactions systematically. In this study, we extend the NEF to include synaptic computation of nonlinear multivariate functions, such as controlled shunting, multiplication, and the Euclidean norm. We present a theoretical framework that provides sufficient conditions under which nonlinear synaptic interaction yields a similar precision compared to traditional NEF methods, while reducing the number of layers, neurons, and latency in the network. The proposed method lends itself to increasing the computational power of neuromorphic hardware systems and improves the NEF's biological plausibility by mitigating one of its long-standing limitations, namely its reliance on linear, current-based synapses. We perform a series of numerical experiments with a conductance-based two-compartment LIF neuron model. Preliminary results show that nonlinear interactions in conductance-based synapses are sufficient to compute a wide variety of nonlinear functions with performance competitive to using an additional layer of neurons as a nonlinearity.

II-15. Current- vs. inhibition-based theta oscillators: applications to speech parsing.

Benjamin Pittman-Polletta[1] BENPOLLETTA@GMAIL.COM
David Stanley[1] DAVEARTHUR.STANLEY@GMAIL.COM
Miles Whittington[2] MILES.WHITTINGTON@HYMS.AC.UK
Charles Schroeder[3] CHARLES.SCHROEDER@NKI.RFMH.ORG
Nancy Kopell[1] NK@BU.EDU

[1]Boston University
[2]Hull York Medical School
[3]Nathan Kline Institute

Experimental evidence suggests that hierarchically nested intrinsic oscillators at frequencies such as delta (~1-3 Hz), theta (~3-9 Hz), and beta/gamma (>=30 Hz) support the parsing of quasi-rhythmic auditory stimuli evolving on multiple timescales, such as speech. By tracking periodic features of such stimuli, intrinsic oscillators are hypothesized to subdivide the sensory input stream into comprehensible chunks. In particular, neurophysiological and psychophysical evidence suggests that the ability of an intrinsic auditory cortex theta oscillator to phase lock to the syllabic structure of speech stimuli—parsing speech into packets ("theta syllables") of 110-330 ms to be encoded upstream—is crucial for speech intelligibility (Ghitza 2014). How such a flexible-frequency theta oscillator is implemented remains unknown. Motivated by in vitro data (Carracedo et al. 2013), we have developed a single-compartment Hodgkin-Huxley model of a layer 5 cortical pyramidal cell in which intrinsic currents lead to resonances on both theta and delta timescales, and delta-nested theta-rhythmic spiking. Changing the level of excitation of this cell toggles its dynamics between rhythmicity at either or both of the delta and theta frequencies, with an intermediate regime of delta-theta coupling. In contrast to inhibition-based theta oscillators (IBTs), this conductance-based theta oscillator (CBT) exhibits phase-locked spiking to periodic input pulses over a broad yet restricted range of frequencies (~1-9 Hz), as well as to frequency modulated sinusoids whose instantaneous frequency ranges over this interval (which corresponds to the duration of theta syllables permitting accurate speech encoding). Phase-response curves (PRCs) indicate that while IBTs exhibit only phase advances in response to pulsatile input, the CBT's PRC exhibits regions of both advance and delay, and is able to phase-lock to rhythmic inputs at its intrinsic frequency within a single cycle. Our work thus suggests that the dual resonance exhibited by our CBT—and the nested rhythmicity it produces—may play a key role in the parsing of speech signals into syllabic units.

II-16. Spiking allows neurons to estimate their causal effect

Ben Lansdell BEN.LANSDELL@GMAIL.COM
Konrad Kording KOERDING@GMAIL.COM

University of Pennsylvania

Learning is typically conceptualized as changing a neuron's properties to improve reward. This ultimately is a problem of causality: to learn, a neuron needs to estimate its causal influence on reward. A typical solution in artificial neural networks is stochastic gradient descent through back-propagation. This is a challenge in biological neural networks due to physiological constraints on information transmission and the discontinuous nature of spiking activity. However viewing a neuron's influence on a reward signal as a measure of causal effect suggests alternative methods from causal inference.

We show that a method commonly employed in econometrics, regression discontinuity design (RDD), can be used by a neuron to estimate its causal effect. It works by comparing reward at instances where a neuron is driven to be marginally below threshold to instances where it is driven to be marginally above threshold. At these instances whether the neuron spikes or not is randomized, so observed differences in reward are attributable only to its spiking, and not to any other variables that the neuron's activity may otherwise be correlated with. Here we demonstrate the method in a simple network of leaky integrate and fire neurons. We derive a learning rule and show it can be used by a neuron to maximize reward.

Other related methods, such as node-perturbation, rely on the existence of independent noise sources for each neuron. These methods are confounded by noise correlations. This limits their applicability, as noise correlations occur in many settings. We demonstrate that the RDD method is robust to noise correlations in inputs. Further, the RDD approach exploits the spiking nature of neural activity rather than smoothing it out, like many other approaches to learning in spiking neural networks. By focusing on the underlying causal inference problem we obtain new ways to understand learning.

II-17. Morphological error detection for connectomics

David Rolnick[1] DROLNICK@MIT.EDU
Yaron Meirovitch[1] YARON.MR@GMAIL.COM
Jeffrey W. Lichtman[2] JEFF@MCB.HARVARD.EDU
Nir Shavit[1] SHANIR@CSAIL.MIT.EDU
Toufiq Parag[2] PARAGT@G.HARVARD.EDU
Hanspeter Pfister[2] PFISTER@SEAS.HARVARD.EDU
Viren Jain[3] VIREN@GOOGLE.COM
Edward S. Boyden[1] ESB@MEDIA.MIT.EDU

[1]Massachusetts Institute of Technology
[2]Harvard University
[3]Google

Deep learning algorithms for connectomics rely upon localized classification, rather than overall morphology. This leads to a high incidence of erroneously merged objects. Humans, by contrast, can easily detect such errors by acquiring intuition for the correct morphology of objects. Biological neurons have complicated and variable shapes, which are challenging to learn, and merge errors take a multitude of different forms. We present an algorithm, MergeNet, that shows 3D ConvNets can, in fact, detect merge errors from high-level neuronal morphology. MergeNet is able to detect merge errors with high accuracy within a three-dimensional segmentation and to pinpoint their locations for correction. The algorithm can be trained using any reasonably accurate segmentation, without the need for any additional annotation. It is even able to correct errors within its own training data. MergeNet can be applied irrespective of the segmentation algorithm or imaging method; it can be trained on one dataset and run on another with high performance. The algorithm also runs faster than state-of-the-art membrane prediction pipelines.

II-18. Within-trial dynamics of noise correlations imply binarized feedback of internal beliefs

Richard Lange[1] RLANGE@UR.ROCHESTER.EDU
Adrian Bondy[2] ADRIAN.BONDY@GMAIL.COM
Bruce Cumming[3] BGC@LSR.NEI.NIH.GOV
Ralf Haefner[1] RALF.HAEFNER@GMAIL.COM

[1]University of Rochester
[2]Princeton University
[3]NIH NEI

Correlated variability in populations of sensory neurons reflects a combination of noise and variable internal states that modulate the population collectively [1]. For example, it has been suggested that some of these internal states may be understood as probabilistic samples of a subject's "belief state" in an inference model of decision-making [2,3], or as continuously-varying attentional states [4,5]. If a single time-varying latent factor

both induces correlated variability in the population and is indicative of the subject's upcoming choice (as in latent belief state models), then one might expect that the time-course of choice probabilities (CPs) would match the time-course of induced correlations. However, a dissociation has recently been found in data from V1 of macaques during an orientation-discrimination task [5]. We show that because CP reflects the choice-triggered mean of the latent factor conditioned on a single choice, and correlations reflect the latent's variance across both choices, this result is evidence for feedback of a binarized (or sharply bimodal) "belief state" to sensory areas, while a graded representation is maintained in decision-making areas. To clarify this result, we study the more general relationship between the within-trial dynamics of latent factors and measured correlations. Using this framework, we present an intuitive model of a recent empirical finding on the timescale of attention-induced correlations in V1 [7], highlighting that the importance of the auto-correlation of a latent process on the magnitude of its effect on a population. Finally, we observe that in a neural sampling model, intermediate levels of the cortical hierarchy low-pass filter the effect of variability of a decision-related variable on sensory responses leading to the surprising prediction that the measured latent's autocorrelation should decrease as one records from areas further up the cortical hierarchy.

II-19. Long-term imaging of sensory representations reveals ongoing recombination of cell assemblies

Jens-Bastian Eppler[1,2] EPPLER@FIAS.UNI-FRANKFURT.DE
Dominik Aschauer[3] DASCHAUER@UNI-MAINZ.DE
Luke Ewig[2,1] EWIG@FIAS.UNI-FRANKFURT.DE
Matthias Kaschube[2,1] KASCHUBE@FIAS.UNI-FRANKFURT.DE
Simon Rumpel[3] SIRUMPEL@UNI-MAINZ.DE

[1]Goethe University Frankfurt
[2]Frankfurt Institute for Advanced Studies
[3]Gutenberg University Mainz

Cell assemblies are believed to form essential functional units within neural networks. How stable are such cell assemblies in the light of evidence from chronic imaging of dendritic spines indicating that cortical networks—even without any explicit learning paradigm—show substantial remodeling on the time scale of days (e.g. Loewenstein et al, 2015)? Using chronic two-photon imaging in awake mouse auditory cortex, we confirm previous observations that neuronal responses to short complex sounds typically cluster into a near discrete set of cell assemblies (Bathellier et al, 2012). Moreover, we find that cell assemblies show significant remodeling over several days: The mapping of sound stimuli onto fixed cell assemblies can undergo dynamic changes, but at the same time new cell assemblies emerge and old cell assemblies get eliminated. Stimuli that at one time-point have been mapped to the same cell assembly tend to re-map together. A simple circuit model shows that strong inhibitory and heterogeneous recurrent connections are sufficient to explain the observation of a near discrete set of cell assemblies. We then use this model to study how these cell assemblies are affected by random synaptic drift fitted to observations from mouse auditory cortex (Loewenstein et al, 2011). Gradual changes in the underlying circuitry can account for the rich dynamics in sensory representations that we observe in the experiment: some auditory responses remain fairly stable, while others show drastic changes within several days, overall resulting in a broadly distributed magnitude of stimulus response changes. We conclude that even subtle ongoing changes in synaptic connections can have a highly nonlinear effect on the stability of cortical representations. We speculate that the ongoing recombination of cell assemblies can provide a mechanism for the formation of associations that can occur also temporally separated from immediate learning situations.

II-20. A neural mechanism for determining allocentric locations of sensed features

Marcus Lewis MLEWIS@NUMENTA.COM
Jeff Hawkins JHAWKINS@NUMENTA.COM
Numenta

The neocortex can learn and recognize objects using input from independently moving sensors but the underlying neural mechanisms are not obvious. We posit that the neocortex accomplishes this by detecting each sensed feature's location relative to the object—the allocentric location—and by learning objects as sets of features-at-allocentric-locations (Hawkins et al, 2017). In this new work, we describe a model inspired by grid cell modules that can efficiently compute the egocentric-allocentric transform. The network consists of multiple cortical columns, each receiving independent sensory input. The key is to represent the allocentric location of each sensed feature as the vector sum of the "location of the feature relative to the body" and a global "location of body relative to the object". During inference, when a column senses a feature, it recalls all allocentric locations where it has previously sensed this feature on objects. The columns then collaborate to iteratively narrow down the unknown (but global) body location and the individual allocentric locations of each sensed feature. The locations can be initially ambiguous; in this case multiple movements and sensations are required to converge to a unique interpretation. Locations in the model are represented using modules similar to grid cell modules. We show a circuit that takes advantage of grid cell properties to perform the required metric computations. We discuss the learning rules, and propose a mapping to the known anatomy of cortical columns. This work shows how cortical columns can use multiple independent moving sensors to identify and locate objects.

II-21. Imbalanced amplification: A mechanism of amplification from local imbalance of excitation and inhibition in cortical circuits

Christopher Ebsch CEBSCH@ND.EDU
Robert Rosenbaum RROSENB1@ND.EDU
University of Notre Dame

Understanding the relationship between external stimuli and the spiking activity of cortical populations is a central problem in neuroscience. Dense recurrent connectivity can produce counterintuitive response properties, raising the question of whether there are simple arithmetical rules for relating circuits' connectivity structure to their response properties. One such arithmetic is provided by the mean-field theory of balanced networks, which is derived under an assumption that excitatory and inhibitory synaptic currents precisely balance on average. However, balance may not be so precise in cortex and balanced network theory is not applicable to some biologically relevant connectivity structures. We show that cortical circuits with such structure are susceptible to an amplification mechanism caused by a break in excitatory-inhibitory balance at the level of local subpopulations that does not necessarily break global balance. We find that a linear correction to the classical mean-field theory of balanced networks corrects for imprecise balance and quantifies this amplification. We show that this mechanism of "imbalanced amplification" explains several response properties observed in somatosensory and visual cortical circuits.

II-22. Learning with precise spike times: A new approach to select task-specific neurons

Dorian Florescu DORIAN.FLORESCU@SHEFFIELD.AC.UK
Daniel Coca D.COCA@SHEFFIELD.AC.UK
The University of Sheffield

There is great interest to understand the circuits that modulate behaviour such as the mirror neurons that fire during action observation as well as action execution. Typically, given the neuron firing rates measured from a large pool of candidate neurons, the challenge is to identify a subset of neurons that selectively respond during the execution of a task [1]. Given the considerable evidence that biological neurons can generate spikes with millisecond temporal precision [2], we argue that the problem of identifying task-specific neurons can be better addressed by considering the exact timings of the spike sequences instead of the firing rates. Here we propose a new Liquid State Machine (LSM) architecture, where the Readout unit is spike time based, and a new training algorithm that implements orthogonal forward selection to identify the best synaptic connectivity for the Readout. The learning algorithm, which is formulated in the Hilbert space of spike trains [3] with the LSM Readout defined as an inner product in this space, can be used not only for task-specific neuron identification but also for conventional machine learning applications. A machine learning example, involving the classification of jittered spike trains, demonstrates that the introduction of the neuron selection step improves classification accuracy compared to the standard method [4] and requires fewer synaptic connections to the Readout. A task-specific neuron identification example shows that using precise spike timings instead of firing rates leads to the identification of a much smaller set of key neurons that are relevant to the task.

II-23. Adaptive stimulus selection for optimizing neural population responses

Benjamin Cowley[1] BCOWLEY@CS.CMU.EDU
Ryan Williamson[2] RCW30@PITT.EDU
Katerina Acar[2] KAC216@PITT.EDU
Matthew Smith[2] SMITHMA@PITT.EDU
Byron Yu[1] BYRONYU@CMU.EDU
[1]Carnegie Mellon University
[2]University of Pittsburgh

Adaptive stimulus selection methods in neuroscience have primarily focused on maximizing the firing rate of a single recorded neuron. When recording from a population of neurons, it is usually not possible to find a single stimulus that maximizes the firing rates of all neurons. This motivates optimizing an objective function that takes into account the responses of all recorded neurons together. We propose “Adept,” an adaptive stimulus selection method that can optimize population objective functions. In simulations, we confirmed that population objective functions elicited more diverse stimulus responses than single-neuron objective functions. We then tested Adept in a closed-loop electrophysiological experiment in which population activity was recorded from macaque V4, a cortical area known for mid-level visual processing. To predict neural responses, we used the outputs of a deep convolutional neural network model as feature embeddings. Natural images chosen by Adept elicited mean neural responses that were 20% larger than those for randomly-chosen natural images, and also evoked a larger diversity of neural responses. Such adaptive stimulus selection methods can facilitate experiments that involve neurons far from the sensory periphery, for which it is often unclear which stimuli to present.

II-24. Dynamical structure of cortical taste responses revealed by precisely-timed optogenetic perturbation

Narendra Mukherjee NARENDRA@BRANDEIS.EDU
Joseph Wachutka WACHUTKA@BRANDEIS.EDU
Donald Katz DBKATZ@BRANDEIS.EDU
Brandeis University

Variability is a hallmark of neural responses to sensory stimuli. This variability is almost ubiquitously treated as nuisance noise; it is common practice to average neural responses across repeated presentations of the same stimulus, thereby minimizing and “dismissing” this trial-to-trial variability. However, in the context of a naturalistic decision to swallow or expel a taste in the mouth, we recently demonstrated that inter-trial variability in responses of taste cortex neural ensembles correlates strongly with the latency of ingestion-egestion orofacial (mouth) movements (Sadacca, Mukherjee, et al., 2016). Sudden transitions of sensory cortical ensemble activity into a previously-described “palatability-related response state” despite emerging as early as 0.5 sec post-stimulus in some trials and 1.5 sec post-stimulus in others, reliably preceded choice behavior by 0.2-0.3 sec. For the current work, we probe the rich and behaviorally-relevant dynamic structure in these ensemble taste responses using a combination of precisely timed optogenetic perturbations, chronic multi-electrode recordings, jaw EMG and probabilistic graphical modelling. We show that brief (i.e., 0.5-sec) optogenetic perturbation of a random subset of taste cortex neurons affects the timing of the animal’s orofacial expression of the swallow/expel decision—but only on trials where the perturbation arrives before the neural population shifts into decision-related firing. Early perturbations delay this decision whereas late perturbations have no impact; identical perturbations delivered during the heart of taste processing, meanwhile, had no impact on trials in which the ensemble state had already emerged, and strongly delayed decisions on trials in which it had not. These results provide evidence for a distributed sensory-motor attractor network in taste processing, characterized by stochastic shifts into a behaviorally-relevant stable state. Perturbations to this network can only delay the inevitable relaxation into the stable state and expression of behavior, pointing at a truly dynamical role for cortex in consumption decisions.

II-25. Orchestration of cortico-cortical synchronization by the visual thalamus during visual processing

Wei Huang ANGELVV@EMAIL.UNC.EDU
Ehsan Negahbani NEGAHBAN@EMAIL.UNC.EDU
Charles Zhou ZHOUZ@EMAIL.UNC.EDU
Flavio Frohlich FLAVIO_FROHLICH@MED.UNC.EDU
University of North Carolina at Chapel Hill

Visual processing requires a concerted effort from a network comprised of both cortical and subcortical structures. As a key player in this network, posterior parietal cortex (PPC) interacts with multiple cortical regions and serves as an essential interface between sensory processing and motor response. This cortico-cortical interaction is further gated by thalamo-cortical rhythms. However, the mechanism underlying how the thalamus affects functional interactions between the cortical areas during visual processing remains unclear. Here, we studied this question by simultaneously recording from three reciprocally connected nodes in the thalamo-cortical visual circuit: PPC, the primary visual cortex (V1) and the lateral nucleus of the lateral posterior thalamic nucleus (LPl), in ferrets performing visual detection tasks. Ferrets were chosen for their well-developed visual system. To our knowledge, this is the first study to perform extracellular simultaneous recording from these three regions in freely moving animals. Our recordings revealed an increase in phase-locking value around the theta frequency band (4-8 Hz) between the three nodes during the visual stimulation period. Conditional Granger causality further demonstrated that this functional connectivity resulted from LPl driving PPC and V1. In addition, the PPC-V1 coherence in the gamma frequency band (25-40 Hz) was modulated by the thalamic theta phase. As a whole, our study revealed that LPl enhanced the cortical functional connectivity between PPC and V1 during visual processing possibly via theta-gamma coupling. Our findings provide a circuit-level mechanistic explanation on how a higher-order visual thalamic structure modulates cortical communication, and may represent a fundamental link between neural network activity and behavior.

II-26. Intermingled ensembles coding stimulus identity and expected value in visual association cortex

Rohan Ramesh RAMESH@FAS.HARVARD.EDU
Christian Burgess CBURGESS@BIDMC.HARVARD.EDU
Arthur Sugden ARTHUR.SUGDEN@GMAIL.COM
Michael Gyetvan MGYETVAN@BIDMC.HARVARD.EDU
Mark Andermann MANDERMA@BIDMC.HARVARD.EDU

Harvard University

The selective enhancement of motivationally-relevant learned stimuli in the brain is conserved across many species. Motivational states such as hunger guide attention and bias perception to food-cues that help maintain homeostasis. In human fMRI studies, visual association cortex shows enhanced responses to food-cues in hungry, but not sated, participants. Previously, in rodents, we showed that postrhinal cortex, a retinotopically-organized visual association cortex, contains neurons encoding stimulus identity in naive animals and, following training in a Go-NoGo visual discrimination task, shows a selective hunger-dependent bias towards a motivationally valuable food-cue. Training on simple cue-outcome associations, however, cannot differentiate coding of stimulus identity from predicted outcome. To address this, using chronic two-photon calcium imaging, we tracked the same neurons in mouse visual association cortex across successful re-learning of a novel combination of cue-outcome associations in a Go-NoGo visual discrimination task. Previous studies in naive, non-behaving animals, suggest that stimulus identity can be encoded in the joint activity of neurons with similar visual response characteristics. We found that the encoding of stimulus value in visual association cortex is not mapped onto these ensembles that encode stimulus identity. Instead, a unique, intermingled ensemble of visually-responsive neurons track predicted outcome. These 'value-coding' neurons formed an ensemble that exhibited elevated within-ensemble noise correlations and spontaneous correlations. This ensemble represented value on the timescale of seconds (reward history), hours (across states of hunger/satiety), and days (during remapping of contingencies). Furthermore, neurons that become responsive during or following learning of the new cue-outcome associations were integrated into this value-coding ensemble. We propose that distinct ensembles representing stimulus identity and value in visual association cortex may achieve dual goals of maintaining a veridical representation of a visual stimulus while regulating its motivational salience in a state-dependent manner.

II-27. Two-photon imaging of offline reactivation in visual association cortex following associative learning

Arthur Sugden[1] ARTHUR.SUGDEN@GMAIL.COM
Lauren Sugden[2] LAUREN.V.SUGDEN@GMAIL.COM
Andrew Lutas[1] ALUTAS@BIDMC.HARVARD.EDU
Rohan Ramesh[1] RAMESH@FAS.HARVARD.EDU
Kelly McGuire[1] KMCGUIRE@G.HARVARD.EDU
Osama Alturkistani[1] OALTURKI@BIDMC.HARVARD.EDU
Christian Burgess[1] CBURGESS@BIDMC.HARVARD.EDU
Mark Andermann[1] MANDERMA@BIDMC.HARVARD.EDU

[1]Harvard University
[2]Brown University

Our understanding of sensory and association cortices is largely limited to neuronal responses to external stimuli, but spontaneous, offline cortical processing likely serves important functions. Here, we have identified reactivation of visual cues associated with appetitive, aversive, or neutral outcomes in layer 2/3 of mouse lateral visual association cortex. Our previous work demonstrated that visual association cortex exhibits motivation-dependent response enhancement of specific sensory cues, and encodes both visual cues and their associated outcomes (e.g. delivery of reward or punishment).

Using two-photon calcium imaging recordings of visual association cortex, we trained a Bayesian classifier on single-trial neuronal population responses to presentations of different visual cues that can identify moments of reactivation of those cues. We applied this classifier to long periods of post-task spontaneous activity in darkness and found reactivation of cue-specific response patterns during quiet waking, which were not present in shuffled controls. Remarkably, we found that visual cue representations associated with the appetitive or aversive outcomes are reactivated at a rate of more than two-fold higher than those associated with neutral outcomes. This finding persisted after changing cue-outcome contingencies, suggesting that reactivation tracks recent salient experiences. We are testing whether these behaviorally-relevant reactivations of cue representations drive offline changes in functional connectivity between neurons encoding cues and those encoding associated outcomes.

II-28. Population-level but not neuron-level similarity during movement of the contra- vs ispi-lateral hand

Katherine Cora Ames KCA2120@COLUMBIA.EDU
Mark Churchland MC3502@COLUMBIA.EDU

Columbia University

The motor system has lateralized outputs: the motor cortex (M1) of each hemisphere projects primarily to the contralateral body. Lateralization of computation is less clear-cut: M1 neurons can be active during movements of either arm. Does each hemisphere contribute differently to movements of the contralateral (driven) and ipsilateral (non-driven) arm?

We trained a monkey to progress through a virtual environment by cycling a right or left pedal with the corresponding hand. Simultaneous recordings (24-channel V-probes) yielded 178 left-hemisphere and 190 right-hemisphere units across days. Median firing rate range was almost as high (86%) during movements of the non-driven versus driven arm. However, a neuron's activity pattern often differed between movements of each arm. Furthermore, neurons that shared response patterns when moving one arm had little tendency to do so when moving the other arm. In fact, population activity occupied largely orthogonal subspaces during movements of each arm. Thus, while M1 was active during movements of either arm, there was a dramatic remapping of individual unit activity.

Might that remapping reflect a change in the information present in the population? We first asked whether the driving (contralateral) cortex was more informative about muscles than the non-driving (ipsilateral) cortex. In fact, the driving and non-driving cortices predicted muscle activity equally well (generalization R^2: 0.85 and 0.83 respectively). We used a variety of other methods to ask whether there were signals in the driving cortex that were absent in the non-driving cortex. We found no evidence for that hypothesis. For example, a given unit's activity could be equally well-predicted by other units from the same or opposite hemisphere. Thus, even during unimanual movements, all major patterns in the driving cortex are also present in the non-driving cortex. This suggests that while motor cortical outputs are lateralized, the computation may be more distributed.

II-29. Neural mechanisms of flexible internal representations

Timothy Muller[1] TIMOTHY.MULLER@MAGD.OX.AC.UK
Tim Behrens[2,1] BEHRENS@FMRIB.OX.AC.UK
Jill O'Reilly[1] JILL.OREILLY@PSY.OX.AC.UK

[1]University of Oxford
[2]University College London

Flexible adaptation to the world is essential for survival and a hallmark of intelligent behaviour. This flexibility in turn requires the internal representations of the environment that guide our choices to be flexible to changes in the world. While signals associated with detecting change in the world have been identified, much less is known about the neural representations of the current task and how they are updated when change is detected. Here we show that different brain regions are responsible for detecting change and representing the current task. We additionally show responses of pupil-linked systems—thought to index noradrenaline—reflect detection of change in the environment and are explained by signals in brain regions associated with change detection. Finally, we show changes in pupil diameter following evidence that the world has changed and the current belief may require revision explain changes to the strength of the neural representation of the current belief. These findings offer insight in to how organisms can adapt their internal representations to accurately reflect a changing world.

II-30. Sparse linear recombination using most retinal output channels yields highly diverse visual representations in mouse dLGN

Yannik Bauer[1] Y.BAUER@BIO.LMU.DE
Miroslav Roman Roson[2] MIROSLAV.ROMAN.ROSON@GMAIL.COM
Philipp Berens[2] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Thomas Euler[2] THOMAS.EULER@CIN.UNI-TUEBINGEN.DE
Laura Busse[1] BUSSE@BIOLOGIE.UNI-MUENCHEN.DE

[1]Ludwig-Maximilians-University Munich
[2]University of Tuebingen

More than 30 functional types of retinal ganglion cells (RGCs) compute in parallel distinct features of the visual world and send this information to the brain. Little is known, however, about which RGC types project to the dorsolateral geniculate nucleus (dLGN) of the thalamus, and how the different RGC channels recombine there. Interest in these questions has been fuelled by recent estimates of retinogeniculate convergence obtained by anatomical work, which far exceeded those obtained in electrophysiological recordings.

To get insights into the nature of retinal input to dLGN, we conditionally expressed the genetically encoded Ca2+ indicator GCaMP6f in dLGN-projecting (dLGN-p) RGCs, followed by in vitro retinal two-photon Ca2+ imaging of light-evoked responses. Using the same stimulus set as an earlier RGC classification by Baden et al. (2016), we compared the responses of each dLGN-p RGC to those of the previously described RGC types and identified the RGC population cluster with the best-matching response properties. We found that most functional RGC types seem to innervate dLGN, with certain types, such as ON- and OFF alpha cells or OFF suppressed cells, showing clear overrepresentations.

Using in vivo extracellular multi-electrode recordings in awake, head-fixed mice, we then recorded the responses of dLGN neurons to the same visual stimuli. We quantitatively assessed the degree of diversity in the dLGN responses by using sparse non-negative matrix factorization (NNMF), which decomposed the dLGN population response into a rich and highly diverse set of 25 components.

Finally, we modelled dLGN cell responses as a sparse linear combination of retinal input types. We found that responses of dLGN neurons could be best predicted by inputs from 1-7 RGC types.

In summary, our study reveals that most mouse RGC types project to dLGN, which yields an unexpectedly diverse representation that can be modelled by a sparse feedforward model.

II-31. Synergistic processing of stimulus- and state-dependent features in two corticothalamic cell types

Ross Williamson ROSS.S.WILLIAMSON@GMAIL.COM
Daniel Polley DANIEL_POLLEY@MEEI.HARVARD.EDU

Harvard University

Neurons in the deep layers of the auditory cortex (ACtx) give rise to a massive subcortical projection that innervates all levels of the central auditory pathway as well as non-auditory areas including the amygdala and striatum. While these targets play distinct roles in behavior and cognition, the information conveyed to them from upstream neurons is not well understood. Here, we performed anatomical tracing, cell-type-specific single-unit recording, and in-vivo calcium imaging from two classes of ACtx projection neurons; L5-corticocollicular neurons (L5CCol), and L6-corticothalamic neurons (L6CT). We used an intersectional viral strategy to selectively express a fluorophore in both neuron classes. While L6CT neurons primarily targeted the thalamus, L5CCol neurons had more divergent connectivity, with collateral axons also targeting the thalamus, and multiple downstream structures. To characterize the information conveyed to these structures, we then used antidromic optogenetic "phototagging" to isolate individual L5CCol and L6CT units on high-density multichannel probes in awake, head-fixed mice. We found that L6CT neurons were more selective to sound features and had more linear receptive fields, when compared with L5CCol neurons. A connectivity analysis revealed that L6CT neurons powerfully regulate response gain across the deep layers of the ACtx via dense local connections with excitatory and inhibitory subnetworks, while L5CCol neurons connectivity was much weaker. Recent studies have shown that an animal's "internal state" exerts a powerful influence over ACtx excitability. We expressed GCaMP6s in L5CCol and L6CT neurons and noted that motor signals suppressed the responses of L5CCol neurons, but enhanced the responses of L6CT neurons. This suggests differences in both local and long-range inputs onto L6CT and L5CCol neurons, which will be the subject of future work. Collectively, these studies show that each class of projection neuron performs distinct operations on internal and external signals, which likely impart distinct effects on their downstream targets.

II-32. Local online learning with random feedback in recurrent networks

James Murray JM4347@COLUMBIA.EDU
Laurence F. Abbott LFA2103@COLUMBIA.EDU

Columbia University

According to numerous experimental studies, the brain performs computations that unfold over time. Building on this observation, a longstanding challenge for computational neuroscience has been the development of biologically plausible learning rules for recurrent neural networks (RNNs) enabling the production of time-dependent signals such as those that might drive movement or facilitate working memory. Classic gradient-based algorithms such as backpropagation through time (BPTT) have been available for decades, but are inconsistent with known biological features of the brain, such as causality and locality. More biologically plausible learning rules have been proposed, but these retain some undesirable requirements such as symmetry between the synaptic readout and feedback weights, as well as restrictions on allowable mappings between RNN inputs and outputs.

We derive an approximation to gradient-based learning that overcomes these limitations while comporting with biological constraints. Specifically, the online learning rule for the synaptic weights involves only local information about the pre- and postsynaptic activities, in addition to a random feedback projection of the RNN output error, and enables an RNN to produce arbitrary time-dependent outputs. In addition to providing mathematical arguments for the effectiveness of this new learning rule, we show through simulations that it attains similar performance to standard algorithms such as BPTT on a variety of tasks.

II-33. Using deep learning to reveal the neural code for images in primary visual cortex

William Kindel[1] KINDEL.WILL@GMAIL.COM
Elijah Christensen[2] ELIJAH.CHRISTENSEN@UCDENVER.EDU
Joel Zylberberg[1] JOEL.ZYLBERBERG@GMAIL.COM

[1]University of Colorado
[2]University of Colorado School of Medicine

Primary visual cortex (V1) is the first stage of cortical image processing, and a major effort in systems neuroscience is devoted to understanding how it encodes information about visual stimuli. Within V1, many neurons respond selectively to edges of a given preferred orientation: these are known as simple or complex cells. Other neurons respond to localized center-surround image features. Still others respond selectively to certain image stimuli, but the specific features that excite them are unknown. Moreover, even for the simple and complex cells, it is challenging to predict how they will respond to natural image stimuli. Thus, there are important gaps in our understanding of how V1 encodes images. To fill this gap, we train deep convolutional neural networks to predict the firing rates of 355 V1 neurons in response to natural image stimuli, and find that the network predicts firing rates that are highly correlated ($r = 0.56 \pm 0.02$) with the neurons' actual firing rates. In contrast to shallow models, such as linear-nonlinear models that can only describe simple cells, we find that our convolutional neural network can describe a broad range of cells. We find that the firing rates of both orientation-selective, and non-orientation-selective neurons can be predicted with high accuracy. To advance our understanding of the visual processing that takes place in V1, we invert the network to find visual features that cause individual cells to spike. In this process, we identify canonical V1 features and new features that warrant further investigation.

II-34. Ongoing, rational calibration of reward-driven perceptual biases

Yunshu Fan YUNSHUF@PENNMEDICINE.UPENN.EDU
Joshua Gold JIGOLD@PENNMEDICINE.UPENN.EDU
Long Ding LDING@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Decision-making is often interpreted in terms of normative computations that can achieve goals like maximizing accuracy or a particular reward function. Here we show that these normative principles alone do not explain how individuals form decisions that require combination of both sensory and reward information. We analyzed the choice and reaction time (RT) performance of three monkeys on an asymmetric-reward perceptual decision-making task. For this task, the monkey decided the global motion direction of a random-dot kinematogram and indicated its choice with a saccade made at a self-determined time. We manipulated perceptual uncertainty in each trial by presenting motion stimuli with varying strengths and manipulated reward associations in each block of trials by assigning a large reward for one direction and small reward for the other. Although the monkeys' choice and RT behavior is well fitted by an accumulate-to-bound decision process that can maximize reward rate, each monkey used different reward-dependent adjustments that were: 1) suboptimal but with a consistent, common pattern of deviations from the optimal solution; 2) good enough to allow each monkey to obtain nearly optimal reward outcomes, given its perceptual sensitivity; and 3) adaptive to variations in reward asymmetry and perceptual sensitivity across sessions. We propose a rational, satisficing process that uses heuristic knowledge of reward function properties to reach these good-enough solutions. Our simulations show that such a process can capture key features of the monkeys' adjustments and raise the possibility that individual differences in decision-making could reflect variations in heuristic knowledge in otherwise rational subjects. These results suggest new dimensions for assessing the rational and idiosyncratic nature of flexible decision-making and constraints on the underlying neural implementation.

II-35. Higher-order response statistics is shaped by stimulus content in the visual cortex

Gergo Orban[1] ORBAN.GERGO@WIGNER.MTA.HU
Mihaly Banyai[1] HAGYMASBAB@GMAIL.COM
Andreea Lazar[2] ANDREEALAZ@GOOGLEMAIL.COM
Hanka Klon-Lipok[2] JOHANNA.KLON-LIPOK@BRAIN.MPG.DE
Wolf Singer[2] WOLF.SINGER@ESI-FRANKFURT.DE

[1]MTA Wigner Research Center for Physics
[2]Ernst Strungmann Institute for Neuroscience

Spike count correlations (SCCs) in the visual system have a rich structure and display intricate stimulus dependence. Yet, it has remained unclear how stimulus-dependence can be predicted by the content of stimuli, defined as higher-order statistical structure. We argue that in a hierarchical model of natural images, reflecting the hierarchical organization of the ventral stream, stimulus-dependence of SCCs is a natural consequence of perceptual inference. Assuming that neuronal activities along the hierarchy represent the presence or absence of increasingly complex stimulus features, statistical inference of high-level features influences the inference of low-level features by establishing a context, or in statistical terms, a local prior. This implies the top-down modulation of both means and correlations in the response statistics. We designed experiments to measure the fine structure of SCCs and to assess the dependence of patterns in correlations on stimulus identity and on stimulus statistics. Measuring multiunit activity from a population of V1 neurons in awake monkeys, we show that the fine structure of SCCs is specific to the identity of natural images. We develop a non-parametric procedure to control for confounds in SCC specificity introduced by modulations in firing rates. Further, we demonstrate that, as predicted by hierarchical inference, stimulus-dependent SCCs are characteristic of natural images, but less so of synthetic images only characterized by low-order statistical structure. Using raster marginal models, we show that measured SCC specificities cannot be accounted for by finite data effects. Finally, we demonstrate that the higher stimulus-specificity of SCCs can be recovered by introducing texture-like higher-order structure in synthetic stimuli. Our results demonstrate that top-down influences resulting from hierarchical inference can predict stimulus-dependent modulations of fine-scale correlations in V1.

II-36. Reorganization of cortical population neuronal activity following auditory fear conditioning

Katherine Wood KATHERINECWOOD@GMAIL.COM
Danielle Bassett DSB@SEAS.UPENN.EDU
Maria Geffen MGEFFEN@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Auditory perception relies on learning-driven neuronal plasticity within the auditory pathway. Here, we investigated how associative learning, differential auditory fear conditioning (DAFC), affects neuronal population responses to sounds in auditory cortex (AC). In DAFC, the subject is presented with two different frequency tones, one of which is paired with a footshock. Previously, we found that AC is required for expression of DAFC-driven changes in sound-frequency discrimination acuity (Aizenberg and Geffen, 2013) and that modulating inhibitory neuronal activity in AC leads to similar bi-directional changes in discrimination acuity (Aizenberg, 2015). However, how DAFC affects tone-evoked population neuronal activity remained unknown. We hypothesized that DAFC would drive changes in population tone-evoked neuronal activity corresponding to either and increase or a decrease in neurometric frequency discrimination acuity, as a function of fear learning specificity. To understand the transformation of sound representation in AC before and after DAFC (Figure 1), we imaged calcium activity in hundreds of neurons simultaneously in AC of awake, head-fixed mice, tracking the same neurons over days under a two-photon microscope (Figure 2) before and after two DAFC sessions. We quantified changes in tone frequency-dependent responses of individual neurons, as well as in population functional connectivity. DAFC drove heterogeneous

changes in individual neuronal responses for either paired or unpaired tone frequencies (Figure 3). At the same time, mean population neuronal response strength to tones across frequencies was preserved (Figure 2). However, neuronal responses to tones following DAFC became more consistent after DAFC (Figure 4). Neuronal populations formed clusters driven by correlated activity, neurons within clusters exhibit heterogeneous response patterns. The neuronal cluster structure changed between days in the absence of DAFC, but the network structure became more consistent over days following DAFC (Figure 5). These findings suggest that DAFC drives cortical population activity toward a more stable state.

II-37. Orthogonal preparatory and movement subspaces in monkeys, mice, and an inhibition-stabilized network

Ta-Chu Kao TCK29@CAM.AC.UK
Mahdieh Sadabadi MSS69@CAM.AC.UK
Guillaume Hennequin G.HENNEQUIN@ENG.CAM.AC.UK

University of Cambridge

In what dynamical regime does the motor cortex operate? Population recordings in monkey motor cortex during delayed reaching tasks have revealed a surprising relationship between preparatory and movement-related activity trajectories: they unfold in orthogonal subspaces (Elsayed et al., 2016). Specifically, the activity subspace that captures most of the across-movement variance in neural activity during movement preparation captures little variance during movement execution - and vice versa.

To begin to understand whether this effect is a useful diagnostic feature of the collective dynamics of motor cortical areas and not a mere epiphenomenon, we examined the relationship between preparatory and movement-related activity in rodents during a delayed tactile discrimination task (data courtesy of Karel Svoboda). Mice were trained to lick left or lick right depending on the location of a pole presented during a brief “sampling“ epoch. Mice had to wait for a certain duration before licking, presumably using this delay period to prepare their motor output. We analysed recordings made in the anterior lateral motor cortex (ALM) using the same population analysis techniques as previously applied to monkey data. We found the same effect: a strong decoupling of ALM delay and movement-related activities, in the eight mice that we analysed.

Next, we show that this effect is naturally captured by the dynamics of recurrent networks with strong and intricate excitatory (E) connections stabilised by detailed inhibitory (I) feedback. In previous work (Hennequin et al, 2014), we have shown that such networks produce naturalistic transients following initialisation in specific states characterized by broken E/I balance. Upon movement onset, activity rotates away from these “sensitive preparatory states“ as it grows bigger in amplitude during the movement epoch, rapidly entering an orthogonal subspace where excitation and inhibition re-balance. This is a highly robust feature of this class of models.

II-38. State-aware control of neural activity: design & analysis

Adam Willats[1,2] AWILLATS3@GATECH.EDU
Michael Bolus[1,2] MFBOLUS@GATECH.EDU
Clarissa Whitmire[1,2] CLARISSA.WHITMIRE@GATECH.EDU
Garrett Stanley[1,2] GARRETT.STANLEY@BME.GATECH.EDU
Christopher Rozell[2] CROZELL@GATECH.EDU

[1]Emory University
[2]Georgia Institute of Technology

State-dependent variability in neural coding impacts activity from single cell spiking to behavior, but nonstationary coding properties make it challenging to study the relationships between external stimuli, neural responses and perception. Closed-loop stimulation strategies promise to provide causal manipulations that allow more direct investigation of neural circuit function, and ongoing work from our lab demonstrates the ability to precisely control time-varying firing rates of single-units in-vivo using optogenetics and extracellular recording. However, the nonstationarities of state-dependent neural coding also present significant challenges to designing closed-loop stimulation systems.

Here we present and characterize state-aware optogenetic control of time-varying spike rates that compensates for state-dependent neural responses. This approach consists of three interacting components: (1) a controller to generate new inputs (2) an observer to estimate latent dynamics from spikes and (3) a decoder for inferring state. We designed these components around switched linear dynamical systems models fit to spiking responses from single units recorded in-vivo under optogenetic input.

Through a combination of simulation and models fit to experimental data, we outline scenarios where state-aware control is likely to outperform state-naive approaches. Specifically, we parameterized state-switches as changes in gain and latency of a cell's spiking response. We find that state-aware control is more robust than state-naive control, facilitating effective control across switches seen in neural systems.

We present a systematic characterization and derive principles for tuning the decoder, controller, and observer to achieve tracking of time-varying targets (with frequency content up to 20Hz) that is robust to state-switching and amenable to implementation in existing real-time hardware. Taken together this work extends the toolbox of closed-loop approaches to neuroscientific investigation and facilitates the study of how state-dependent neural activity influences perception and behavior.

II-39. A method for finding neural correlates of behavior in regions and networks from large-scale recordings

Macauley Breault[1] MBREAUL1@JHU.EDU
Juan Bulacio[2] BULACIJ@CCF.ORG
Pierre Sacre[1] P.SACRE@JHU.EDU
Matthew Kerr[3] MATTHEW.SD.KERR@GMAIL.COM
Jorge Gonzalez-Martinez[2] GONZALJ1@CCF.ORG
Sridevi Sarma[1] SSARMA2@JHU.EDU
John Gale[4] JOHN.T.GALE@EMORY.EDU

[1]Johns Hopkins University
[2]Cleveland Clinic
[3]Abbott
[4]Emory University

Recent technology has allowed us to record more neural data than neuroscientists are able to analyze using traditional statistical methods. In particular, finding neural correlates of behavior from large data sets, a goal of many systems neuroscientists, is challenging. When encountered with a large data set generated from high-resolution whole brain recordings, the typical analysis approach is to limit the investigation to single brain regions and study spectral characteristics of neural activity in specific frequency bands. However, one may miss important interactions with such limited analyses [1]. Ideally, one would like to extract all the brain regions that are related to behavior without predefining frequency bands. Here, we propose a technique that relates local field potentials (LFPs) from whole brain recordings to dynamic behavioral information to quantitatively extract a subset of brain regions or networks of interest. This method uses singular value decomposition (SVD) to reduce thousands of spectrograms (matrices) of neural data from each contact into scalar time-varying "reduced spectral" signals. Each contact‘s reduced spectral signal can then be correlated with a time-varying behavioral signal per trial and per subject where highly correlated regions can be further studied. This method can reduce searching over thousands of recordings taken from hundreds of brain structures to a short list of regions that encode behavior. We demonstrate the application of this method on a data set consisting of LFPs using stereoelectroencephalography (SEEG) from 9 epileptic subjects implanted with 680 total depth electrodes covering 79 brain areas executing a goal-directed center-out reaching task [2]. The goal is to find neural correlates of the dynamic error of movement

during the reach. Using our methodology, we identified subsets of highly-correlated contacts, primarily located in limbic and visual regions. This preliminary result suggests the presence of emotional integration in visual feedback to correct movements that go off-course.

II-40. Disentangling evidence integration from memoryless strategies in perceptual decision making

Gabriel Stine GABRIEL.STINE@COLUMBIA.EDU
Ariel Zylberberg AZ2368@CUMC.COLUMBIA.EDU
Michael Shadlen SHADLEN@COLUMBIA.EDU

Columbia University

The neurobiology of perceptual decision making offers an experimental approach to the study of cognition in part because it involves the integration of evidence over time-scales that are longer than those of sensory neurons. However, many perceptual decisions do not require prolonged integration times, and even if evidence integration would be helpful (or optimal), it is possible that a decision maker would employ alternative strategies. This concern is exacerbated in animal models because explicit instructions cannot be given. Therefore, an essential logical step before interpreting neural activity in the context of evidence accumulation is to confute these alternative strategies. We considered a number of testing paradigms, behavioral analyses, and model-fitting exercises to identify factors that implicate integration. We found that integration and non-integration strategies often make surprisingly similar quantitative predictions. For example, we show that fixed-duration tasks combined with stimulus-aligned psychophysical kernels are unable to distinguish integration from non-integration strategies. However, we also identified conditions in which these strategies make qualitatively different predictions. We leveraged this finding to test the hypothesis that subjects would employ non-integration strategies when a stream of evidence contains compelling examples. Naive subjects discriminated motion in both a random dot stimulus and a recently developed stimulus consisting of pulsed Gabor elements (Katz et al, 2016; Yates et al, 2017). Preliminary data from the pulsed Gabor task suggest some subjects may use non-integration strategies. The observation supports the hypothesis that decision-makers can adopt non-integration strategies depending on the task being performed, even if such a strategy is suboptimal. Our results have important implications for those interpreting neural activity in the context of evidence accumulation tasks and serve as a guide for ruling out alternative strategies.

II-41. Monitoring behavior and neural activity in freely moving mice with head-mounted cameras and implants

Arne Meyer[1,2] ARNE@GATSBY.UCL.AC.UK
Jasper Poort[1] J.POORT@UCL.AC.UK
John O'Keefe[1] J.OKEEFE@UCL.AC.UK
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK
Jennifer Linden[1] J.LINDEN@UCL.AC.UK

[1]University College London
[2]Gatsby Computational Neuroscience Unit

Detailed behavioral monitoring is crucial to understand the link between neural activity and natural behavior in freely behaving animals. The mouse is a prominent model animal because of the availability of unique methods for measuring and manipulating neural circuits, but detailed behavioral control in freely moving mice has been limited by the absence of video tracking methods in head-centered coordinates. Here we overcome this limitation by developing a miniature ultra-lightweight head-mounted camera system for mice (weight 1.3 grams) combined with movement sensors and chronic electrophysiology recordings. The camera does not affect neural recording quality and generates stable video recordings, and mouse behavior is similar with and without the tracking system.

We demonstrate the potential of the system in a series of experiments in freely moving mice. First, we show that behavioral variables such as whisking frequency and pupil size vary systematically with behavioral state, and that these changes are correlated with neural activity, thereby generalizing results obtained in head-restrained mice to natural behaviors (Reimer et al. 2014; McGinley et al. 2015). Second, we demonstrate that a large fraction of variability in eye position in freely moving mice is explained by head movements, as has also been observed in rats (Wallace et al. 2013). Our data further indicates that mice stabilize their gaze with respect to the horizontal plane, and that this stabilization does not depend on visual input. Third, we describe a novel form of modulation of neural activity in primary visual cortex (V1), where responses are strongly modulated by head movements even in the absence of visual input. This effect was not explained by whisking or variability in eye movements. These results demonstrate how the new camera system can give novel insights into the interactions between different behaviors and their relation with neural activity.

II-42. Computational model for memory as a log-compressed timeline constrained by cognitive and neural data

Zoran Tiganj ZORAN.TIGANJ@GMAIL.COM
Inder Singh MAILINDERDEEP@GMAIL.COM
Marc Howard MARC777@BU.EDU

Boston University

Behavioral experiments suggest that working memory can express not only the identity of recently presented stimuli, but also an estimate of the elapsed time since each of the stimuli was presented. Previous studies of the neural underpinnings of working memory have focused on sustained firing and recurrent dynamics. Sustained firing obscures information about the elapsed time. Recurrent dynamics can retain information about the identity of the stimulus and elapsed time, but in general decoders must be learned, which can be computationally costly. Here we use cognitive and neural data to constrain a computational model of working memory. We first show that in the short-term judgment of recency (JOR) task, the time to access an item in the past depends only on the recency of the more recent probe, suggesting that subjects have access to a temporally ordered memory representation. Moreover, the relationship between the response time and the item recency was well described with a logarithmic function, indicating the compression of the memory representation. Furthermore, we show neural evidence suggesting that the brain maintains a conjunctive representation of what and when, with the temporal dimension showing logarithmic compression. Reanalysis of single-unit recordings from the macaque lateral prefrontal cortex (lPFC) during performance of a delayed-match-to-category task shows that each sample stimulus triggered a consistent sequence of neurons that encoded both stimulus identity and the elapsed time. The encoding of the elapsed time became less precise as the sample stimulus receded into the past. Finally, we provide a neurally plausible computational model that gives rise to a scale-invariant log-compressed timeline. The model is constructed as a two-layer feed-forward neural network with analytically computed weights. Because this network is linear, it is not necessary to learn a new decoder for each stimulus. This property is essential in constructing models of human working memory.

II-43. Inhibitory control of neuronal tuning in an attractor model of visual cortex

Angus Chadwick[1,2] ANGUS.CHADWICK@UCL.AC.UK
Thomas Mrsic-Flogel[3] THOMAS.MRSIC-FLOGEL@UNIBAS.CH
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]Gatsby Computational Neuroscience Unit
[3]University of Basel

The tuning of cortical neurons to sensory stimuli varies with behavioural context and past experience, but it remains poorly understood how these changes arise from reorganisation in the cortical circuit. We have recently shown that learning-related changes in neuronal selectivity in mouse visual cortex are associated with the emergence of tuned excitatory-inhibitory interactions in the local microcircuit (Chadwick et al., Cosyne 2017). However, the statistical models of population dynamics used to infer such interactions are not biologically interpretable, and the consequences of modifying the tuning of inhibitory interactions in recurrent cortical circuits are not clear. In particular, orientation tuning has been extensively studied in simplified ring attractor networks comprising a single cell type (Ben-Yishai et al., 1995), but little is known about the relationship between connectivity, inputs and tuning curves in attractor networks obeying Dale's principle with separate excitatory and inhibitory populations. To address this, we varied the inhibitory connectivity and top-down inputs within such networks and measured the tuning of responses to external inputs. We found that, compared to networks with uniform inhibition, tuning curves are sharpened when interneurons preferentially target pyramidal cells with orthogonal orientation preferences to their own, but instead are broadened when interneurons target pyramidal cells with similar orientation preferences to their own. In networks with uniform inhibitory connectivity, we found that adding a uniform, task-dependent drive to interneurons sharpens tuning curves, consistent with a previous optogenetic study (Lee et al., 2012). However, in networks with tuned inhibition, top-down inputs must target local subpopulations of interneurons if tuning curves are to be sharpened. Our findings constrain the inhibitory connectivity profiles and top-down inputs that can account for selectivity changes with learning and behavioural context, and make testable predictions for neuronal tuning curves across behavioural conditions that further distinguish amongst alternative mechanisms.

II-44. Persistent neurons drive stable population-level working memory representations

Klaus Wimmer[1] WIMMER.KLAUS@GMAIL.COM
Joao Barbosa[2] PALERMA@GMAIL.COM
Adria Galan[3] GALAN95@GMAIL.COM
Christos Constantinidis[4] CCONSTAN@WAKEHEALTH.EDU
Gianluigi Mongillo[5] GIANLUIGI.MONGILLO@GMAIL.COM
Albert Compte[3] ACOMPTE@CLINIC.CAT

[1]CRM
[2]IDIBAPS
[3]IDIBAPS Barcelona
[4]Wake Forest Univ.
[5]CNRS Paris

Neurophysiological experiments in primates have found that during the delay period of working memory tasks, a fraction of neurons in the prefrontal cortex carries information about the stimulus as sustained activity, therefore supporting a stable code during the whole delay period. However, many neurons show strong temporal dynamics, which has given rise to the dynamic coding model for working memory. This model proposes that due to the time-varying dynamics of single neurons, a stable memory representation can only be achieved at the population level through a linear combination of individual neural responses of a sufficiently large population of neurons. Here we set out to investigate how prefrontal neurons with different delay-period dynamics contribute to population dynamics during an oculomotor delayed response task (Constantinidis et al. 2001). We first characterized the delay dynamics of single neurons based on their firing rate autocorrelation. Autocorrelation decays were heterogeneous, ranging from persistent neurons with slow decay to dynamic neurons with more transient delay activity autocorrelation. We extended the result of Murray et al. (2017) by analyzing how different neurons contribute to the principal components of the pseudo-population responses and found that the persistent neurons, but not the dynamic neurons, span a stable, low-dimensional mnemonic subspace. We then used linear decoders on single neurons and compared stimulus information during different time points throughout the whole trial period. Persistent neurons carried more information than dynamic neurons on any tested time point during the delay. Moreover, by combining single neuron recordings to pseudo-population responses we found that a small subset of 32/541 neurons with the highest individual cue and delay selectivity provides a stable representation throughout the trial, as accurate as the whole population of 541 neurons. In sum, we conclude that persistent neurons are the main drivers of memory-selective delay period dynamics in our data.

II-45. Generalisation of structural knowledge in the hippocampal-entorhinal system

James Whittington[1] JAMES.WHITTINGTON@MAGD.OX.AC.UK
Timothy Muller[1] TIMOTHY.MULLER@MAGD.OX.AC.UK
Tim Behrens[2,1] BEHRENS@FMRIB.OX.AC.UK
Caswell Barry[2] CASWELL.BARRY@UCL.AC.UK

[1]University of Oxford
[2]University College London

A central problem to understanding intelligence is the concept of generalisation. This allows previously learnt structure to be exploited to solve tasks in novel situations differing in their particularities. Here we propose that in order to generalise knowledge, the representations of the structure of the world, i.e. how entities in the world relate to each other, need to be separated from the representations of the entities themselves. We propose a novel model where structural information provided by higher order cortex (grid cells) and sensory information provided by sensory cortex is combined to form a conjunctive representation in hippocampus (place cells). The model is trained end-to-end in an artificial neural network that resembles a variational autoencoder, where the task is to predict the next vertex on multiple graphs that share the same structure but with shuffled vertices. For graphs with 2D structure, grid and place cell representations naturally emerge, without any training for navigation. Using data of simultaneously recorded grid and place cells in a remapping experiment, we test the prediction that grid cells impose a structural constraint on place cells and thus grid-place relationships are preserved across environments.

II-46. Decoding social information from population codes in the prefrontal cortex of behaving mice

Tal Tamir TAL.TAMIR@WEIZMANN.AC.IL
Dana Rubi Levy DANA-RUBI.LEVY@WEIZMANN.AC.IL
Maya Kaufman MAYA.KAUFMAN@WEIZMANN.AC.IL
Aharon Weissbrod AHARON.WEISSBROD@WEIZMANN.AC.IL
Elad Schneidman ELAD.SCHNEIDMAN@WEIZMANN.AC.IL
Ofer Yizhar OFER.YIZHAR@WEIZMANN.AC.IL

Weizmann Institute of Science

The medial prefrontal cortex (mPFC) plays a prominent role in regulating diverse functions in the mammalian brain. Functional and morphological deficits in this region are associated with several neuropsychiatric disorders, specifically ones involving social deficits. Yet, little is known about how the mPFC encodes social information, and how changes in these representations relate to impaired social behavior. We used a novel behavioral setup to simultaneously record the activity of 10-30 single units from the mPFC of behaving mice, presented with precisely-timed social and nonsocial olfactory stimuli. To explore these representations in social impairment, we studied both wild type (WT) mice and Cntnap2-/- mice, an established genetic model of autism. We compared the encoding of different stimuli using pairwise maximum entropy models of the population responses to each stimulus. We found that male and female odors evoked similar encoding distributions, which were markedly different from those evoked by various non-social cues, regardless of odor valence. These models also allowed us to accurately de-

code single trial data and determine whether a given stimulus was social or not. Interestingly, the spatio-temporal responses to social stimuli and non-social ones became more distinct in consecutive recording sessions, reflecting experience-based refinement of social and non-social representations. In the Cntnap2-/- mice, population responses to social and non-social stimuli were less distinct than those of WT littermates, and showed only small changes with time. We further found that the variance of ongoing activity in Cntnap2-/- mice was significantly higher than in WT mice and that the level of this 'noise' presented strong negative correlation with the distinguishability of responses to social and non-social stimuli. Our results identify unique representations of social information in the mouse mPFC and their refinement with experience. Moreover, they suggest autism-associated differences in neuronal social representations, which are correlated with elevated neural population noise.

II-47. Neuromorphic computing algorithms and hardware for low-power neural decoding and brain-machine interfaces

David Clark[1,2] DAVIDCLARK@BERKELEY.EDU
Jesse Livezey[1] JLIVEZEY@LBL.GOV
Edward Chang[3] EDWARD.CHANG@UCSF.EDU
Kristofer E. Bouchard[1] KEBOUCHARD@LBL.GOV
David Donofrio[1] DDONOFRIO@LBL.GOV
Paolo Calafiura[1] PCALAFIURA@LBL.GOV
Jose Carmena[2] JCARMENA@BERKELEY.EDU

[1]Lawrence Berkeley National Laboratory
[2]University of California, Berkeley
[3]University of California, San Francisco

Decoding neural population activity is a ubiquitous task in systems neuroscience. A common application is brain-machine interfaces (BMIs), which restore lost function by mapping neural recordings to control signals in real-time. However, challenges remain in deploying such systems. For example, decoding algorithms are often computationally demanding and thus dissipate significant energy operating continuously. We address this problem by mapping the Kalman filter (KF) onto a low-power neuromorphic architecture: IBM's TrueNorth. The resulting decoder consumes only ~10-100 mW of power running in real-time and reproduces the KF with high accuracy. In contrast to previous research, our decoder is run on neuromorphic hardware. We demonstrate the utility of our decoder in two tasks: decoding monkey reaches from spiking data and decoding human vocal pitch from ECoG. We show that a characteristic tradeoff between numerical precision and latency in spiking neuromorphic architectures is mitigated by including more on-chip neurons, at the cost of increased power consumption. To reduce power consumption, we perform feature selection using the scalable CUR matrix decomposition, which addresses a variant of the column subset selection problem. In contrast to principal component analysis (PCA), CUR selects specific input features, not linear combinations thereof. With fewer input features, fewer on-chip neurons are required for the decoder, reducing power consumption. Interestingly, we find that decoding accuracy is only slightly worse using a small number of electrodes/neurons as compared to a similar number of PCs. This suggests that the information necessary for decoding is not broadly distributed across all electrodes/neurons, but rather is contained within a small subset. Overall, our findings demonstrate that a combination of neuromorphic technology and subset selection-based dimensionality reduction form a promising and practical path forward for low-power, portable neural decoding.

II-48. Cerebellar climbing fibers can signal learned sensory prediction errors

Court Hull HULL@NEURO.DUKE.EDU
Jake Heffley JAKE.HEFFLEY@DUKE.EDU
Ziye Xu ZIYE.XU@DUKE.EDU
Duke University

Classical models of cerebellar learning posit that climbing fibers operate according to a supervised learning rule to instruct changes in motor output by signaling the occurrence of movement errors. However, recent evidence has challenged this view by suggesting that climbing fiber-driven complex spiking can exhibit characteristics consistent with a reinforcement learning rule after the acquisition of learning in an aversive conditioning paradigm. To test whether sensory prediction error provides a generalizable model to explain the behavior of climbing fibers (CFs) in other behavioral paradigms and across other cerebellar regions, we have adapted a different classical conditioning paradigm that utilizes an appetitive stimulus to evoke CF inputs to Purkinje cells of superficial lobule simplex. Specifically, we have measured CF activity in a behavioral paradigm where head-fixed mice learn to associate a visual cue with an upcoming reward. In this regime, we have used a combination of mesoscale, single-photon imaging and resonant scanning two-photon imaging of virally expressed GCaMP6f in Purkinje cells, and measured CF activity both at the population level and within individual Purkinje cell dendrites across multiple learning sessions. This approach also allows us to measure CF input to the same neurons both before and after learning. Data from both single and multiphoton imaging sessions suggest that climbing fibers can signal learned prediction errors, suggesting a revised model of cerebellar learning for some behaviors where a reinforcement learning rule is appropriate.

II-49. Mapping perceptual decisions to cortical regions

Peter Zatka-Haas P.ZATKA@UCL.AC.UK
Nicholas Steinmetz NICK.STEINMETZ@GMAIL.COM
Matteo Carandini M.CARANDINI@UCL.AC.UK
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
University College London

Perceptual decisions involve a complex interaction of several brain areas. The neocortex is thought to play a major role in this process, but it is unclear which cortical areas are causally involved, and what their individual roles are.

To identify the contributions of specific cortical areas in a visual decision task, we trained Ai32xPV-cre mice in a two-alternative unforced visual discrimination task (Burgess et al, Cell Reports 2017). Mice were rewarded with water for turning a wheel to indicate which of two stimuli had higher contrast, or for holding the wheel still if no stimuli were present. To investigate the causal role of cortical areas in this task, we used laser-scanning optogenetics to inactivate each of 52 cortical sites in randomly interleaved trials.

Inactivation of visual and secondary motor areas (VIS and M2) decreased choices towards stimuli contralateral to the inactivated side, while increasing ipsilateral choices. Inactivation of VIS was maximally effective in a critical time-window, -40ms to +75ms relative to stimulus onset. The critical window for M2 peaked ~40ms later. These results are consistent with M2 having a causal role after VIS areas.

To ascertain what type of computation VIS and M2 might be performing in the behavioral task, we measured their activity with widefield calcium imaging and electrophysiological recordings with Neuropixels probes. Activity in both areas correlated with the visual stimulus during the corresponding critical window. However, while activity in both regions could be used to predict accurately Go vs. NoGo choices, it was barely informative about Left vs. Right choices. Moreover, the time-window in which choice could best be decoded from activity was later than the critical inactivation window.

These results indicate that visual and secondary motor cortical areas are necessary for relaying, in succession, sensory information to downstream targets where the decision is made.

II-50. Mental model complexity, information geometry and the resolution of observations

Kamesh Krishnamurthy KAMESH.KRISHNAMURTHY.W@GMAIL.COM
Alexandre Filipowicz ALSFILIP@GMAIL.COM
Vijay Balasubramanian VIJAY@PHYSICS.UPENN.EDU
Joshua Gold JIGOLD@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

The ability to make effective predictions depends on learned, internal models of the environment. A key characteristic of such models is their complexity, which accounts for the balance between the flexibility and robustness of their predictions. Here we synthesize a novel set of approaches for defining and measuring the complexity of internal models in a principled way. We start with notions of complexity from classical model-selection techniques using Bayesian approaches with a non-informative Jeffreys prior. However, these methods require defining a functional form of the model to be assessed, which in the case of internal models is likely unknown. An alternative approach, called predictive information (PI), is an empirical measure of the complexity of an unspecified model that is based on the growth of the subextensive component of the entropy in data. We show that for data generated with the uninformative prior over the parameters from a given model, classical model-selection and PI approaches yield closely aligned metrics of complexity. However, both these methods assume that there are enough data to probe all the degrees of freedom of the model. As shown recently, this assumption is often violated in typical scenarios in which measurements are not tuned to the details of the model generating the data. Under these conditions, many parameter combinations may have little effect on data measured at a chosen resolution, implying that a simpler model is sufficient. We extend the classical complexity metrics by showing that this dependency of complexity on measurement resolution is also naturally embodied in lower-order terms of Bayesian model-selection methods. We then discuss measuring and interpreting subjective model complexity in predictive tasks in a measurement dependent manner. Together, these ideas provide a unifying and practical framework for measuring and interpreting the complexity of mental models.

II-51. Changes in effective network coupling mediate memory across time in the hippocampus

James Priestley JBP2150@CUMC.COLUMBIA.EDU
Mohsin Ahmed MSA2011@CUMC.COLUMBIA.EDU
Angel Castro AC2955@COLUMBIA.EDU
Fabio Stefanini FS2545@CUMC.COLUMBIA.EDU
Elizabeth Balough EMB2162@CUMC.COLUMBIA.EDU
Erin Lavoie EL2927@COLUMBIA.EDU
Luca Mazzucato LM3356@COLUMBIA.EDU
Stefano Fusi SF2237@COLUMBIA.EDU
Attila Losonczy AL2856@CUMC.COLUMBIA.EDU

Columbia University

Episodic memory requires linking discontiguous events in time, a function dependent on the hippocampus. This temporal association learning is often modeled using trace fear conditioning, where a conditioned stimulus (CS) and aversive, unconditioned stimulus (US) are mnemonically linked across delays of tens of seconds. Trace associations are nontrivial and preclude simple Hebbian mechanisms, due to non-overlapping external sensory information about the cues. To bridge this gap, persistent activity has been proposed to propagate information forward in time, accomplished either through activity which is approximately constant, as in attractor models, or evolving in time, along a trajectory in neural space which reliably connects the CS and US representations.

Here we investigate ensemble activity in hippocampal area CA1 during trace fear learning to resolve the underlying

representation. To do so, we have developed a head-fixed, auditory trace fear memory assay that is dependent on dorsal CA1. We integrate this paradigm with 2-photon calcium imaging to chronically record activity of hundreds of CA1 pyramidal cells over the course of learning. Mice exhibit robust learned fear responses and reliably discriminate between fearful (CS+) and neutral (CS-) stimuli.

We find that neither static nor dynamic persistent activity is congruent with the observed population code, because stimulus and temporal information cannot be decoded from the neural activity. However, CS identity can be reliably decoded from the covariance of neural activities, which defines the effective network couplings. Our analysis suggests CS identity is encoded by combinatorial patterns of cell activation, which are temporally unreliable across trials. These patterns develop with learning, and are consistent across both cue presentation and the 'trace' interval. Thus we propose a new model for trace fear conditioning whereby transient synaptic modification may bias network states to efficiently store cue information, without maintaining an energetically expensive persistent representation throughout the trace interval.

II-52. Local, reinforceable and information-optimal learning in growing networks

Sensen Liu LIUS@ESE.WUSTL.EDU
ShiNung Ching SHINUNG@WUSTL.EDU

Washington University in St. Louis

A central problem in learning is to enable behavior that achieves good outcomes in new environments. Several theoretical frameworks have been proposed for such learning at the neuronal level, i.e., in terms of synaptic plasticity. Here, we bring together notions of reinforcement learning (RL) and information optimization to construct a local, dynamic synaptic plasticity rule that has versatility in achieving different learning regimes. Our model follows from our prior work, where we developed a learning algorithm that optimizes information retention in a recurrent network. Here, we introduce an augmented version of this learning rule by employing a reinforcement mechanism that can modulate the alternations between Hebbian and anti-Hebbian regimes on the basis of a prescribed objective function. In parallel, we consider not only plasticity, but also the issue of network growth. That is, the idea that learning arises not only from the rules of synaptic plasticity, but also from activity-dependent construction and elimination of synapses. Thus, the second contribution of our work is to supplement the learning rule with a strategy that enables both synapse production and pruning to design an information-efficient, and low-cost network. More specifically, we: (i) develop a local, dynamical plasticity rule that incorporates a reinforcement mechanism within an already information-optimal paradigm, in a recurrent, stochastic spiking network, (ii) derive an algorithm to enable online network growth/pruning that leverages the learning dynamics and, (iii) illustrate the learning rule through simple examples.

II-53. How do calcium indicator properties affect spike inference?

Thomas Delaney TD16954@BRISTOL.AC.UK
Cian O'Donnell CIAN.ODONNELL@BRISTOL.AC.UK
Michael Ashby M.C.ASHBY@BRISTOL.AC.UK

University of Bristol

The use of fluorescent calcium indicators, such as GCaMP, to monitor neuronal activity is widespread. But the relationship between action potential firing and the fluorescence signal is poorly understood. For example, it is known that genetically encoded indicators accumulate within neurons over weeks and months. This makes comparison of activity levels at different time points difficult. Furthermore, the effects of the indicator characteristics on this fluorescence signal are unknown. As a result, it remains unknown if spike train inference is always possible using fluorescent calcium indicators.

The aim of this project was to model the fluorescence traces produced by a fluorescent calcium indicator in a neuron soma, given parameters such as binding rate, dissociation rate, and molecular concentration from a specific spike train. The ultimate goal of the model is to allow benchmarking of the various spike inference algorithms that have been developed, and to understand how indicator characteristics affect the quality of spike train inference.

The fluorescence traces produced by the model were calibrated to reproduce the signal-to-noise ratio observed in experimental data. We tested three leading spike inference algorithms (Pnevmatikakis et al. 2016, Jewell et al. 2010, Friedrich et al. 2017) on the real versus modelled calcium imaging traces. We varied the values of model parameters to determine the effect on system dynamics and on spike inference. This framework has two uses, firstly for helping experimentalists optimise their calcium imaging experiment design choices for best spike inference, and also for helping computational researchers optimise and understand their spike inference algorithms.

II-54. A model of neuronal circuit development through activity-dependent plasticity

Yayoi Teramoto YAYOI.TERAMOTO@CNCB.OX.AC.UK
Simon J.B. Butt SIMON.BUTT@DPAG.OX.AC.UK
Tim Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Brain circuitry changes dramatically throughout the first weeks of life, but the rules governing those changes – genetic, activity-driven, or else – remain elusive. One example of drastic restructuring is the transient set of connections between GABAergic interneurons and spiny stellate neurons in the developing barrel cortex of mice that are present exclusively in the first two weeks. While electrophysiological and genetic tools enable us to take snapshots of these early connectivity changes, such approaches cannot easily answer questions about the mechanisms driving such structural changes. Here, we build networks of spiking neurons to study what mechanisms could govern the assembly and dis-assembly of transient circuits, as well as the functional role they could play in establishing the adult brain circuitry. In our models, activity-dependent mechanisms can account for the correct assembly of functional networks. The reversal potential of GABAA – known to change over development – and the amount of input activity to the network play a crucial role in driving connectivity modifications, but there is no single parameter value that alone determines the final connectivity of the simulated network. This poses a challenge in the design of validation experiment for our model. To compare model and experiment, we used the initial and final connectivity of each simulation to map all parameter combinations to one the experimental conditions used by Marques-Smith et al. (2016) (wildtype, genetic knockout and sensory perturbation). Using this mapping we were able to propose a minimal set of experiments to validate (or falsify) the models, and continue the dialogue between modeller and experimentalist. The model highlights features that may direct experimental work. Vice versa, any experimental results are informative about the model's predictive ability and applicable to the model itself.

II-55. Multiscale modeling and decoding of spike-field activity during a naturalistic reach-to-grasp task

Han-Lin Hsieh[1] HANLINHS@USC.EDU
Yan Wong[2] YAN.WONG@MONASH.EDU
Bijan Pesaran[3] BIJAN@NYU.EDU
Maryam Shanechi[1] SHANECHI@USC.EDU

[1]University of Southern California
[2]Monash University

[3]New York University

Motor behavior is encoded across spatiotemporal scales, from small networks of spiking neurons to larger networks measured using local field potentials (LFP) and electrocorticogram (ECoG). Precise modeling and decoding of multiscale activity can both help investigate neural encoding across scales and improve brain-machine interface (BMI) performance and longevity. However, multiscale modeling introduces new computational challenges because of significant differences in statistical profiles and time-scales across scales of activity. In particular, spikes are binary-valued with a millisecond time-scale while LFP/ECoG are continuous-valued with slower time-scales (e.g., milliseconds to tens of milliseconds). Here, we develop a novel multiscale state-space encoding model that characterizes how movements are simultaneously represented across scales. We also derive the corresponding multiscale filter (MSF) for decoding movement. The MSF runs at the fast time-scale of spikes and models their inherent properties as a point process, while simultaneously adding information from fields at a slower time-scale. We apply our multiscale models and decoders to estimate 7 arm joint angles in a non-human primate (NHP) from motor cortical spike-LFP activity during a 3D naturalistic reach-to-grasp task. We find that multiscale decoding improves estimation accuracy compared with single-scale decoding of spikes with a point process filter (PPF) as well as single-scale decoding of LFP with a Kalman filter (KF). This improvement is observed regardless of the number of spike and LFP channels included in the decoder. Thus MSF can combine information at multiple time-scales. Moreover, the improvement in MSF is greatest in the low-information regime where fewer LFP and spike channels are available, thus suggesting that multiscale decoding can improve BMI longevity as recording quality degrades over time. Taken together, the multiscale modeling and decoding framework could improve emerging neurotechnologies and has potential as a scientific tool to understand how behavioral processes are represented across large-scale networks and across multiple time-scales.

II-56. The structure of whole fly brain spontaneous activity mirrors the structure of fly behavior

Stephane Deny SDENY@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Kevin Mann MANNK2@STANFORD.EDU
Thomas Clandinin TRCLANDININ@GMAIL.COM

Stanford University

How the structure of internal spontaneous activity during resting state reflects the structure of behavior is a fundamental question in neuroscience, which we study in the fly for the first time by leveraging two recent advances: (1) the ability to measure dynamic whole-brain fly spontaneous activity (Mann et al. 2017), and (2) the derivation of static brain-behavior relations associating each fly behavior with a brain map, defined as the set of brain regions eliciting the behavior upon stimulation (Robie et al. 2017). By combining these advances through basic statistical analyses, we provide a quantitative framework to study relations between spontaneous activity and behavior in a new model system amenable to exquisite physiological accessibility and extensive behavioral characterization. Although a long-standing hypothesis for the selection of specific behaviors involves mutual inhibition between neural circuits that elicit competing behaviors, we nevertheless found that pairs of voxels, recorded via calcium imaging of fly whole-brain spontaneous activity, were relatively dominated by positive correlations. In contrast, projections of spontaneous activity patterns onto a functional basis set of brain maps associated with different behaviors, exhibited a much richer mixture of positive and negative correlations. These results suggest the potential for rich patterns of mutual inhibition between behaviors, despite the lack of prevalent anatomically organized patterns of mutual inhibition. Moreover, through dimensionality reduction, we embed fly behaviors into a low dimensional space such that two behaviors are close to (distant from) each other if spontaneous activity projected onto the associated brain maps are similar (different). Remarkably we find that nearby behaviors (i.e. wing extension and attempted copulation) also tend to co-occur during fly behavior, while distant behaviors (i.e. walking and resting) tend to be mutually exclusive in fly behavior. Thus, we find the structure of fly whole-brain spontaneous activity directly reflects the co-occurrence structure of fly behavior.

II-57. Hierarchical inference and learning using distributed representations of uncertainty

Eszter Vertes[1,2] VERTES.ESZTER@GMAIL.COM
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]Gatsby Computational Neuroscience Unit

Sensory systems in the brain must make sense of noisy, often ambiguous incoming stimuli. Forming a percept based on receptor activations in the periphery is challenging—the underlying computation is ill-posed, and thus must be tackled probabilistically. The statistical accuracy of inference observed in behavioral experiments points to the capacity of neural circuits to learn about the generative model underlying natural statistics. A number of schemes have been suggested for how populations of neurons may code for, and compute with, uncertainty (e.g. Hoyer & Hyvarinen 2003, Ma et al. 2006, Zemel et al. 1998), but there has been very little work on how such representations could be acquired by neural systems. We propose a new approach, the Distributed Distributional Code (DDC) Helmholtz machine, to learn a causal generative model of sensory stimuli and simultaneously learn to accurately infer the corresponding explanatory (or latent) variables. A key feature of our model is that neural activity encodes uncertainty about the latent causes implicitly. The inferred posterior distribution is represented as a set of expectations distributed across a population of neurons (Zemel et al. 1998; Sahani & Dayan 2003), i.e. in a “distributed distributional code” (DDC). To learn both the generative and the recognition model, that performs inference over latent causes given the incoming sensory observations, we use a wake-sleep-like algorithm inspired by the Helmholtz machine (Dayan et al. 1995). Even for hierarchal models, the learning rules remain local, making our approach biologically appealing. Furthermore, the posterior representation does not impose independence or a rigid parametric structure, thus it is able to capture the statistical dependencies of the latent causes faithfully. We show that the DDC Helmholtz machine accurately learns generative models of olfactory and visual stimuli with hierarchically organized latent variables, where standard approaches relying on factorized posterior approximation fail.

II-58. Modulation and propagation of information in visual pathway

Chengcheng Huang[1] CHENGCHENGHUANG11@GMAIL.COM
Marlene Cohen[1] MARLENERCOHEN@GMAIL.COM
Alexandre Pouget[2] ALEX.POUGET@GMAIL.COM
Brent Doiron[1] BDOIRON@PITT.EDU

[1]University of Pittsburgh
[2]University of Geneva

How neuronal variability impacts neural codes is a central question in systems neuroscience, often with complex and model dependent answers (Kohn et al. 2016). Most population models are parametric (Abbott & Dayan 1999, Ecker et al. 2016), with tacitly assumed structure of neuronal tuning and population variability. While these models provide key insights, they cannot inform how the physiology and circuit wiring of cortical networks impact information flow. Attentional modulation is an often used tool to probe the neural correlates of cortical processing, since attention is well known to improve cognitive performance in discrimination tasks, as well as attenuate population-wide response variability (Cohen & Maunsell 2009, Ruff & Cohen 2014). Attention offers key constraints that have allowed our group to propose and analyze a circuit-based cortical model which recapitulates the attentional modulation of both trial averaged and trial variable response (Huang et al. 2017). In this study, we use this model to investigate how the feedforward and recurrent structure of cortical circuits, and their attentional modulation, shape information flow within the visual system. When the stimulus has trial-to-trial fluctuations that are external to the network, the Fisher information grows sub-linearly with the number of neurons, showing signs of saturation, consistent with past models (Moreno-Bote et al. 2014, Kanitscheider et al. 2015). We show that a network with narrow feedforward and recurrent projections can transmit almost all of the Fisher information across

layers. Moreover, the attentional manipulation in our model increases transmitted Fisher information from a finite network while decreasing pairwise correlations to an extent comparable to that observed in experiments.

II-59. An inhibitory wave produces an omega turn in a model of C. elegans

Charles Fieseler[1] KONDA@UW.EDU
J Nathan Kutz[1] KUTZ@UW.EDU
James Kunert-Graf[2] JKUNERT@PNRI.ORG

[1]University of Washington
[2]Pacific Northwest Research Institute

In a model of C elegans designed to produce forward motion, a posteriorly traveling wave of inhibition on the proprioceptive signal robustly produces an omega turn. This work describes the first biomechanically grounded testable hypothesis for behaviors in C elegans beyond simple motion, and is of particular relevance because the forward motion limit cycle is transiently modulated and no new musculature is needed. Our model includes three layers: (i) muscle structure and activation, (ii) key neural activation circuitry, and (iii) weighted and time-dynamic propioception, which is the main oscillatory mechanism. In combination, we show that these model components can reproduce the complex waveforms exhibited in C elegans locomotive behaviors, crucially including the fast transient "omega turn" that is vital in the escape sequence. This inhibitory wave of modulation is qualitatively similar to extra-synaptic neuromodulator diffusion, and can robustly and realistically produce this behavior. A powerful mathematical intuition about the connection between an organism's behavior and its neural network is that of the fixed point and its cousin, the limit cycle. These tools are tailor-made for probing the structure of relatively stable behaviors, but are less useful for transient or short timescale actions. This work reveals a plausible and testable mechanism for behavior that is produced outside both the connectomic and fixed point paradigms.

II-60. Probabilistic integration and learning in grid cells produces experience dependent distortion and rotational alignment

Talfan Evans TALFAN.EVANS.13@UCL.AC.UK
Daniel Bush D.BUSH@UCL.AC.UK
Neil Burgess N.BURGESS@UCL.AC.UK

University College London

Grid cells (GC; Hafting et al., 2005) found in the medial entorhinal cortex (mEC) have been suggested to play a role in path integration (PI), the process of integrating self-motion cues in order to maintain an estimate of location relative to a previous estimate. However, GC firing patterns are stable over time, suggesting that they also receive stabilising input from their environment. Existing GC models rely on strong stimuli from pre-learned associations with sensory inputs to 'reset' the accumulated error in the grid pattern (Burak and Fiete, 2009; Hardcastle et al., 2015). We argue that this mechanism of 'hard-resetting' leads to localisation errors when sensory and self-motion cues conflict, and does not support simultaneously navigating while learning these sensory associations (the neural 'map') in large novel environments. Instead, we propose a biologically plausible neural architecture that weighs self-motion and sensory estimates according to their certainty and learns the correct mapping of place cell (PC) mediated sensory inputs in novel environments using standard Hebbian-type learning rules under certain conditions. Secondly, we propose that competition between location estimates from pre-learned sensory cues and self-motion may underlie the elliptical shearing and accompanying 7–9^{o} alignment of the GC pattern's axis to the walls of the environment observed experimentally (Krupic et al., 2015; Stensola et al., 2015) in square environments. Our model responds in a scale-dependent 'all or nothing' manner in response to environmental contractions or expansions (Barry et al., 2007; Stensola et al., 2012). Lastly, we demonstrate that our model performs GC driven PC remapping, allowing the hippocampal formation to distinguish between perceptually identical

'local' environments according to their embedding within a 'global' map (Carpenter et al., 2015).

II-61. Model sonification reveals advantages of task-optimized sensory models

Jenelle Feather JFEATHER@MIT.EDU
Josh H McDermott JHM@MIT.EDU
Massachusetts Institute of Technology

Models of neural responses and perceptual judgments have traditionally been built using engineering principles and experimental observations, but modern-day machine learning allows models to be learned from data. We sought to compare hand-engineered and learned models in the auditory domain, and developed a general-purpose optimization method to synthesize sounds from models for this purpose. We synthesized stimuli that produce the same values in a model's representation as a natural stimulus. Such stimuli should evoke the same percept if the model replicates the representations underlying perception, and the same neural response at processing stages accurately described by the model. We utilized automated gradient-based optimization tools such that the same synthesis procedure could be applied to any differentiable model. Instead of conventionally optimizing and then inverting an intermediate spectrogram-like representation, optimization occurred directly on the sound waveform, minimizing artifacts. We compared a common hand-engineered model of spectrotemporal filters to learned filters from the first layer of a task-optimized convolutional neural network. We evaluated model-matched synthetic sounds from both models using realism ratings and an objective recognition task. Sounds generated with the learned filters were more recognizable and realistic than those from the hand-engineered filter bank. To explore the origins of this difference, we additionally constrained sounds to match statistics of the cochlear representation preceding the spectrotemporal filters. The inclusion of cochlear statistics caused sounds from the hand-engineered model to improve to the level of those from the learned filter bank. By contrast, including cochlear statistics did not improve the quality of sounds from the learned filters. The learned filters evidently retain task-relevant information from earlier processing stages that is discarded by conventional filters. The results suggest that this information is perceptually important, illustrating that better models can be obtained by task-optimizing sensory representations. The methodology should be broadly applicable to other models.

II-62. Organization of neural population code in mouse visual system

Kathleen Esfahany[1] KATHLEEN.ESFAHANY@GMAIL.COM
Isabel Siergiej[2] IS278@CORNELL.EDU
Yuan Zhao[1] YUAN.ZHAO@STONYBROOK.EDU
Il Memming Park[1] MEMMING.PARK@STONYBROOK.EDU
[1]Stony Brook University
[2]Cornell University

The mammalian visual system consists of several anatomically distinct areas, layers, and cell types. To understand the role of these subpopulations in visual information processing, we analyzed neural signals recorded from excitatory neurons from various anatomical and functional structures. For each of 186 mice, one of six genetically tagged cell-types and one of six visual areas were targeted while the mouse was passively viewing various visual stimuli. We trained linear classifiers to decode one of six visual stimulus categories with distinct spatiotemporal structures from the population neural activity. We found that neurons in both the primary visual cortex and secondary visual areas show varying degrees of stimulus-specific decodability, and neurons in superficial layers tend to be more informative about the stimulus categories. Additional decoding analyses of directional motion were consistent with these findings. We observed synergy and redundancy in the population code of direction in several visual areas suggesting area-specific organization of information representation across neurons. These

differences in decoding capacities shed light on the specialized organization of neural information processing across anatomically distinct subpopulations, and further establish the mouse as a model for understanding visual perception.

II-63. Dissociable cortical networks encode cue sequences and movement sequences

Patrick Beukema[1] PLB23@PITT.EDU
Timothy Verstynen[2] TIMOTHYV@ANDREW.CMU.EDU

[1]University of Pittsburgh
[2]Carnegie Mellon University

The efficient execution of serially ordered actions is crucial for many everyday tasks. Most paradigms for probing sequential skill learning involve the presentation of a visual stimulus that cues a particular movement. This makes disambiguating systems that learn sequences of goals from systems that learn sequences of actions difficult. Thus, it remains an open question as to whether it is possible to selectively distinguish networks that encode sequences of cues from those that encode sequences of actions (or simply encode individual cues or actions in a non-sequential manner). Here we explicitly dissociated the visual stimulus from the cued movement in a finger sequence task using a novel remapping paradigm that breaks the visual cue and response colinnearity. This enabled us to examine each of two resulting networks, visual cue and movement sequence, in isolation. Using an unbiased decoder based on the cross-validated Mahalanobis distance measure, we observed two distinct networks that were both entirely lateralized to the contralateral hemisphere. As expected, the visual cue sequences were primarily encoded within the visual cortex, especially the ventral visual stream, as well as an anterior aspect of the superior parietal lobule. We also observed encoding of the visual sequences in higher level motor planning areas in premotor dorsal cortex. In contrast, the motoric sequences were not decodable anywhere within visual cortex. Instead, we observed encoding of the motoric sequences in primary motor and primary somatosensory cortex, centered on the hand knob. We also observed motoric encoding within the most anterior aspect of the intraparietal sulcus, supplementary motor area, the inferior parietal lobule, as well as the superior aspect of the medial temporal lobe. Our results highlight the fact that different features of sequences are represented in distinct cortical areas and suggest that learning can proceed along two orthogonal channels.

II-64. A novel framework for dynamic modeling of brain-network response to electrical stimulation

Yuxiao Yang[1] YUXIAOYA@USC.EDU
Omid Sani[1] GHASEMSA@USC.EDU
Kristin Sellers[2] KRISTIN.SELLERS@UCSF.EDU
Edward Chang[2] EDWARD.CHANG@UCSF.EDU
Maryam Shanechi[1] SHANECHI@USC.EDU

[1]University of Southern California
[2]University of California, San Francisco

Model-based closed-loop control of neural activity is critical for establishing functional connectivity in neural circuits and developing precisely-tailored electrical stimulation therapies for neuropsychiatric disorders. Model-based control requires the characterization of brain-network dynamics (output) in response to stimulation (input). However, a principled brain-network input-output (IO) model identification framework, which is also amenable to closed-loop control design, is currently lacking. Here, we develop a control-theoretic data-driven brain-network IO modeling framework and validate it in human epilepsy patients implanted with large-scale electrocorticography (ECoG) electrodes. We first construct a multivariate input-driven linear state-space model (LSSM) that describes the brain-

network IO dynamics and is amenable to controller design. We then develop a data-driven system-identification method to estimate the model parameters. To collect appropriate IO data for model estimation, it is critical to devise an informative input waveform to stimulate the brain. We design a novel input waveform—a pulse train modulated by binary noise (BN) parameters such as pulse frequency and amplitude—that we show is optimal for system-identification and conforms to clinical safety requirements. To validate our framework, we applied the BN electrical stimulation in epilepsy patients and recorded the ECoG response. We found that the estimated model could accurately predict both the dynamic and the steady-state ECoG response to stimulation. Our results show the feasibility of identifying dynamic predictive models of neural response to electrical stimulation and have significant implications for future model-based closed-loop therapies of various neurological disorders.

II-65. Computational and neural mechanisms of goal-directed planning and problem-solving

Joshua Brown JWMBROWN@INDIANA.EDU
Noah Zarr NNZARR@INDIANA.EDU

Indiana University

A significant limitation of both model-based and model free RL is that typically there is only a single ultimate goal. Q-values are thus learned in order to maximize a single reward value. In contrast, real organisms will find differing reward values associated with different goals at different times and circumstances. This implies that goals will change over time, and re-learning Q-values with each goal change would be highly inefficient. Instead, a more flexible mechanism will dynamically assign values to various goals and then plan accordingly. We developed the Goal-Oriented Learning and Selection of Action (GOLSA) model as a new approach to overcome the limitations of less flexible Q-values, while maintaining fidelity to known biological mechanisms. The model learns the structure of state transitions, then plans actions to arbitrary goals via a novel hill-climbing algorithm inspired by Dijkstra's algorithm, and similar to that used in GPS navigation devices. The model provides a domain-general solution to the problem of solving problems and performs well. Moreover, we use model-based fMRI with representational similarity analysis (RSA) to show that in addition to solving complex planning problems, the GOLSA model provides a novel computational account of network interactions of a number of brain regions involved in flexible action planning. In particular, both current and final goal state representations in the model significantly match patterns in visual sensory rather than prefrontal cortical regions. The currently selected goal in the model matches ventral striatal activity patterns. Goal-driven state planning in the model matches activity patterns in premotor cortex. Surprisingly, no significant decodability for goal states was detected in the hippocampal region despite significant results elsewhere. Together our results suggest a novel computational account of how the brain plans actions to solve problems, and how a number of brain regions contribute interacting computational roles to such behavior.

II-66. OpenCortex: an integrated platform for cortex-wide imaging and deep-brain modulation

Ben Huang BSH2002@MED.CORNELL.EDU
Conor Liston COL2004@MED.CORNELL.EDU

Cornell University

Cognitive functions, such as memory encoding/retrieval, involve distributed networks spanning the cortex and concurrent interplay with subcortical structures such as hippocampus. Past work largely focused on separate regions in isolation, while mechanisms of large-scale cortical-subcortical interactions remain unclear. Current barriers in studying cortex-wide interactions with deep-brain regions include: 1) limited chronic whole-cortex recording access, 2) limited cortex-wide recording methods at cellular resolution, and 3) limited techniques for concurrent cortical recording and subcortical manipulation. To overcome the first challenge, we have developed a cortex-wide cranial window in mice that provides optical access to entire dorsal cortical surface, by replacing the skull with 10-mm-diameter glass coverslip, fused with micro-prisms for accessing medial cortical regions. For the second challenge, we took a multiscale imaging approach (combining one-photon widefield imaging and 2-photon microscopy at different magnifications) in the same brain, to analyze network dynamics at both global and single-neuron levels. For the third challenge, we combined the cortex-wide window with chronically-implanted multi-functional fiber-probes (Park et al., 2017), which consist of a central light-guide surrounded by microfluidic channels and recording electrodes. To accommodate these probes while preserving cortex-wide window/imaging access, we developed a novel lateral entry-port to insert the probes from lateral side of the skull, combined with a micro-drive for in vivo repositioning. This integrated OpenCortex platform enables concurrent optogenetic stimulation, drug infusions, and electrical recordings of subcortical regions while cortex-wide activity is recorded. Here we present the development of this new platform and the data demonstrating its use in studying cortical-hippocampal interactions in behaving mice. In sum, OpenCortex enables the cortex-wide recording of single-cell activity with simultaneous monitoring and modulation of subcortical inputs. Long-term stability of the cortex-wide windows and fiber-probes allows repeated recordings from same animal-subject over time, thus enabling rapid theory-experiment iterations to test causal models.

II-67. Extending models of latent dynamics in area LIP during perceptual decision-making

David M. Zoltowski[1] ZOLTOWSKI@PRINCETON.EDU
Kenneth Latimer[2] LATIMER1@UW.EDU
Alex Huk[3] HUK@UTEXAS.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]University of Washington
[3]The University of Texas at Austin

Trial-averaged firing rates in macaque area LIP tend to ramp upwards during sensory evidence accumulation. Recently, (Latimer et al., 2015) reported the surprising finding that the single-trial responses of the majority of these neurons were better described by a discrete state (stepping) model than a continuous diffusion-to-bound (ramping) model. We extended these analyses in several novel and important directions. First, we introduced nonlinear ramping models in which the firing rates were saturating or accelerating functions of a diffusion-to-bound process, and we included a minimum positive firing rate parameter (baseline rate) in the ramping models. Second, we added autoregressive spike history terms to account for non-Poisson firing statistics. We compared the models using the Watanabe-Akaike information criterion (WAIC), a Bayesian model comparison metric. We used Markov chain Monte Carlo methods to fit the models to the responses of 40 LIP neurons recorded during a random dot motion task (Meister et al., 2013). We found that among the ramping models, the WAIC favored a model with a square root nonlinearity, spike history terms, and a non-zero baseline firing rate. An approximately equal number of cells were favored by this model and the stepping model with spike history terms. However, we found that ramping models fit with a baseline rate can exhibit single-trial dynamics that blur the distinction between steps and ramps. Finally, we reproduced our results on responses of LIP neurons during a reaction-time task (Roitman & Shadlen, 2002) and during a motion pulse task (Yates et al., 2017).

II-68. The importance of suppression in mid-level auditory processing

Joel Kaardal[1] JKAARDAL@SALK.EDU
Frederic Theunissen[2] THEUNISSEN@BERKELEY.EDU
Tatyana Sharpee[1] SHARPEE@SALK.EDU

[1]Salk Institute for Biological Studies
[2]University of California, Berkeley

Object detection is possible in sensory systems despite the manipulation and transformation of objects in the environment. This is achieved through hierarchical transformations of sensory information as it passes through the sensory ciruitry. For instance, the visual system recognizes objects by exploiting invariances where individual neurons have been found to compute logical OR functions of invariant inputs. However, it is not clear whether the same behavior is observed in the auditory system. We show, in contrast to vision, that several auditory regions of zebra finches are better modeled by logical AND instead of logical OR functions. This is done by introducing a novel dimensionality reduction method called low-rank maximum noise entropy (MNE) and fitting Boolean functions to predict the neural responses of individual neurons. We further show that suppressive inputs dominate across each auditory region.

II-69. A simplified model of a pyramidal neuron as a Canonical Correlation Analyzer (CCA)

Dmitri Chklovskii[1] MITYA@FLATIRONINSTITUTE.ORG
Cengiz Pehlevan[1] CPEHLEVAN@FLATIRONINSTITUTE.ORG
Anirvan Sengupata[2] ANIRVANS.PHYSICS@GMAIL.COM

[1]Flatiron Institute
[2]Rutgers University

Simplified normative models of neurons are appealing because of their interpretability. In particular, Oja's model of a neuron as a Principal Component Analyzer (PCA) computing the top principal component of the upstream activity vectors demonstrated that an important computational task can be implemented using biologically plausible Hebbian learning rules (Oja, 1982). However, PCA treats all inputs to a neuron equally, despite ample evidence that cortical pyramidal neurons treat feedforward inputs to basal dendrites and feedback inputs to the apical tuft distinctly: The two types of inputs are integrated separately and then combined together to generate the output (Larkum, 2013).

Here, we make a step towards a normative model of a pyramidal neuron, including activity dynamics and synaptic plasticity, as an implementation of an online CCA algorithm. Given two related datasets, CCA finds the subspace which maximizes a correlation between their projections onto that subspace. In the model, the two datasets are streamed to a pyramidal neuron as activity vectors of the upstream neurons in feedforward and feedback pathways. At each time step, the two streamed activity vectors are projected onto the common subspace by multiplying by the synaptic weights and summing in basal and apical dendrites. Then, the pyramidal neuron adds these projections and outputs the sum as its firing rate. In addition, synaptic weights are updated according to biologically plausible Hebbian learning rules. Finally, we propose that an extension of CCA to greater than two sources may model processing in dendritic branches.

II-70. Decoding mood state from multisite ECoG activity in human subjects

Omid Sani[1] GHASEMSA@USC.EDU
Yuxiao Yang[1] YUXIAOYA@USC.EDU
Morgan Lee[2] MORGAN.LEE@UCSF.EDU
Heather Dawes[2] HEATHER.DAWES@UCSF.EDU
Edward Chang[2] EDWARD.CHANG@UCSF.EDU
Maryam Shanechi[1] SHANECHI@USC.EDU

[1]University of Southern California
[2]University of California, San Francisco

Decoding mood state from neural activity is key in developing precisely-tailored treatments for mood disorders, but has not yet been achieved due to several challenges. Mood is likely represented across distributed brain sites but the precise spatiotemporal characteristics of mood-predictive neural dynamics is not well-understood. Further, assessing mood state over time is challenging. Finally, decoding mood state would require novel modeling techniques that can incorporate activity across distributed brain sites and deal with the sparsity in available mood measurements caused by the difficulty of mood assessment. Here, we resolve these challenges and demonstrate that mood state variations over time in individual subjects can be decoded from multisite intracranial recordings. We continuously recorded multisite electrocorticogram (ECoG) from six epilepsy subjects across multiple days, and asked them to self-report their mood states intermittently at discrete time-points. We developed a modeling approach to build subject-specific dynamical neural encoding models of mood state and the corresponding decoders. Our modeling approach identified mood-predictive networks across large-scale recording sites to decode mood state. Despite inter-subject variability in psychiatric conditions and mood state ranges, the decoders significantly predicted mood state variations in every subject across days of recordings. We further studied the spectro-spatial distribution of networks that were selected for decoding. We found that the decoders consistently recruited networks within the limbic regions. Moreover, spectral power features in these networks were tuned to mood variations over time, and all frequency bands contributed to mood state prediction. These results shed light on mood-predictive neural activity and demonstrate the feasibility of real-time mood state decoding in individuals, which may facilitate future precisely-tailored treatment of depression and anxiety.

II-71. Dissecting stability and gain modulation in interneuron circuits

Hannah Bos[1] HTB11@PITT.EDU
Anne-Marie Oswald[2] AMOSWALD@PITT.EDU
Brent Doiron[1] BDOIRON@PITT.EDU

[1]University of Pittsburgh
[2]University College London

Inhibition is involved in two opposing mechanisms: controlling the responsiveness (gain) of excitatory (E) neurons and maintaining network stability. Interneurons subdivide into vasoactive intestinal- peptide (VIP), somatostatin (SOM) and parvalbumin (PV) expressing neuron classes. However, it is not clear how gain modulation and stability mechanisms are simultaneously performed in one circuit. Inhibition can suppress the activity of E neurons by direct projections or increase their activity by inhibiting an intermediate interneuron (disinhibition). For the latter, two main pathways have been suggested: SOM neurons inhibit PV neurons which disinhibit E neurons (Xu et al., 2013) and VIP neurons inhibit SOM neurons which disinhibit E neurons (Fu et al., 2014). In this study, we ask how these disinhibitory pathways perform in a recurrently connected circuit with respect to increasing gain while keeping noise correlations low (which are a reflection of network stability). We investigate potential roles of the interneurons by applying theoretical tools developed for balanced and finite-size network dynamics, as well as simulations of spiking neurons. We find that the SOM->PV->E pathway initially shows a gain maximum for moderate SOM rates. Further modulation induces either pathologically large fluctuations or the transition to a silent state. In contrast, disinhibition via VIP and SOM neurons exhibits smooth gain modulation accompanied

with an initial rise in noise correlations which eventually saturates to a lower value. Stability of this pathway is improved by SOM neurons projecting to both, E and PV neurons. SOM neurons tend to suppress rather than stabilize the activity of their targets and recurrent connections between E and SOM neurons suppresses noise correlations but reduce gain. Thus, we predict that gain modulation and stability are most effective when carried out by separate interneurons subtypes.

II-72. Modulation of visual responses by navigational signals during active behavior

Efthymia Diamanti EFTHYMIA.DIAMANTI.11@UCL.AC.UK
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
A Saleem AMAN.SALEEM@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

A fundamental role of vision is to guide navigation, and indeed neurons in cortical areas such as the posterior parietal cortex are involved in the joint processing of visual and navigational signals. But how early in the visual system are navigational signals found? Are such responses present in one or more areas in visual cortex, and how are they different from responses during passive viewing? We used 2-photon calcium imaging to record neural activity across primary visual cortex (V1) and 6 higher visual areas, while head-restrained mice ran along a corridor in virtual reality (VR). The corridor contained two landmarks (vertical grating or plaid) repeated after 40 cm, creating 2 visually identical sections. Imaging sessions involved 3 conditions: (1) closed-loop, where the speed of the virtual corridor matched the animal's run speed; (2) open-loop, where previous closed-loop visual scenes were played back to the animal regardless of its running speed; (3) presentation of vertical drifting gratings at varying temporal and spatial frequencies. In closed-loop, neurons as early as in V1 did not respond similarly to identical landmarks. Instead they fired more strongly at a single virtual position. A similar trend was observed in all areas. Responsiveness at a specific position could not be explained by a complex-cell model of visual responses, or by running speed. In open-loop, fewer cells were responsive in all areas, and the reliability of their responses was markedly reduced. Finally, cells tuned to passively-presented vertical drifting gratings did not necessarily respond to the vertical gratings in VR. Conversely, cells with reliable responses in VR became unreliable during grating presentation. We conclude that during active navigation responses across visual cortex are modulated by spatial context. These responses are overall stronger than during VR replay and cannot be attributed solely to preference for similar grating stimuli.

II-73. Combining deep-learning RL with fMRI to probe the encoding of state-space representations in the human brain

Logan Cross LCROSS@CALTECH.EDU
Jeff Cockburn JEFF.COCKBURN@GMAIL.COM
Yisong Yue YYUE@CALTECH.EDU
John O'Doherty JDOHERTY@CALTECH.EDU

California Institute of Technology

Models of reinforcement learning detail a computational framework for how agents should learn to take actions in an environment to maximize cumulative reward. Numerous studies have found implementations of components of reinforcement learning algorithms in the brain1,2. However, it is unknown how computational principles like state space representation scale up to high-dimensional environments of real-world complexity. Researchers in computer science began to tackle this problem by developing artificial neural networks, such as the deep Q network, that can learn to play Atari 2600 video games with human level performance3. These networks are loosely based

on biological nervous systems and combine the hierarchical sensory processing of convolutional networks with reinforcement learning algorithms. Thus, similar computational strategies for extracting visual features relevant to reward and action may occur in the interaction between the brain's sensory cortices and reward-learning circuitry during video game play.

Here, we have human subjects freely play three Atari video games during fMRI scanning. Using an encoding model analysis, we map representations in visual, motor, and parietal cortices to representations in the last hidden layer of the deep Q network. These neural network representations have filtered out basic and irrelevant visual information and can be interpreted as a state space for the network, as action values are directly computed as a linear combination of these features. Non-primary sensory areas, such as the precuneus, posterior cingulate cortex, and lateral parietal cortex were found to show correspondence with these models of state space across multiple games. These results provide evidence that these regions play a role in encoding state-space features that are subsequently utilized for the computation of action-values in high dimensional decision-making environments.

II-74. Quantitative analysis of excitatory-inhibitory dynamics in the thalamocortical network

I-Chun Lin[1] I.LIN@UCL.AC.UK
Michael Okun[2] M.OKUN@LE.AC.UK
Matteo Carandini[1] MATTEO@CORTEXLAB.NET
Kenneth Harris[1] KENNETH.HARRIS@UCL.AC.UK

[1]University College London
[2]University of Leicester

Although the cortical circuits are highly complex, it is suggested that their activity can be broadly captured by “mean-field” models such as those involving interactions between excitatory (E) and inhibitory (I) populations. How accurately can such models predict cortical dynamics? Can one describe local cortical dynamics with models treating a cortical region as an isolated dynamical system, or must one include interactions with other brain regions such as the thalamus?

To address these questions, we used optogenetics to independently stimulate E and I neurons in mouse primary visual cortex (V1), and measured the subsequent dynamics of E and I populations in V1, and of neurons in the lateral geniculate nucleus (LGN).

Stimulating E neurons strongly but transiently activated the E population and, after a short delay, also the I population. This activation was followed by a prolonged suppression of both populations, and a recovery above baseline (rebound) $>$100 ms later. Neither the suppression nor the rebound could be explained by I activity alone. Indeed, stimulating I neurons activated only the I population; this was followed by a shorter suppression of E and I populations with no rebound. Moreover, when the two populations were stimulated successively, the rebound was always time-locked to E stimulation.

Next, we paired stimulation of cortical E and I neurons with LGN recordings. The effects of E stimulation on LGN were similar to those seen in V1, whereas the effects of I stimulation on LGN were minimal.

These results indicate that the cortical effects of increased E activity involve the corticothalamic network, whereas the cortical effects of increased I activity may be purely intracortical, possibly operating through slow GABA components. Accordingly, a mean-field model of cortical E and I cells could predict responses to I stimulation, but modeling responses to E stimulation required an additional thalamic term.

II-75. Nonlinear mixed selectivity produces noise-tolerant neural representations

W. Jeffrey Johnston WJEFFREYJOHNSTON@GMAIL.COM
Stephanie Palmer SEPALMER@UCHICAGO.EDU
David J. Freedman DFREEDMAN@UCHICAGO.EDU

University of Chicago

Brain regions that process sensory stimuli must have efficient stimulus representations (i.e., an efficient source code) and downstream brain regions must be able to reliably read out those stimulus representations in noisy conditions (i.e., a reliable channel code). Here, we focus on the second, often overlooked, requirement and analyze a family of neural codes that vary in the degree to which different stimulus features are nonlinearly mixed with each other in the representation. Previous work has shown that nonlinearly mixing stimulus (and task) features produces linearly decodable representations, that these mixed representations exist in the brain, and that they may be behaviorally relevant. However, the previously described benefits of mixed representations assume that neural activity is interpreted by downstream brain regions using linear decoders. Here, we show that mixed representations are advantageous even when using an optimal nonlinear decoder; in particular, mixed codes have a lower probability of decoding error than codes without mixing given the same signal-to-noise ratio (SNR; with three stimulus features, a code without mixing requires 1.33 times the SNR of a mixed code to achieve 1% decoding error). We also show that some mixed codes are within error bars of the optimal trade-off between minimal error and stimulus information, while codes without mixing are up to 1.51 times greater than this minimum. Next, we introduce a cost for code dimensionality and show that more mixing is optimal for stimulus features that take on fewer values (as in areas with categorical representations, such as prefrontal cortex) while less mixing is optimal for features that take on many values (as in primary sensory areas), matching experimental observations of more and less mixed representations. Overall, this work provides a novel argument in favor of mixed stimulus representations and insight into the principles underlying the organization of sensory processing.

II-76. Relative contributions of three mammalian brain systems to computation of exploration/exploitation tradeoffs

Beatriz Godinho[1,2] BEATRIZ.GODINHO@NEURO.FCHAMPALIMAUD.ORG
Eran Lottem[1] ERAN.LOTTEM@NEURO.FCHAMPALIMAUD.ORG
Dario Sarra[1] DARIO.SARRA@NEURO.FCHAMPALIMAUD.ORG
Pietro Vertechi[1] PIETRO.VERTECHI@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[1] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Champalimaud Neuroscience Programme

During foraging animals must choose between exploiting a given patch and leaving for a different one, requiring them to balance a variety of costs, benefits and uncertainties. The mechanisms contributing to optimal foraging are likely part of the standard toolkit of the mammalian brain, but little is known about their neural implementation across brain systems. Here, we studied this problem using a task in which mice exploit water ports by nose poking. Reward delivery was probabilistic and a given port depleted over time. We first established that mouse behavior conformed to the statistics of the task in a manner consistent with optimal foraging theory. For example, when a physical barrier was used to manipulate travel time between sites, the latency to switch sites increased, consistent with normative predictions. Next, to dissect brain systems that underlie exploration/exploitation computations, we used optogenetic techniques to activate or inhibit different brain regions. We first examined serotonin, a neuromodulator implicated in patience. We found that activation of dorsal raphe nucleus serotonin neurons favored persistence of active exploitation over exploration of new sites. We next investigated two prefrontal brain regions that are reciprocally connected to the serotonin system and which are also implicated in reward and foraging behavior, the anterior cingulate cortex (ACC) and the orbitofrontal cortex (OFC). Interestingly, we found

inactivation of both regions increased tolerance to uncertainty (inhibited exploration), consistent with a negative loop with the 5-HT system. Moreover, the detailed pattern of behavioral effects indicated that OFC contributed to the inference of the hidden states of foraging sites, whereas ACC was involved in translating these inferences into switching times. These results elucidate the relative contributions of three key brain regions to the regulation of exploration and exploitation, sketching the beginnings of a map of the canonical systems underlying optimal mammalian foraging behavior.

II-77. High-dimensional representation of texture in the somatosensory cortex of primates

Justin Lieber JUSTINLIEBER@GMAIL.COM
Sliman Bensmaia SLIMAN@UCHICAGO.EDU
University of Chicago

Our sense of touch endows us with an exquisite sensitivity to surface microstructure, such that we can perceive and integrate features that range in size from tens of nanometers to tens of millimeters. In the somatosensory nerves, large surface features are encoded by the spatial patterns of activation evoked in slowly adapting type-1 (SA1) and rapidly adapting (RA) afferents, while finer surface features are encoded by precisely timed spiking patterns in RA and Pacinian-associated (PC) afferents. These spatial and temporal representations must be combined and synthesized by central mechanisms to achieve a unified percept of texture, a process about which little is known. To this end, we scanned a wide range of textures—including fabrics, furs, papers, in addition to the traditional embossed dots and gratings—across the fingertips of Rhesus macaques and recorded the responses evoked in somatosensory cortex, including Brodmann's areas 3b, 1, and 2. First, we found that texture identity is faithfully encoded in somatosensory cortex and that this texture information is distributed across neurons, who each exhibit idiosyncratic texture responses. Second, we showed that the heterogeneity across somatosensory neurons is in part driven by differences in the submodality composition of their input (SA1, RA, PC). We then found the downstream recipients of the spatial and temporal codes observed at the periphery: a subpopulation of SA1-like cortical neurons are selectively responsive to coarse textural features, whereas another PC-like population of neurons encode fine surface features. Finally, we show that texture perception, measured in human subjects, can be accounted for based on the responses of somatosensory neurons.

II-78. Mouse olfactory bulb odor responses are negatively modulated by expectation: evidence for predictive coding

Georg Raiser[1] GEORG.RAISER@RESEARCH.FCHAMPALIMAUD.ORG
Bassam V. Atallah[1] BASSAM.ATALLAH@NEURO.FCHAMPALIMAUD.ORG
Eran Lottem[2] ERAN.LOTTEM@NEURO.FCHAMPALIMAUD.ORG
Solene Sautory[1] SOLENE.SAUTORY@RESEARCH.FCHAMPALIMAUD.ORG
Cindy Poo[1] CINDY.POO@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[2] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG
[1]Champalimaud Research
[2]Champalimaud Foundation

It is widely believed that learning mechanisms allow the brain to generate and use predictions to enhance sensory processing. Yet, such predictions have been proposed to impact sensory responses in two opposite ways: (a) suppressing neural activity that is predicted, hence enhancing efficient coding (i.e. predictive coding theory), or (b) enhancing neural activity that is predicted, effectively increasing sensitivity (attentional enhancement).

To test how predictions modulate olfactory processing, we developed an expectancy-violation paradigm in which

head-fixed mice are trained to expect a sequence of three odor cues as they traversed a virtual path. We subsequently probed for behavioral and neural evidence of odor prediction on 'catch trials' (~15%) by omitting the last odor.

Consistent with odor sequence learning, mice slightly altered their running speed and sniffing frequency during odor omission. Electrophysiological recordings from the olfactory bulb demonstrated that a fraction (6%) of the cells responded during odor omission. We found that omission responses were negatively correlated with a cell's odor response magnitude, i.e. the stronger a cell's response to an expected odor, the more it was inhibited during its omission. Conversely, the stronger a cell was inhibited by an odor, the more it was excited during omission. Thus, omission responses were specific to the learned odor sequence and appear to reflect the violation of learned expectations.

These data show a pattern more consistent with a predictive coding scheme in which sensory expected signals are suppressed and only the surprising and informative signals propagate to the rest of the brain. Olfactory bulb neurons likely compute the difference between feedforward sensory signals from the periphery and top-down predictions from olfactory cortex.

II-79. Deep Nose: Training neural networks to represent the space of molecules

Ngoc Tran NTRAN@CSHL.EDU
Daniel Kepple DKEPPLE@CSHL.EDU
Alexei Koulakov KOULAKOV@CSHL.EDU

Cold Spring Harbor Laboratory

Sensing odor molecules by olfactory receptors is done in a combinatorial manner: one molecule activates many receptors; and one receptor is activated by different molecules. In human, ~350 olfactory receptors have to encode a large space of odor molecules. It is unclear which and how structural elements of odor molecules are represented in the brain. In this study, we simulate this chemical encoding process using neural networks. We hypothesize that olfactory receptors are three dimensional spatial filters whose responses serve to identify molecular features. We trained Deep Nose, a multilayer neural network whose architecture explicitly reflects the geometry of ligand-receptor interaction. First, we tested the deep autoencoder consisting of an encoder and a decoder. The encoder transforms an input molecule into a compressed representation, and the decoder performs the reverse operation to reconstruct the molecule. We used 10^5 molecular structures obtained from PubChem to train autoencoder to about 98% accuracy. Second, we trained the classifier based on annotated datasets. Deep Nose features, when paired with a neural network classifier, can predict molecules' water solubility to 97% accuracy. More generally, Deep Nose acts as a data-driven feature extractor that opens up a new "vocabulary" of chemical properties, in contrast to the knowledge-driven features such as those produced by the popular software packages, such as E-Dragon. We conclude that Deep Nose network can extract chemical features that can be used to predict various bioactivities including human odor percepts.

II-80. All your (data)base are belong to you!

Victor Hernandez-Urbina[1] VICTOR.HERNANDEZ-URBINA@CNCB.OX.AC.UK
Chaitanya Chintaluri[1] CHAITANYA.CHINTALURI@DPAG.OX.AC.UK
Panagiotis Bozelos[1] PANAGIOTIS.BOZELOS@DPAG.OX.AC.UK
Asma Motiwala[2] ASMA.MOTIWALA@NEURO.FCHAMPALIMAUD.ORG
Alexander Seeholzer[3] ALEX.SEEHOLZER@EPFL.CH
William F Podlaski[1] WILLIAM.PODLASKI@CNCB.OX.AC.UK
Tim Vogels[1] TIM.VOGELS@CNCB.OX.AC.UK

[1]University of Oxford
[2]Champalimaud Foundation

[3]Ecole Polytechnique Federale de Lausanne

Rapid progress in data acquisition through experiments and modelling, and ongoing efforts to make such data publicly accessible, have catapulted neuroscience into the age of big data. In addition to having access to data from various labs, it will be crucial to acquire tools that allow scientists to sift through the data—to browse, study and discover, i.e. to extend our knowledge based on these data. Unlike e.g., the Allen Institute or the Blue Brain Project, which can direct substantial resources towards data accessibility, smaller labs often lack the infrastructure and resources required to store and visualize their data effectively. To overcome this limitation, we are developing a web-based 'Big Data Visualizer' (BDV) that separates data from visualization: our web-based engine can be 'plugged into' any data presented in a standardized format. Our framework then processes the information, presents relationships between data items in a graphical fashion, and produces spec-sheets based on its analyses. It showcases the results interactively through a standard web browser, allowing in-depth exploration and data discovery by anyone. The BDV is based on the engine of the Ion Channel Genealogy database (Podlaski et al., 2017), which provides visualization and comparison of more than 2000 publicly available ion channel models. We adapted this framework to a database of experimentally obtained neuron morphologies, to browse their metadata and reveal their similarities. Moreover, we demonstrate the tool's generic utility by analysing a dataset of past Cosyne abstracts, creating a searchable map of topics, authors, and collaborators. Importantly, the BDV is a work in progress, and we recently assembled a new team to continue its development. If you are interested in using our engine, or you have ideas of how to optimize its utility, come by and discuss.

II-81. Learning predictability in the input with balanced spiking networks

Lyudmila Kushnir KUSHNIRLV@GMAIL.COM
Sophie Deneve SOPHIE.DENEVE@ENS.FR

Ecole Normale Superieure

To be efficient, a network should only encode what it cannot predict. How can a spiking network learn this in an unsupervised manner? We show that a single principle, namely excitation-inhibition balance, applied on several time scales, leads to learning the spatial and temporal correlation structure of the networks' input. Specifically, we propose a model network of integrate-and-fire neurons based on (Boerlin, Machens, Deneve 2013) and (Brendel, Bourdoukan, Vertechi, Machens, Deneve 2017) that learns the low-dimensional structure of its feedforward input and its temporal correlations. As fast recurrent connections learn the low-dimensional structure, the firing rate dramatically decreases. The proposed learning rule is biologically plausible. We comment on its relation to the learning rule of (Brendel et al 2017). When the input goes out of the learned low-dimensional subspace, the network fires at higher firing rate until the connectivity is adjusted. In the same way, the network learns to predict future input in order to decrease its firing even further. The matrix of slow connections that allows for this is also learned in an unsupervised way via a biologically plausible learning rule, whose aim is, as in the case of fast connections, to keep the membrane voltage at zero. In this case, only the deviations from the predicted input on the slower time-scale are represented by the neuronal spiking. To capture predictability of the input on the even longer time scale we introduce a population of interneurons that learn its dynamics and keep the principal network below firing threshold until the learned dynamics is violated. The full feedforward input is still encoded in the activity of the principal network. To summarize, our model solves the problem of learning the repeating input pattern and encoding the deviations from it by enforcing excitation-inhibition balance.

II-82. Spatially constrained model of a mesoscopic whole-brain connectivity: insights from network dynamics

Hannah Choi[1] HANNAHCH@UW.EDU
Stefan Mihalas[2] STEFANM@ALLENINSTITUTE.ORG

[1]University of Washington
[2]Allen Institute for Brain Science

How do we represent the weighted, mesoscopic whole-brain network using minimal information? Structural neural connectivity and its implications on brain function have been a long-sought subject in neuroscience. However, previous studies have been limited either to small networks or coarser connectivity, often binarized and without spatial information. Recent development of Allen Mouse Brain Connectivity Atlas provides us the unique opportunity to investigate precise weighted anatomical connectivity of the mammalian whole brain network. Using the latest mesoscopic connectivity data from a new mapping algorithm, we seek a parsimonious representation of the weighted whole-brain network that captures the network properties to the full extent. We find that the connectivity has a significant spatial dependence—the connection strength decreases with distance between the regions following a power law. However, we found a few positive residuals, indicating strong connections between distal brain regions unpredicted by the power law relationship. To probe possible implications of the residual connections on the network dynamics, we constructed a network of phase oscillators with the data-driven adjacency matrix, and compared its dynamics to those of the oscillator network with the connections following the power-law spatial dependence. Surprisingly, a small perturbation causes a rapid transition between localized and global synchronies in the data-driven network, but fails to do so in the artificial network. Such steep phase transition can be introduced to the artificial network by adding a small fraction of the strong residual connections, suggesting that these residuals may underlie brain's ability to rapidly switch between global and modularized computations, a feature impaired in pathological conditions such as Alzheimer's disease. The mesoscopic whole-brain network, therefore, is not fully described by its spatial embedding; additional complexity of the network plays a critical computational role. A spatially-constrained model plus an idiosyncratic sparse matrix provides a parsimonious representation of the measured connectivity.

II-83. Revealing the neural correlates of behavior without relying on behavioral measurements

Alon Rubin ALON.RUBIN@GMAIL.COM
Liron Sheintuch LIRONSHEINTUCH@GMAIL.COM
Or Pinchasof OR.PINCHASOV@GMAIL.COM
Noa Brande-Eilat N.BRANDE.EILAT@GMAIL.COM
Nitzan Geva NITZGEVA@GMAIL.COM
Yoav Rechavi RECHAVI@GMAIL.COM
Yaniv Ziv YANIV.ZIV@WEIZMANN.AC.IL

Weizmann Institute of Science

Revealing the neural correlates of behavior has been a leading approach to study the neural code. Traditionally, this has been done by calculating tuning curves, i.e., neuronal responsiveness to an examined external variable. While measurement of neuronal tuning curves has been instrumental to many discoveries in neuroscience, it requires an a-priori selection of external variables, which limits the findings to these variables only. Recent technological advancements enable simultaneous readout of activity from large neuronal populations, promoting a shift in the analysis of neuronal activity data from a single-neuron to a population-level perspective. Here, we applied unsupervised learning to study large-scale Ca2+ imaging data recorded from the hippocampus of freely behaving mice, and exposed the internal structure of neuronal activity. This structure allowed defining for each neuron an internal tuning curve that characterizes its activity relative to the network activity, rather than relative to any pre-defined external variable. We found that internal tuning curves closely matched "classic" tuning curves of place cells, supporting the reconstruction of spatial representation without behavioral measurements. Because internal tuning curves are defined irrespective of behavioral measurements, their calculation circumvents the need to focus on pre-defined external variables. Using similar investigation of neuronal activity recorded from the pre-frontal cortex, we exposed a schematic representation of locations and actions, and discovered the encoding of a schematic variable, the 'trajectory phase'. Additionally, we found that the internal structure of neuronal activity was conserved across days and across mice. Overall, our results suggest that the internal structure of neuronal activity may serve as a fingerprint of computations undertaken by the network, enabling the discovery of variables hidden within a neural code.

II-84. Robust dendritic computations with sparse distributed representations

Subutai Ahmad SAHMAD@NUMENTA.COM
Max Schwarzer MAXA.SCHWARZER@GMAIL.COM
Jeff Hawkins JHAWKINS@NUMENTA.COM

Numenta

Empirical evidence suggests that the neocortex represents information using sparse distributed patterns of activity. There exist a variety of sparse coding algorithms demonstrating how to compute sparse representations, and a number of mathematical results on the capacity of sparse representations. Here we focus on dendritic computations and analyze properties of sparse representations from a machine learning viewpoint. Are sparse representations useful for neuronal pattern recognition, and under what conditions? We propose a formal mathematical model for recognition accuracy of binary sparse representations using active dendrites. We derive scaling laws that characterize the chance of false positives and false negatives when detecting patterns under adverse conditions. We describe three primary results. First, we show that using very high dimensional sparse representations, a network of neurons can reliably classify a massive number of patterns under extremely noisy conditions. The results hold even when synapses subsample a tiny subset of the target patterns or when individual neurons themselves are unreliable. Second, the equations predict optimal dendritic NMDA spiking thresholds that closely match experimental findings. Finally, we consider two existing computational models of active dendrites: the Poirazi/Mel neuron and the HTM neuron. Through simulations we show that the scaling behavior of these two models closely matches the theory. We show dramatically improved recognition accuracy over published results when "good parameters" (as predicted by the theory) for sparsity and dimensionality are applied. The theory presented here complements existing work and represents a practical mathematical framework for understanding the accuracy and robustness of sparse representations in cortical networks.

II-85. Non-spatial neural replay in building and updating world models in humans

Yunzhe Liu[1] YUNZHE.LIU.16@UCL.AC.UK
Zeb Kurth-Nelson[2] ZEBKURTHNELSON@GMAIL.COM
Ray Dolan[1] R.DOLAN@UCL.AC.UK
Tim Behrens[1,3] BEHRENS@FMRIB.OX.AC.UK

[1]University College London
[2]Google DeepMind
[3]University of Oxford

Flexible behaviour is thought to be underpinned by neural models of the world that account for the relationships between experiences. In rodent spatial navigation, hippocampal and neocortical replay and preplay are considered important for building such models. In replay, patterns of cellular firing during rest spontaneously play out past spatial trajectories in both forward and reverse directions, and reverse replays are increased for rewarded

trajectories. This computation is also implied by Dyna-type learning algorithms. In preplay, future trajectories are played out, perhaps constrained by structural knowledge of possible relationships in the world. However, despite the theoretical importance of neural replay, its study has so far been largely restricted to spatial navigation tasks in rodents. Furthermore, it is unclear whether knowledge of task structure can indeed impact on the sequences that are played during rest. Using magnetoencephalography (MEG), our goal here was to test 1) is there non-spatial replay during rest in humans, and 2) can replay generalise learnt structure to new stimuli? We exploited a revised sensory preconditioning paradigm to build two distinct sequences: participants were presented with pairwise associations where the correct sequence is jumbled in time, and pre-trained to re-assemble them in the right order. By building pattern classifiers for MEG sensor activity for each stimulus, we could detect sequences of their reactivations during rest. These reactivations recapitulated known features of hippocampal replay but in a non-spatial task. Transitions were rapid (40-70ms lag) and forward replay transitioned to reverse replay after new learning (including after reward). Notably, the replayed sequences reflected the correctly re-assembled sequence, not the viewed sequences, implying that task structural knowledge can impose constraints on replay. Our data extend known spatial replay mechanisms to human reinforcement learning and suggest such replay can reflect new sequences beneficial for future behaviour rather than simply recapitulating past events.

II-86. A neurocomputational approach of controllability

Romain Ligneul[1] ROMAIN.LIGNEUL@RESEARCH.FCHAMPALIMAUD.ORG
Roshan Cools[2] ROSHAN.COOLS@GMAIL.COM

[1]Champalimaud Foundation
[2]Donders Institute for Cognitive Neuroimaging

In order to guide behavior efficiently and to optimize the allocation of computational resources, agents must monitor their degree of control over the environment in a wide variety of contexts. Controllability monitoring is key to alleviate the credit assignment problem, to predict future events and to arbitrate between proactive and reactive behavioral strategies. However, very little is known about the computations underlying this ability beyond the context of simple tasks such as the learned helplessness paradigm. Here, I will present a computational model able to track controllability estimates immune to confounding sources of statistical regularity. By comparing prediction error signals generated by two first-order learning modules tracking respectively state-state (SS') and state-action-state (SAS') contingencies, this model can derive a second-order variable approximating the average information associated with one's own actions. A behavioral study demonstrated the superiority of the SS'SAS' architecture over classical model-based learning algorithms. Participants were invited to explore an environment whose dynamics was covertly governed by controllable or uncontrollable transition rules which had to be learnt to perform accurate predictions regarding upcoming events. In an fMRI study based on the same paradigm, a left-sided parietofrontal network correlated with controllability estimates when participants made explicit predictions, whereas the left angular gyrus and the ipsilateral sensorimotor cortex correlated with controllability prediction errors during exploration. More importantly, BOLD responses in the subgenual anterior cingulate cortex (sACC) reflected the difference in SS' and SAS' prediction errors, a key quantity involved for the updating process. This effect likely induced a downward bias in controllability estimates which exerted a detrimental influence on prediction performances. Taken together, these results shed light on the mechanisms through which humans monitor the causal impact of their actions and point towards sACC activity as the putative neural origin of distorted controllability representations, a cardinal feature of several psychiatric disorders.

II-87. Disentangling neural population variability using time-warped point-process GPFA

Lea Duncker[1] DUNCKER@GATSBY.UCL.AC.UK
Maneesh Sahani[2] MANEESH@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit
[2]University College London

Neural firing across repeated trials is inherently variable. This variability may partly reflect noise at the level of individual neurons or synapses; however, variability shared across a population may also arise from internal processes such as decision-making, movement preparation, or attention that might differ in content or time course across trials even when external stimuli and behavioural task are held constant. In recent years, there has been substantial interest in statistical methods that can leverage simultaneous recordings of neural population activity in order to extract meaningful structure on a single-trial basis. While these methods have facilitated analyses that do not rely on averaging over repeated trials, questions about the nature of population-level inter-trial variability remain to be addressed. Here, we extend continuous-time point-process GPFA in order to extract single-trial dynamical trajectories from neural population spike trains collected over repeated trials with different experimental conditions and variable time courses. Our approach disentangles contributions to the population activity on each trial made by an overall mean latent process, condition-specific latents, and latent processes capturing additional trial-to-trial variation. We allow each of these latent components to evolve in potentially separate subspaces. To take into account differences in the temporal evolution of each trial, our method incorporates probabilistic time-warping and thereby automatically aligns trials in the latent space. We apply our method to macaque motor cortical population activity from a variable-delay centre-out reaching task. The inferred time-warping automatically resolves the variable time-course of different trials, so that behaviourally-salient events like movement onset are reliably anchored to a neurally-defined time point. After time warping, the remaining inter-trial variability is of low dimension and falls mostly outside the subspaces related to the overall mean and condition-specific activity.

II-88. The neural circuit basis of feature-binding in working memory

Joao Barbosa[1] PALERMA@GMAIL.COM
Ainsley Temudo[2] AINSLEY.TEMUDO@NYU.EDU
Vahan Babushkin[2] VAHAN.BABUSHKIN@NYU.EDU
Tim Buschman[3] TBUSCHMA@PRINCETON.EDU
Kartik Sreenivasan[2] KARTIK.SREENIVASAN@NYU.EDU
Albert Compte[4] ACOMPTE@CLINIC.CAT

[1]IDIBAPS
[2]New York University, Abu Dhabi
[3]Princeton University
[4]IDIBAPS Barcelona

Binding (or swap) errors occur in working memory tasks when a wrong response is in fact accurate relative to a non-target stimulus [1]. These errors reflect the failure to maintain bundled in memory the conjunction of features that define one object, and the mechanisms implicated remain unknown. Here, we tested the mechanism of synchrony across feature-specific neural assemblies [2]. We built a biophysical neural network model for working memory items defined by one color and one location. The model is composed of two one-dimensional attractor networks for working memory (as in [3]), one representing colors and the other one locations. These two networks are then connected via weak cortico-cortical excitation. Gamma-oscillations were induced during bump attractor activity through the interplay of fast recurrent excitation and slower feedback inhibition [3]. Binding between color and location was accomplished through the synchronization of pairs of bumps across the two networks via weak cortico-cortical excitation. As a result, different memorized items were held at different phases of the network's intrinsic oscillation. In some simulations, swap errors arose: "color bumps" abruptly changed their

phase relationship with “location bumps”. The model makes specific testable predictions that we addressed experimentally. Firstly, a uniform drive pulsating at the natural frequency of the networks stabilizes the bumps and reduces the incidence of swap errors. This was validated in behavioral experiments with oscillating visual placeholders, with a specific swap-reducing effect at 10 Hz. Secondly, swap errors in the model are associated with a lower phase consistency of oscillatory activity in the delay period. We validated this prediction in MEG experiments, finding alpha-band phase changes specific to swap trials in fronto-parietal sensors.

Significance We propose a plausible mechanism for working memory binding based on neural synchronization in spiking neural networks, and we support it with behavioral and neurophysiological (MEG) experiments in humans.

II-89. Breakdown of spatial coding and neural synchronization in epileptic mice

Tristan Shuman[1] TRISTAN.SHUMAN@MSSM.EDU
Daniel Aharoni[2] DBAHARONI@GMAIL.COM
Denise Cai[1] DENISECAI@GMAIL.COM
Christopher Lee[1] CHRISTOPHER.LEE1@MSSM.EDU
Peyman Golshani[2] PGOLSHANI@MEDNET.UCLA.EDU

[1]Icahn School of Medicine at Mount Sinai
[2]University of California, Los Angeles

Epilepsy causes significant cognitive deficits in both humans and rodent models yet the circuit mechanisms leading to cognitive dysfunction remain unknown. We developed a wireless miniature microscope (Miniscope) for in vivo calcium imaging to examine the spatial representation of freely behaving epileptic and control mice running on a linear track. Following pilocarpine-induced status epilepticus, chronically epileptic mice had reduced spatial information and stability of CA1 neurons both within and across days. Deficits in place cell stability emerged after just 30 minutes between imaging sessions and degraded to chance levels by 7 days. Both individual place cells and population coding was severely disrupted in the epileptic mice as well as decoding accuracy within and across sessions. To further examine how the hippocampal circuit is dysfunctional in epileptic mice we used silicon probes to record hippocampal interneurons during head-fixed virtual navigation. We found that epileptic mice had profound deficits in theta and gamma power and coherence, and altered phase preferences to ongoing theta oscillations. In particular, dentate hilar interneurons had a similar magnitude of theta phase modulation of their firing rate, but the preferred phase of these cells as a group was highly dispersed. This led to a desynchronization between the firing of dentate gyrus and CA1 interneuron populations in the epileptic mice. Finally, we transplanted embryonic interneuron precursors from the medial ganglionic eminence into the hippocampus of epileptic mice and found a partial rescue of theta power in the transplanted side. Together, these results demonstrate that spatial processing is severely disrupted in chronically epileptic mice and this dysfunctional circuit likely contributes to the cognitive deficits associated with epilepsy.

II-90. Neural population dynamics underlying motor learning transfer

Saurabh Vyas[1] SMVYAS@STANFORD.EDU
Nir Even-Chen[1] NIREC@STANFORD.EDU
Sergey Stavisky[1] SERGEY.STAVISKY@STANFORD.EDU
Stephen Ryu[2] SEOULMANMD@GMAIL.COM
Paul Nuyujukian[1] PAUL@NPL.STANFORD.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU

[1]Stanford University
[2]Palo Alto Medical Foundation

Understanding motor-related covert mental processes (e.g., mental rehearsal) is tantalizing as decades of human behavioral studies have shown that such internal behavior can exhibit varying degrees of motor learning transfer to overt performance. Working theories posit that such transfer is a result of covert learning engaging neural population activity similar to that employed during overt practice. These results, however, are still debated, and do not propose mechanistic hypotheses about why neural similarity is helpful for learning transfer. The challenge lies in the fact that covert processes are open-loop and hidden, where neither the experimenters nor the subjects observe the trial-by-trial progression of learning. Here, we present a covert process that enables a direct and real-time probe into this evolution, by “closing the loop” via a brain-machine interface (BMI). We use a BMI that mathematically regresses neural activity from dorsal premotor and primary motor cortex onto two-dimensional cursor kinematics. Using this BMI, we describe a “covert rehearsal” paradigm whereby subjects can, in a sense, “rehearse” visuomotor rotation (VMR) tasks, without overt movements, using directly their neural activity. We then evaluate the degree of learning transfer by having the subjects repeat the same task via overt arm reaches. If transfer is observed across these contexts, the corresponding neural activity would provide a glimpse at its mechanism. Note that we do not equate covert rehearsal to mental rehearsal (though this may well be the case). Instead, we evaluate two key scientific questions underlying most covert processes: 1.Can covert processes (which covert rehearsal is a type of) facilitate overt motor learning? (Yes) 2.If so, what neural mechanism mediates this transfer of learning? (Shared neural preparatory states)

II-91. Amygdala-cortical projection pathways that enable hidden state expectations

Nina Lichtenberg NLICHTENBERG@UCLA.EDU
Zachary Pennington ZACHPEN87@GMAIL.COM
Sandra Holley SHOLLEY@MEDNET.UCLA.EDU
Venuz Greenfield VNZGREENFIELD@UCLA.EDU
Linnea Sepe-Forrest LINNEASF25@GMAIL.COM
Carlos Cepeda CCEPEDA@MEDNET.UCLA.EDU
Michael Levine MLEVINE@MEDNET.UCLA.EDU
Kate Wassum KWASSUM@UCLA.EDU

University of California, Los Angeles

To make an appropriate decision, one must anticipate potential future rewarding events, even when they are not readily observable. These expectations are generated by using observable information (e.g., stimuli or available actions) to retrieve detailed memories of available rewards. The basolateral amygdala (BLA) is one identified key node in the circuitry supporting reward-related behavior. But whether the BLA contributes to adaptive behavior and choice online and the output pathways through which it might achieve this function are both unknown. To address this, we identified BLA projections to both the lateral (lOFC) and medial orbitofrontal cortex (mOFC) that were largely anatomically distinct. Using a chemogenetic approach, we evaluated the unique function of each pathway in expectation-guided behavior. Inactivation of BLA terminals in the lOFC was found to disrupt the influence of cue-triggered reward expectations over both reward-seeking decisions and adaptive conditional goal-approach responding. Inactivation of BLA terminals in the mOFC disrupted only the latter, leaving cue-directed decision making intact. Neither projection was necessary when actions were guided by reward expectations generated based on learned action-reward contingencies. Correspondingly, BLA to OFC projections were activated by reward-predictive cues. If, as suggested, the OFC represents the current, not fully observable state, then these results suggest that BLA projections enable predictive stimuli to provide the OFC with detailed expectations of potential rewards available in that state. BLA to lOFC projections might serve a more primary function in this respect, while BLA to mOFC projections are only required when appropriate responding requires an understanding that, although things have not perceptually changed (e.g., CS presence), the state is nonetheless different because the anticipated reward is no longer valuable. The cognitive symptoms underlying many psychiatric disorders, including addiction, result from a failure to appropriately anticipate potential future events. These data, therefore, have implications for the understanding and treatment of these conditions.

II-92. Corticostriatal activity is target cell type-specific during a skilled movement

Haixin Liu LHX1986@GMAIL.COM
Keelin O'Neil KEELINONEILWORK@GMAIL.COM
Madalyn DeViso MDEVISO@UCSD.EDU
Varoth Lilascharoen VLILASCH@UCSD.EDU
Byung Kook Lim BKLIM@UCSD.EDU
Takaki Komiyama TKOMIYAMA@UCSD.EDU

University of California, San Diego

Cortical projection to the striatum is one of the major inputs to the basal ganglia and is considered critical for motor learning and skilled movements. The striatum consists of spiny projecting neurons (SPNs) and local interneurons, which are intermingled with each other. SPNs can be divided into two subtypes; direct pathway SPNs (dSPNs) project to the basal ganglia output nuclei directly and express dopamine D1 receptor (D1R), while indirect pathway SPNs (iSPNs) project to the intermediate nuclei and express dopamine D2 receptor and adenosine A2a receptor (A2A). Local interneurons include GABAergic and cholinergic cells. All of these cell types receive cortical inputs and serve distinct functions. However, whether and how different cell types in the striatum receive distinct information from the cortex is unknown. Here, we combined two-photon Ca2+ imaging with cell-type specific monosynaptic retrograde labeling using a modified rabies virus (EnvA) to investigate the activity of corticostriatal neurons in the motor cortex presynaptic to dSPNs, iSPNs and cholinergic interneurons during a skilled movement. We found that target-defined corticostriatal neurons across the three groups contained similar response profiles. However, the distribution of the modulation profiles were distinct among these groups. Neurons targeting iSPNs had the highest percentage of neurons that were suppressed during movements and had more neurons modulated when mice cease moving, compared to neurons targeting dSPNs. Movement-modulated neurons that are presynaptic to cholinergic interneurons were almost exclusively activated during movements. Our results provide the first functional study of the presynaptic inputs of three different neuronal subtypes in the striatum and suggest a synergistic scheme of the corticostriatal pathways during a skilled movement, in which all pathways are activated but with distinct modulation profile distributions.

II-93. Receptive field estimation from spikes via modeling of calcium-related fluorescence

Peter Ledochowitsch PETERL@ALLENINSTITUTE.ORG
Nicholas Cain NICHOLASC@ALLENINSTITUTE.ORG
Gabriel Ocker GABEO@ALLENINSTITUTE.ORG
Clay Reid CLAYR@ALLENINSTITUTE.ORG
Stefan Mihalas STEFANM@ALLENINSTITUTE.ORG
Michael Buice MICHAELBU@ALLENINSTITUTE.ORG
Josh Siegle JOSHS@ALLENINSTITUTE.ORG
Xiaoxuan Jia XIAOXUANJ@ALLENINSTITUTE.ORG
Ulf Knoblich KNOBLICH@GMAIL.COM
Lawrence Huang LAWRENCEH@ALLENINSTITUTE.ORG
Lu Li LUL@ALLENINSTITUTE.ORG
Daniel Millman DANIELM@ALLENINSTITUTE.ORG
Michael Oliver MICHAELO@ALLENINSTITUTE.ORG
Hongkui Zeng HONGKUIZ@ALLENINSTITUTE.ORG
Shawn Olsen SHAWNO@ALLENINSTITUTE.ORG
Saskia de Vries SASKIAD@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Historically, most studies of receptive fields were undertaken using electrophysiology, with a recent shift towards calcium imaging as a method of choice. Electrophysiology and Ca2+ imaging both constitute noisy observations of the underlying neural activity. However, the statistical properties of the raw signals, as well as the sources and effects of the introduced noise, are quite distinct. To interpret receptive fields derived from Ca2+ imaging in the context of existing literature, it is crucial to understand how these differences are reflected in the detectability and properties of receptive fields derived from different recording modalities, for it is impractical to perform every experiment with each modality. To address this issue, we first computed receptive fields from extracellular spike trains recorded in mouse visual cortex, directly, using the spike-triggered average (STA). Then, we calibrated a biophysically inspired model that relates spiking activity to observed fluorescence (MLSpike, [1]) on 'ground truth' data, in vivo Ca2+ recordings paired with juxtacellular electrophysiology, where the Ca2+-dependent fluorescence was consistent with the Allen Brain Observatory (http://observatory.brain-map.org/visualcoding/), a public resource providing standardized in vivo characterization of single neuron activity in mouse visual cortex based on Ca2+-imaging. Following calibration, we computed model calcium activity from above spike trains, and analyzed the synthesized Ca2+ data using techniques developed for mapping of classical receptive fields based on responses to locally sparse noise in the Allen Brain Observatory data processing pipeline. We found that such analysis readily yielded receptive fields, which largely agreed with those identified directly from the electrophysiological recordings via STA, and investigated the sensitivity of the obtained receptive field structure to the parameterization of the Ca2+ forward model. In the future, this data-driven modeling approach may provide a Rosetta Stone for receptive field comparison across recording modalities, as well as inspire improvements to algorithmic receptive field fitting procedures.

II-94. Emergent elasticity in the neural code for space

Sam Ocko SOCKO@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Kiah Hardcastle KHARDCAS@STANFORD.EDU
Lisa Giocomo GIOCOMO@STANFORD.EDU

Stanford University

Upon encountering a novel environment, an animal must construct a consistent map of the environment by combining information from self-motion cues and sensory landmark cues. How do known aspects of neural circuit dynamics and synaptic plasticity conspire to accomplish this feat? Here we show analytically how a neural attractor model that combines path integration of self-motion with Hebbian plasticity in synaptic weights from landmark cells can self-organize a consistent map of space as the animal explores an environment. Intriguingly, the emergence of this map can be shown mathematically to be an elastic relaxation process between landmark cell synapses mediated by the attractor network, yielding a self-consistent map even for arbitrary environments. Moreover, our model makes several experimentally testable predictions about spatial representations in the medial entorhinal cortex, including: (1) systematic deformations in the firing fields of grid cells in irregular environments, akin to elastic deformations of solids forced into irregular containers, (2) systematic path-dependent shifts in the firing fields of grid cells towards the most recently encountered landmark, even in a fully learned environment, and (3) the creation of topological defects in grid cell firing patterns through precise environmental manipulations. Taken together, our results conceptually link known biophysical aspects of neurons and synapses to an emergent solution of a fundamental computational problem in navigation, while providing a unified account of disparate experimental observations.

II-95. Memory compression in the hippocampus leads to the emergence of place cells

Marcus Benna MKB2162@COLUMBIA.EDU
Stefano Fusi SF2237@COLUMBIA.EDU

Columbia University

The observation of place cells in the hippocampus has suggested that this brain area plays a special role in encoding spatial information. However, several studies show that place cells do not only encode position in physical space, but that their activity is in fact modulated by several other variables, which include the behavior of the animal (e.g. speed of movement or head direction), the presence of objects at particular locations, their value, and interactions with other animals. Consistent with these observations, place cell responses are reported to be rather unstable, indicating that they encode multiple variables, many of which are not under control in experiments, and that the neural representations in the hippocampus may be continuously updated. Here we propose a memory model of the hippocampus that provides a novel interpretation of place cells and can explain these observations. We hypothesize that the hippocampus is a memory device that takes advantage of the correlations between sensory experiences to generate compressed representations of the episodes that are stored in memory. We have constructed a simple neural network model that can efficiently compress simulated memories. This model naturally produces place cells that are similar to those observed in experiments. It predicts that the activity of these cells is variable and that the fluctuations of the place fields encode information about the recent history of sensory experiences. Our model also suggests that the hippocampus is not explicitly designed to deal with physical space, but can equally well represent any variable with which its inputs correlate. Place cells may simply be a consequence of a memory compression process implemented in the hippocampus.

II-96. Ga-based parameter optimization of DWI-based global fiber tracking with neuronal tracer signal as a reference

Carlos Gutierrez[1] CARLOSENGUTIERREZ@GMAIL.COM
Henrik Skibbe[2] HENRIK.SKIBBE@GMAIL.COM
Ken Nakae[2] NAKAE-K@SYS.I.KYOTO-U.AC.JP
Alexander Woodward[3] ALEXANDER.WOODWARD@RIKEN.JP
Akiya Watakabe[3] AKIYA.WATAKABE@RIKEN.JP
Junichi Hata[3] J.HATA@A3.KEIO.JP
Hideyuki Okano[3] HIDOKANO@A2.KEIO.JP
Tetsuo Yamamori[3] TETSUO.YAMAMORI@RIKEN.JP
Yoko Yamaguchi[3] YOKOY@BRAIN.RIKEN.JP
Shin Ishii[2] ISHII@I.KYOTO-U.AC.JP
Kenji Doya[1] DOYA@OIST.JP

[1]Okinawa Institute of Science and Technology Graduate University
[2]Kyoto University
[3]RIKEN Brain Science Institute

The Brain Mapping by Integrated Neurotechnologies for Disease Studies (Brain/MINDS) project maps the brain of the common marmoset (Callithrix Jacchus) across multiple scales (Okano et al., 2016). As a part of this project, we validate and optimize diffusion weighted MR image (DWI)-based fiber tracking results (Freiburg Fibertools global tracking algorithm (Reisert et al., 2011, 2013)) by comparison with information obtained from neural AAV-based fluorescent tracer injected into the prefrontal cortex of the same subject. To evaluate a parameter set, we generated and evaluated density maps from the two image modalities. We assumed that parameters were good when both density maps are similar at the voxel level. The procedure was as follows: (1 Two photon microscopy). Based on a tracer injection into a specific brain region (injection site), we obtained an axon density map. (2 DWI). We selected a set of tracking parameters to compute a full brain DWI tractogram. We extracted a subset of fibers originating from the injection site of (1) and convert it into a fiber density map. (3 Evaluation). We compared the density maps from (1) and (2) to obtain an objective function based on the true positive rate and the false positive rate. We implemented an evolutionary algorithm for search of the DWI fiber tracking parameters in (2) and optimized our objective function. Best results from DWI tractography can complement the sparse structural connectivity obtained from tracer injections, improve connection quantification between source areas and targets, and provide a reliable nondestructive 3D brain-wide connectivity mapping method.

II-97. Internal models of sensorimotor integration regulate cortical dynamics

Seth Egger SWEGGER@MIT.EDU
Evan Remington EREMING@MIT.EDU
Chia-Jung Chang CHIAJUNG@MIT.EDU
Mehrdad Jazayeri MJAZ@MIT.EDU

Massachusetts Institute of Technology

Theoretical considerations and psychophysical studies of sensorimotor integration describe the dynamic regulation of behavior in terms of three computational building blocks: a controller (i.e., inverse model), a simulator (i.e., forward model) and a state estimator (i.e., Bayesian estimator). Although this framework has had a profound impact on our understanding of motor control, its utility for understanding the brain's control principles in cognitive tasks has not been verified. We tackled this problem by designing a timing task that required control of internal states in the absence of any movement. Recording from the frontal cortex of monkeys performing this task revealed the interplay of a controller, a simulator and a Bayesian estimator during the evolution of the underlying neural states. Our findings provide direct evidence that the nervous system controls internally-generated dynamics by establishing task-relevant internal models.

II-98. A modular neural network model of the primate grasping circuit

Jonathan A Michaels[1] JONATHAN.ASHER.MICHAELS@GMAIL.COM
Stefan Schaffelhofer[2,3] SSCHAFFELH@MAIL.ROCKEFELLER.EDU
Andres Agudelo-Toro[2] AAGUDELO-TORO@DPZ.EU
Hansjorg Scherberger[4] HSCHERBERGER@DPZ.EU

[1]Stanford University
[2]German Primate Center
[3]Rockefeller University
[4]University of Gottingen

Grasping objects is an essential part of primate behavior. In macaque monkeys, the core of the grasping circuit is formed by the interconnected anterior intraparietal area (AIP), the hand area (F5) of the ventral premotor cortex, and the hand area of the motor cortex (M1). Generating appropriate delayed grasping movements involves many inter-related steps, from identification of visual target identity and spatial location, to the determination and maintenance of the appropriate movement plan, and finally the control of muscles. We hypothesized that the grasping circuit could be effectively modeled by training a modular recurrent neural network on visual object features to output muscle dynamics. To train and test our model, we recorded from neural populations simultaneously from AIP, F5, and M1 using floating microelectrode arrays while two macaque monkeys performed a delayed grasping task in which ~50 objects of distinct shape, size, and orientation had to be grasped and lifted. During every trial, arm and hand kinematics were recorded and transformed into a 50-dimension muscle length space using a musculoskeletal model. The network model was successfully trained to produce single-trial muscle velocities during grasping (normalized error: <5 %). Interestingly, the internal dynamics of the model matched the recorded neural data (canonical correlation, mean r=0.7 over 12 dimensions). Furthermore, biological regularizations were

implemented to encourage simplistic solutions, which resulted in a strong alignment between the contributions of modules of the model and the recorded brain areas to the canonical variables (r=0.80) that was not present in untrained networks (r=-0.06). Our model therefore provides a simplistic and accurate representation of the primate grasping circuit and suggests that the combined processing of these areas can be well understood as a network optimized to transform object information into the muscle dynamics required to grasp each object.

II-99. Predicting the emotional content of images with convolutional neural networks and visual cortex activity

Philip Kragel PHILIP.KRAGEL@COLORADO.EDU
Marianne Reddan MARIANNE.REDDAN@COLORADO.EDU
Tor Wager TOR.WAGER@COLORADO.EDU

University of Colorado, Boulder

Developments in computational neuroscience have used deep convolutional neural networks (CNN) to model population responses in visual cortex. This approach has produced computational models of object categorization that are both highly accurate and consistent with brain representations of visual images. However, it remains unclear how the emotional content of images is represented in the brain. To this end, we fine-tuned the last three layers of a CNN to classify images along 20 emotion categories. We validated this model in three ways: testing it on hold-out images from the same stimulus set; applying it to an established stimulus set in psychological research, the International Affective Picture System (IAPS); and using it to predict the genre of films based on their trailers. To establish a mapping between the CNN and brain responses to emotional images (with varied content ranging from negative to positive), we developed a multivariate brain-based model to predict activations in the last fully-connected layer of the CNN using patterns of visual cortex activity measured via fMRI (n = 18). The brain-based model predicted significant trial-to-trial variability in CNN activations in data from independent subjects (cross-validated r = 0.27, $P<.001$, permutation test). We additionally conducted a follow-up generalization test on data from an independent study (n = 127) where individuals were presented IAPS images that conveyed scenes with neutral or negative content. The outcome of interest was a brain signature developed to measure the unpleasantness of visual scenes, the Picture Induced Negative Emotion signature (PINES). The brain-based model explained 34.7% of the trial-to-trial variability in the PINES response ($P<.001$, permutation test). These findings establish a novel computational account of emotional processing that relates abstract image features to distinct emotion categories and ground the model in distributed population-level activity in visual cortex.

II-100. Uncertainty-dependent exploration accounts for lapses in perceptual decisions

Sashank Pisupati SPISUPAT@CSHL.EDU
Lital Chartarifsky LCHARTAR@CSHL.EDU
Anne Churchland CHURCHLAND@CSHL.EDU

Cold Spring Harbor Laboratory

During perceptual decisions, highly trained subjects can make errors on easy stimuli, often referred to “lapses”. Proper treatment of these lapses is crucial for estimating decision parameters, assessing training progress and interpreting inactivations. However, the factors that modulate and cause lapses remain poorly understood; current models treat them as stimulus-independent coin-flips made on some trials due to lack of task engagement. Here, we demonstrate that perceptual uncertainty modulates lapse rates, and propose uncertainty- driven exploration as the underlying cause. To achieve this, we manipulated uncertainty on an audiovisual rate discrimination task in rats using 2 strategies: (1) varying the signal-to-noise ratio of individual sensory events and (2) presenting events in a unisensory vs. multisensory context. Reduced uncertainty decreases the probability of guessing

in both cases. This novel effect, not due to parameter trade-offs in fitting, is captured with a reduced model in which guessing probabilities across conditions are constrained to be proportional to uncertainty. This relationship between guessing probability and uncertainty was even better explained by an alternate model, not normally used for perceptual decisions: uncertainty-guided exploration. This model, well-known in reinforcement learning, balances exploration and exploitation. Surprisingly, the model favored by BIC was a softmax model of exploration with optimal perceptual inference, whose 'exploratoriness' was proportional to posterior uncertainty. To evaluate how putative decision-making structures drive uncertainty-dependent lapses, we inactivated (muscimol) secondary motor cortex (FOF) and posterior striatum (pStr). FOF inactivation increased guessing probability, suggesting a role in biasing exploratory behavior; pStr inactivation impaired the uncertainty dependence and optimality, suggesting a role in processing uncertainty. In summary, we argue that uncertainty-dependent exploration, not stimulus-independent coin flips, drive lapses. Inactivation experiments uncover the roles of FOF and Striatum in modulating this historically mysterious feature of perceptual decisions.

II-101. Sparse attention for long-term credit assignment

Anirudh Goyal[1,2] ANIRUDHGOYAL9119@GMAIL.COM
Nan Rosemary Ke[1,2] ROSEMARY.NAN.KE@GMAIL.COM
Olexa Bilaniuk[1,3] OBILANIU@GMAIL.COM
Jonathan Binas[1,3] JBINAS@GMAIL.COM
Laurent Charlin[1,4] LCHARLIN@GMAIL.COM
Christopher Pal[1,2] CHRIS.J.PAL@GMAIL.COM
Yoshua Bengio[1,3] YOSHUA.UMONTREAL@GMAIL.COM

[1]Montreal Institute for Learning Algorithms
[2]Polytechnique Montreal
[3]University of Montreal
[4]HEC Montreal

Learning long-term dependencies in extended temporal sequences requires credit assignment to events far back in the past. The most common method for training recurrent neural networks, backpropagation through time (BPTT), suffers from vanishing and diverging gradients, and is computationally expensive when used with long sequences, rendering it non-ideal for such tasks. Importantly, biological brains are unlikely to perform detailed reverse replay of long sequences of internal states, as required by BPTT, but rather seem to capture causal relationships among events and use only the relevant events to perform credit assignment. Based on this principle, we design a novel training method, mitigating the mentioned problems with long-term credit assignment. In particular, the RNN to be trained is augmented with an attention network, which simultaneously learns to do instantaneous credit assignment, and is used to skip to relevant past states in the backtracking phase, without having to replay irrelevant intermediate states. We demonstrate in simulations that our method outperforms regular BPTT in tasks involving particularly long-term dependencies.

II-102. Experience-dependent formation of a perceptual category for maternal behavior

Jennifer Schiavo JK.SCHIAVO@GMAIL.COM
Robert Froemke ROBERT.FROEMKE@MED.NYU.EDU

New York University

Learning requires generalization from exemplars to enable animals to respond reliably to variable stimuli with the same behavioral significance. Categorization of sounds enables generalization to novel stimuli, such as communicative vocalizations that vary along a continuum of sound features. While the neural correlates of categorization are well studied, how animals form categories in order to generalize to novel stimuli is poorly understood. Here, we take advantage of a maternal behavior in mice, pup retrieval, in which mothers retrieve nest-isolated pups based on distress ultrasonic vocalizations (USVs). Although these calls are variable across individuals, the significance is the same: to signal the mother. Mothers therefore categorize pup USVs to generalize over a range of calls and reliably retrieve pups. Pup-naive virgins typically do not retrieve pups, but begin to respond following maternal experience. This behavior therefore serves as a model to 1) Examine how a naturally variable stimulus is reliably encoded in the auditory cortex and 2) Assess the experience-dependent plasticity underlying category formation in naive females as these calls gain relevance. Using a Y-maze, we assessed the perceptual boundary for distress USVs. We confirmed that dams approach speakers playing USVs, and we found that mothers categorize USVs based on inter-motif duration. Using in vivo two-photon calcium imaging, we found that USV-responsive excitatory neurons in retrieving, but not naive, virgins respond invariantly to calls morphed in the temporal domain. In addition, there is a mismatch in the width of temporal tuning between excitatory and inhibitory populations in naive, but not retrieving, virgins. Finally, we tracked these populations during learning and found that the excitatory and inhibitory responses match 24 hours post-retrieval onset. These results demonstrate that the ability of retrieving females to generalize over a range of vocalizations may be based on invariant responses to USVs in core auditory cortex.

II-103. Local and long range patterns of neural coordination in cortex

Juan Alvaro Gallego[1] GALLEGO.JUANALVARO@GMAIL.COM
Carsen Stringer[2] CARSEN.STRINGER@GMAIL.COM
Michalis Michaelos[2] MICHALIS.MICHAELOS@GMAIL.COM
Marius Pachitariu[2] MARIUS10P@GMAIL.COM

[1]Spanish National Research Council
[2]HHMI Janelia Research Campus

The dynamics of networks of neurons gives rise to perception and action. These networks coordinate through local mechanisms, such as synaptic connectivity, but also global mechanisms such as long-range connectivity and neuromodulation. What are the functional signatures of local and global patterns of coordination? To answer this question, we recorded from populations of ~10,000 neurons distributed over an area of 5mm diameter in mouse dorsal cortex, using a 2-photon random access microscope (“mesoscope”) while mice were locomoting or were stationary on an air-floating ball. We found that the average pairwise correlations were not strongly dependent on relative distance between cell somas. Instead, the population dynamics were dominated by multi-dimensional global coordination patterns, which allowed the activity of each neuron to be predicted from that of neurons recorded much further away. However, in addition to these global coordinated patterns, we also observed that strongly-correlated pairs (< 1 % of all pairs) tended to be significantly more co-localized (~5x) than expected by chance. To quantify the effect of shared coordination, we predicted each neuron from all others (peer prediction). We found that ~256 principal components were sufficient to saturate peer prediction accuracy. However, this accuracy was nearly matched by predicting from the top 64 most strongly correlated pairs. A model that combines both predictors achieved best performance with only 32 principal components and 32 most correlated pairs. Our results suggest a model of the structure of neural activity in which dense, global inputs are distributed over large areas of cortex, while sparse inputs are shared locally in a highly specific manner.

II-104. Hierarchical recurrent models reveal latent states of neural activity in C. elegans

Scott Linderman[1] SCOTT.LINDERMAN@GMAIL.COM
Annika [2,3] ANNIKA.NICHOLS@IMP.AC.AT
David Blei[1] DAVID.BLEI@COLUMBIA.EDU
Manuel Zimmer[2,3] MANUEL.ZIMMER@IMP.AC.AT
Liam Paninski[1] LIAM@STAT.COLUMBIA.EDU

[1]Columbia University
[2]Research Institute of Molecular Pathology
[3]Vienna Biocenter

Recent advances in neural recording technologies have enabled simultaneous measurements of the majority of head ganglia neurons in both immobilized and freely-behaving C. elegans. The dynamics of neural activity shed light on how C. elegans processes sensory information and generates motor activity. To understand these dynamics, we develop recurrent switching linear dynamical systems, probabilistic models that decompose complex time-series into segments with simple, linear dynamics. We incorporate these models into a robust, hierarchical framework for combining information across whole-brain recordings of many worms. Using this framework, we reveal latent states of population neural activity, along with the discrete behavioral modes that drive dynamics in this latent state space. We find stochastic transition patterns between these discrete behavioral modes, and we see that transition probabilities are determined by a combination of current brain state and environmental cues. In addition to quantifying neural dynamics, this probabilistic framework aids in neural identification—currently a laborious, manual task—and reveals clusters of neurons that are similarly tuned in latent state space. Finally, we find a significant overlap between our inferred modes and the manually-labeled modes of Kato et al. and Nichols et al., which were shown to correspond to different behaviors, like forward crawling, reversals, and turns. Our methods automatically discover and quantify these behaviorally-meaningful states directly from neural activity, yielding powerful new tools for neuro-behavioral analysis.

II-105. Place field translocation by bidirectional behavioral time-scale synaptic plasticity

Aaron Milstein[1] AARONMIL@STANFORD.EDU
Katie Bittner[2] BITKAT@GMAIL.COM
Christine Grienberger[2] GRIENBERGERC@JANELIA.HHMI.ORG
Ivan Soltesz[1] ISOLTESZ@STANFORD.EDU
Sandro Romani[2] SANDRO.ROMANI@GMAIL.COM
Jeff Magee[2] MAGEEJ@JANELIA.HHMI.ORG

[1]Stanford University
[2]HHMI Janelia Research Campus

We recently reported a novel form of synaptic plasticity that allows a single, brief (~100 ms – 1 s) event, called a dendritic plateau potential, to rapidly and potently potentiate any synaptic input to a hippocampal CA1 pyramidal neuron that is active within a prolonged (~5 s) time window surrounding the dendritic event. We showed that this behavioral time-scale synaptic plasticity (BTSP) can completely change the spatial selectivity of a hippocampal neuron in a single shot, converting spatially nonselective cells into “place cells.” This constitutes an enhanced representation of the sequence of sensory and behavioral events that preceded and followed the onset of the dendritic event. We now report in vivo electrophysiological and behavioral data and computational modeling results demonstrating that BTSP is bidirectional and state-dependent. If a dendritic plateau potential occurs in a CA1 neuron that already expresses a place field at some location along a linear trajectory, the neuron shifts its preferred firing location in the direction of the location in space where the plateau occurred. Whether the synaptic strength of a given input is increased or decreased depends on the time delay between the dendritic plasticity

event and the synaptic activation, as well as the initial strength of the input. Synaptic strength is widely accepted to scale with the number of glutamate receptors incorporated into each synapse. We are able to account for the bidirectional, state-dependent nature of BTSP with a simple kinetic model of glutamate receptor trafficking (Figure 2). Given that dendritic excitability and the probability of dendritic plateau potential initiation is under control by dendritic inhibition and various neuromodulators, we speculate that when an animal encounters reward, this learning rule is engaged by a disinhibitory circuit mechanism to mediate the translocation of hippocampal place fields towards the rewarded location.

II-106. Exponentially decaying temporal integration of stimulus history in the primary auditory cortex

Monzilur Rahman MONZILUR.RAHMAN@KEBLE.OX.AC.UK
Ben Willmore BENJAMIN.WILLMORE@DPAG.OX.AC.UK
Andrew King ANDREW.KING@DPAG.OX.AC.UK
Nicol Harper NICOL.HARPER@DPAG.OX.AC.UK

University of Oxford

For a linear non-linear model (LN) to achieve a high level of performance (normalized correlation coefficient, CCnorm 0.71; Hsu et al. 2004 Network: Comput Neural Syst; Schoppe et al. 2016 Front Comput Neurosci) in predicting single unit neural responses to natural sounds in ferret primary auditory cortex, we found that it is necessary to include stimulus history going back up to about 200 ms in the past in the spectrotemporal receptive field (STRF). Given this finding, we asked how much of this dependence on stimulus history can be explained by certain simple dynamical aspects of neurons. We constructed a neural-network-like model whose output is the weighted sum of the response of multiple units, each unit being modified by a dynamic firing-rate equation. The dynamic aspect low-pass filters the unit's response (by convolving with an exponential decay impulse response) providing a simple exponentially decaying memory governed by a time constant individual to each unit. This integrative characteristic can be related to the capacitance and resistance of neural membranes (Dayan & Abbott 2001 Theoretical Neuroscience). The results of our work show that this dynamic-network (DNet) model, when fitted to the neural data using component STRFs that are of only 25 ms in duration, can achieve prediction performance on a held-out dataset (CCnorm 0.70) comparable to the best performing LN model. When the network model with 25 ms STRFs is fitted without the dynamic aspect, its predictive performance is substantially impaired. Also, when the units in the dynamic network that have long time constants are silenced, the predictive performance of the model is affected significantly. These findings suggest that membrane time constants and/or other simple exponentially-decaying memory processes may underlie much of the dependence of the neural responses on stimulus history beyond 25 ms.

II-107. The songbird VTA integrates opponent evaluative signals for vocal learning

Matthew Kearney[1] MATTHEW.GENE.KEARNEY@GMAIL.COM
Erin Hisey[1] EEHISEY@GMAIL.COM
Timothy Warren[2] TIMLWARREN@GMAIL.COM
Richard Mooney[1] MOONEY@NEURO.DUKE.EDU

[1]Duke University
[2]University of Oregon

Animals must evaluate the consequences of their actions to adaptively guide future behaviors, negatively reinforcing actions with adverse behavioral outcomes and positively reinforcing those that lead to beneficial results. Studies in mammals have implicated the Ventral Tegmental Area (VTA) to Basal Ganglia (BG) projection as a

driver of reinforcement learning, advancing the VTA as a site for integration of evaluative signals. However, dissociating the specific contribution of evaluative inputs into the VTA has proven challenging as the VTA receives a complex suite of inputs that help to reinforce many different types of behaviors. To simplify this task, we turned to the songbird which contains a specialized VTA ’ BG circuit that enables song learning. Using closed loop optogenetic methods, we find that two song system inputs into the VTA, the ventral intermediate arcopallium (Aiv) and the Ventral Pallidum (VP), play opposing roles in vocal learning. Specifically, pitch-contingent, i.e. stimulating on low or high pitch syllable renditions, optogentic activation of Aiv-VTA terminals makes the bird subsequently less likely to sing targeted variants, while pitch-contingent stimulation of VP-VTA terminals makes the bird subsequently more likely to produce variants paired with stimulation. Song learning requires exploration, mediated by premotor variability signals, and performance evaluation through auditory feedback. To explore which of these processes different components of the BG circuitry contribute to, we used pitch-contingent noise to drive learning in a target syllable while optogenetically disrupting activity at different nodes in the BG circuit during either the premotor or auditory feedback window. We find that interfering with Aiv-VTA activity during the auditory feedback but not premotor window blocks learning, consistent with an evaluative role. Conversely, in LMAN, a premotor input to the BG that contributes to vocal variability, learning was blocked when we disrupted activity during the premotor but not auditory feedback window.

II-108. Amygdala-TRN projections amplify tone-evoked activity in auditory thalamus and cortex.

Solymar Rolon-Martinez ROLONMAR@PENNMEDICINE.UPENN.EDU
Mark Aizenberg MARK.EISENBERG@GMAIL.COM
Maria Geffen MGEFFEN@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Many forms of behavior require selective amplification of neuronal representation of relevant signals. Here we identify a novel pathway between the basolateral amygdala (BLA), an emotional learning center in the mouse brain, and the inhibitory nucleus of the thalamus (TRN), and demonstrate that activation of this pathway amplifies sound-evoked activity in the central auditory pathway. We stimulated BLA using channelrhodopsin (ChR2) with a laser via implanted optic cannulas, while recording neuronal activity in the auditory cortex (AC) in response to a presentation of random tone sequences in awake, head-fixed mice. Optogenetic activation of the BLA suppressed spontaneous activity (Fig.1C, paired t-test, p=0.0007), while amplifying tone-evoked response magnitude in AC (Fig.1D, paired t-test, p=8.5e-5). Inspection of fluorescence following retrobead injections in BLA revealed direct projections from BLA to TRN. These projections were further confirmed by retrograde labeling of neurons in the BLA using a CAV-2 virus in TRN. We next directly activated projections from the BLA to TRN by repeating the initial experiment, but positioning the optic cannula over TRN. We found that there was a significant suppression of spontaneous activity (Fig. 2C, paired t-test, p=0.003), and a significant increase in tone-evoked responses in AC (Fig. 2D, paired t-test, p=3.9e-8). We found that activation of the BLA projections to TRN also led to inhibition of spontaneous activity (Fig. 3C, paired t-test, p=4.3e-9) and an increase in tone-evoked responses in auditory thalamus (Medial Geniculate Body, MGB) (Fig. 3D, paired t-test, p=3.4e-7), consistent with the hypothesis that the changes in AC responses with BLA activation are a result of projections from BLA to TRN via MGB. These results demonstrate a novel circuit mechanism for amplification of sensory representation of behaviorally relevant signals and provide a potential target for treatment of neuropsychological disorders, in which emotional control of sensory processing is disrupted.

II-109. Motor cortical control of vigor but not reach direction in freely moving mice

Teja Pratap Bollu TB387@CORNELL.EDU
Samuel Whitehead SCW97@CORNELL.EDU
Itai Cohen IC64@CORNELL.EDU
Jesse Goldberg JHG285@CORNELL.EDU

Cornell University

A mainstay of motor neuroscience is the center-out reach task, in which primates hold a hand still at a center position prior to reaching out to various target locations. The task is simple yet provides millisecond timescale information on movement direction, velocity, timing, and variability, as well as immobility during 'hold-still' periods. Representation of these kinematic parameters in motor cortex enables the decoding of ongoing movements and brain machine interfaces. Yet, causal roles of neural circuits controlling movement remain poorly understood, in part because reversible trial-by-trial silencing of precise cell types and pathways remains difficult in primates. Genetic tools in mice provide opportunities to study how specific motor circuits control the forelimb. Rodents naturally use forelimbs for appetitive behaviors and can learn to reach for food, press levers and move manipulanda. However, to our knowledge it remains unknown if mice can learn a center-out reach task in which they first actively 'hold-still' and then reach to target directions. We trained mice in a center-out reach task optimized for high-throughput dissection of motor circuits. We designed ultra-low torque touch-sensing joysticks that resolve mouse forelimb kinematics with micron-millisecond spatiotemporal resolution. These joysticks are then integrated into computer-controlled, rack-mountable homecages that automate behavioral training and closed-loop optogenetics. We use this system to show that mice can learn, with no human handling, a direction specific center-out reach task. Next, we inactivate caudal (CFA) and rostral forelimb areas (RFA) of motor cortex and show that while the probability of reaching out is reduced by CFA contralateral inactivation, the direction of the reach is unimpaired. Finally, we specify the effect of inactivations by decomposing the trajectories into constituent sub-movements. We find that the duration of sub-movements is unaffected but the peak speed and distance are reduced across all sub-movements, a classic signature of reduced vigor.

II-110. The hippocampus provides an internal source of evidence for value-based decisions

Akram Bakkour AB4096@COLUMBIA.EDU
Ariel Zylberberg AZ2368@COLUMBIA.EDU
Michael Shadlen SHADLEN@COLUMBIA.EDU
Daphna Shohamy DS2619@COLUMBIA.EDU

Columbia University

The speed and accuracy of many decisions suggest they arise through a process that involves accumulation of independent samples of evidence until satisfying a threshold. For perceptual decisions, the samples arrive from sensory input. What gives rise to the sequence of samples in value-based decisions? We theorized that samples of evidence bearing on value preference are derived through a process that involves episodic memory retrieval and prospection and depends on the hippocampus. To test this, we scanned 30 human participants with fMRI while they performed a value-based decision task (snack preference) and, for comparison, a perceptual decision task (dynamic random dot color dominance). Choice and RT functions were consistent with sequential sampling models for both tasks, but fMRI revealed differences in the correlates of RT: Value-based RT was correlated with significantly greater BOLD activity in the hippocampus. Moreover, a separate localizer indicated that this effect in the hippocampus overlapped with voxels that are active during memory retrieval in the same participants. We reasoned that if memory contributes to constructing preference, then the items may undergo revaluation as a consequence of the decision. We tested this using an algorithm that supposes the chosen and unchosen items change value by +/- delta. The revised values better accounted for choices and RT than the original values, and

were more strongly correlated with BOLD activity in the ventromedial prefrontal cortex, a brain region implicated in valuation, suggesting that they provide a more veridical representation of value than the ones ascertained at the beginning of the experiment. Additional analyses support the idea that value is partially constructed during value-based decisions. These findings may help explain: 1) why value-based decisions take time; 2) stochasticity in choice behavior and failures of transitivity; and 3) alterations of value due to choice as an alternative to cognitive dissonance theory.

II-111. Shared neuronal variability accounts for behavioral variability in count discrimination tasks

Mikio Aoi[1] MIKIOAOI@GMAIL.COM
Ben Scott[1] BBSCOTT@PRINCETON.EDU
Christine Constantinople[1] CMC9@PRINCETON.EDU
Carlos Brody[2,1] BRODY@PRINCETON.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Howard Hughes Medical Institute

The psychophysics of numerosity, or number sense, is the study of how discrete numbers of stimuli are perceived. Empirical studies have shown that the number sense conforms to Weber's law of perceptual discriminability. This means that the standard deviation of the perceptual noise on the perceived number of objects scales linearly with the number of objects, indicating perceptual noise is not independent of the percept. This property, observed in both humans and non-human animals, is commonly known as "scalar variability." The generation of large scale behavioral data sets ($\approx 10\hat{}6$ trials) from animals trained to perform numerosity tasks, allows us to determine whether or not judgements about count stimuli coincide with deviations from precise scalar variability (Scott & Constantinople, et al. Elife, 2015). Here we propose a probabilistic behavioral model to account for variability in count discrimination tasks. Our model, based on a model of shared stochastic gain in neuronal populations, can closely fit the psychometric function and the inferred uncertainty of count perception from behaving animals. We compare the performance of our model to a previously proposed 16-parameter model based on signal detection theory and show that our model fits data better than the previous model with just two parameters. This work draws a direct connection between neurophysiology and behavior by demonstrating that perceptual psychophysics can be explained by the statistical properties of neuronal populations.

II-112. Brain Modeling ToolKit (BMTK): an open-source package for multiscale modeling of brain circuits

Yazan Billeh YAZANB@ALLENINSTITUTE.ORG
Michael Buice MICHAELBU@ALLENINSTITUTE.ORG
Nicholas Cain NICHOLASC@ALLENINSTITUTE.ORG
Stefan Mihalas STEFANM@ALLENINSTITUTE.ORG
Sergey L. Gratiy SERGEYG@ALLENINSTITUTE.ORG
Kael Dai KAELD@ALLENINSTITUTE.ORG
David Feng DAVIDF@ALLENINSTITUTE.ORG
Nathan W. Gouwens NATHANG@ALLENINSTITUTE.ORG
Ramakrishnan Iyer RAMI@ALLENINSTITUTE.ORG
Jung Hoon Lee JUNGL@ALLENINSTITUTE.ORG
Christof Koch CHRISTOFK@ALLENINSTITUTE.ORG
Anton Arkhipov ANTONA@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Neuroscientists today model the nervous system at multiple levels of abstraction. The different levels of resolution have their own benefits and drawbacks and the optimal choice is dependent on the question asked by the scientists. Nevertheless, it is at times necessary to bridge these network models. To this end, we have developed the Brain Modeling ToolKit (BMTK, github.com/AllenInstitute/bmtk) that, in a single package, allows neuroscientists to model brain networks at five levels of resolution: as biophysically detailed neuronal networks, point neuronal networks, firing rate network models, filter networks, and machine intelligence networks.

The package has a modular design that separates the building, simulation, and analysis steps; all implemented as Python application programming interfaces (APIs). The Builder component of BMTK enables construction of highly complex, large-scale networks at different levels of resolution, all using a similar graph structure that translates across the levels. The modeling platform builds on this same efficient graph representation format and provides scaling of the code to handle models involving 100's of thousands of cells and 10's of millions of connections. It is implemented as a wrapper to a number of commonly-used software packages. For the most detailed resolution of biophysically detailed models, the developed API uses NEURON, for point-neuron models it uses NEST, for rate models—DiPDE, and for machine intelligence—TensorFlow. For the filter models, we have developed code that allows users to create filters that convert movies and images into firing rates and/or spike trains. Together, these tools provide a convenient interface for modeling workflows, enabling standardized and easy ways for model sharing, data exchange, and comparison of results across levels of resolution.

BMTK is developed by the Allen Institute for Brain Science and is freely available at github.com/AllenInstitute/bmtk.

II-113. Quantitative assessment of long-term, multi-session recordings from area MT of the awake marmoset

Aaron J Levi ALEVI@UTEXAS.EDU
Hayden C. Carney HAYDENCCARNEY@GMAIL.COM
Jens-Oliver Muthmann OLLIMUH@UTEXAS.EDU
Alex Huk HUK@UTEXAS.EDU

The University of Texas at Austin

The marmoset has drawn attention as a complementary nonhuman primate model system for computationally-oriented visual neuroscience due to their smooth (lissencephalic) cortex, availing virtually all cortical areas to the use of chronically implanted electrode arrays. However, despite applications in the anesthetized marmoset, no reports to our knowledge have provided quantitative characterizations of long-term recordings from chronically implanted arrays in the awake marmoset. Of additional concern is the tendency of marmosets to perform far fewer trials per session compared to macaques, resulting in correspondingly smaller amounts of neural data per experiment—and potentially undermining the advantages of the species. Here, we report successful long-term recordings using a novel “3D” array implanted in area MT of an awake, behaving marmoset. This array complements 2D (“Utah”) arrays (the standard chronic arrays used in macaque) by providing recordings from different depths within the gray matter. Their slow (non-ballistic) insertion also appears better suited to the delicate and small marmoset brain. These arrays provide stable recordings over several months, allowing for combinations of experiments over time frames far longer than an individual recording session. By concatenating recording sessions and spike sorting the resulting combined file using KiloSort, we found that tuned single- and multi-unit clusters are stable over multiple successive sessions, where waveform shape, spatial receptive fields, inter-spike interval distributions, and direction preference persisted over time frames on order of weeks to months. The ability to splice together datasets from a large, stable population over several days can therefore circumvent the behavioral limitations of marmosets. This approach enables long-term study of individual neurons and populations in the awake primate, providing statistical power for detailed quantitative analyses and opening prospects for studying the same neural population over long time frames.

II-114. Function and dysfunction: VIP interneuron contributions to state-dependent cortical computation

Katie Ferguson KATIE.FERGUSON@YALE.EDU
Renata Batista-Brito RENATA.BRITO@YALE.EDU
Jessica Cardin JESS.CARDIN@YALE.EDU

Yale University

Inhibitory interneurons are key regulators of cortical processing, but it remains unclear how state-dependent modulation of distinct interneuron populations affect sensory circuit dynamics. GABAergic interneurons expressing vasoactive intestinal peptide (VIP-INs) are strongly activated by arousal, regulate cortical excitability by disinhibiting local interneuron populations, and receive serotonergic and cholinergic afferents—neuromodulatory systems affected during distinct behavioral states. Inhibitory-inhibitory interactions between VIP-INs and their major synaptic target, somatostatin expressing interneurons (SOM-INs), have been suggested to play a critical role in the regulation of cortical circuit activity. VIP-INs are thus uniquely situated to be key contributors to cortical function and dysfunction. However, their role in cortical development has not been investigated and little is known about interneuron-interneuron interactions in vivo, making it unclear how VIP-INs affect behavioral state-dependent regulation of sensory processing.

To test the role of VIP-INs in cortical development and dysfunction, we used a conditional deletion model to remove ErbB4, a key developmental gene for interneurons, from VIP-INs to decrease their excitatory input gain. To assess the contribution of VIP-IN deficits in arousal-mediated cortical processing, we expressed the genetic calcium indicator gCaMP6 in either a large population of VIP-INs, or, using intersectional genetic tools, in SOM-INs in the mouse primary visual cortex (V1). We used large-scale two-photon imaging and high-throughput analyses to determine VIP-IN and SOM-IN activity in V1 cortex of awake mice across behavioral states (e.g., quiet wakefulness, running).

In contrast to controls, developmentally disrupted mutants exhibited a loss of cortical response to behavioral arousal. The state-dependent modulation of VIP-IN activity was significantly decreased in mutants. Although SOM-INs in V1 show dramatic responses to visual stimuli in healthy mice, these responses were eliminated in VIP-dysregulated mutants. Instead, locomotion-dependent activity was uncovered in SOM-INs. This suggests an unanticipated role for VIP-INs in regulating large-scale changes to sensory processing across behavioral states.

II-115. Non-responsive frontal and sensory cortical cells encode behavioral variables via consensus-building

Michele Insanally[1] MNI1@NYU.EDU
Ioana Carcea[1] CARCEA.IOANA@GMAIL.COM
Rachel Field[1] RACHEL.FIELD@NYUMC.ORG
Badr Albanna[2] BALBANNA@FORDHAM.EDU
Robert Froemke[1] ROBERT.FROEMKE@MED.NYU.EDU

[1]New York University
[2]Fordham University

Spike trains recorded from the cortex of behaving animals can be complex, highly variable from trial-to-trial, and therefore challenging to interpret. A fraction of cells exhibit trial-averaged responses with obvious task-related features such as pure tone frequency tuning in auditory cortex. However, a substantial number of cells (including cells in primary sensory cortex) do not appear to fire in a task-related manner1 and are often neglected from analysis. Even classically-responsive cells lose their stimulus representation during task-engagement without impairing behavior2,3. These results suggest that nominally non-responsive cells may play an underappreciated role in sensory processing and cognition. At Cosyne 2017, we presented a novel single-trial, spike-timing-based analysis to evaluate whether the single-unit activity recorded from auditory and frontal cortex encode task variables in behaving rats. Here we expand our investigation to decoding population activity and demonstrate: 1) Nominally non-responsive cells reveal hidden task information complementary to responsive cells. The activity of cells that seem unresponsive when trial-averaged often encode additional task-relevant information at levels comparable to responsive cells. 2) Stimulus information is more prevalent and pervasive in frontal cortical ensembles. When tones become behaviorally significant, stimulus information is encoded more accurately in frontal cortex suggesting it is critical for extracting task-relevant stimulus information. Furthermore, stimulus decoding in frontal cortex improves dramatically when using small ensembles demonstrating this information is ubiquitous. 3) Ensemble 'consensus-building' dynamics underlie hidden task information. On correct trials only, ensemble members coordinate the behavioral meaning of their spiking activity moment-to-moment over the course of the trial to reach 'consensus' on a common representation of task-variables. These hidden dynamics demonstrate how non-responsive frontal cortical cells are modulated by behavioral stimuli and all ensembles coordinate leading to behavioral response. Intriguingly, the shared dynamics of nominally non-responsive ensembles suggests their contribution to the functional coordination of these cortical regions.

II-116. Dynamical structure of socio-vocal network in marmoset monkeys

Daniel Takahashi[1] TAKAHASHIYD@GMAIL.COM
Ahmed El Hady[1] AHADY@PRINCETON.EDU
Gabriel Montaldo[2,3] GABRIEL.MONTALDO@NERF.BE
Alan Urban[2,3] ALAN.URBAN@NERF.BE
Asif A. Ghazanfar[1] ASIFG@PRINCETON.EDU

[1]Princeton University
[2]Neuro-Electronics Research Flanders
[3]KU Leuven

Vocal communication is the quintessential form of social interaction. Humans and other animals coordinate their social behaviors by producing and perceiving distinct vocalizations. Brain networks related to vocal communication include areas at the intersection of social behavior and vocal production-perception networks. Recent studies of primate vocal communication focused on lateral cortical regions, despite the fact that medial cortical and subcortical areas constitute the main vocal production and social behavior network (SBN). Hence, we aim to unravel the brain-wide network underlying social communication focusing on the role played by medial cortical and subcortical areas. We use as our model the marmoset monkey, a highly vocal New World species. To image large-scale

neural activity, we use functional ultrasound imaging which has a large spatial coverage and high spatio-temporal resolution. Furthermore, we built a stochastic dynamical systems model of vocal behavior that interacts with the marmoset in a closed-loop to fully control the vocal interaction and make quantitative predictions about brain dynamics during communication. We first show the existence of a medial brain system at the intersection of vocal production-perception and SBN; we call it the socio-vocal network (SVN). These areas differentially respond to affiliative vocalizations—contact, trillphee, and trill calls—produced in different contexts, exhibiting the highest and quickest response to contact calls. Given that the contact calls reflect the highest arousal state of the vocalizing animal, this is consistent with the hypothesis that SVN is related to the monitoring of others motivational state through vocalization. Second, through a closed-loop interaction between the computational model and a marmoset, together with large-scale functional imaging, we found that the marmoset anterior cingulate cortex (which is part of SVN) and the model's "SVN" are entrained. These results demonstrate what the SVN encompasses and its roles in vocal communication.

II-117. A connectome derived hexagonal lattice convolutional network model of the fruit fly visual system accurately predicts direction selectivity

Srini Turaga[1] TURAGAS@JANELIA.HHMI.ORG
Fabian Tschopp[2] TSCHOPFA@STUDENT.ETHZ.CH

[1]HHMI Janelia Research Campus
[2]University of Zurich and ETH Zurich

What can we learn from a connectome? An electron microscopic reconstruction of a circuit provides rich detail regarding the 3d morphologies of individual neurons, and the locations and counts of synapses between pairs of neurons. On the one hand the connectome contains too much information, since it is yet unclear how important the precise morphological details are to the functioning of the neural circuit. On the other hand there is too little information, since many important biophysical details regarding the input-output transformations and the dynamics of neurons and synapses cannot be gleaned from such reconstructions.

In this work, we show how a connectome, when combined with a hypothesis about the function of the circuit, can be used to infer missing parameters and perform error correction. 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 convolutional network was trained using backpropagation through time to perform object tracking in natural scene videos. Networks initialized with weights from connectome reconstructions automatically discovered well-known orientation and direction selectivity properties in T4 neurons and their inputs, while networks initialized at random did not.

Our results suggests that drosophila visual system can be well approximated by a simple point neuron threshold linear network, and that synapse counts can be a good proxy for synaptic strength. Our work is the first demonstration, that knowledge of the connectome can enable in silico predictions of the functional properties of individual neurons in a circuit, leading to an understanding of circuit function from structure alone.

II-118. Neural network trained with supervision represents uncertainty by non-linear moments

Li Wenliang[1,2] KEVINWLI@OUTLOOK.COM
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]Gatsby Computational Neuroscience Unit

Making optimal inferences about the state of the world from noisy sensory information requires that accurate and flexible representations of the concomitant uncertainty be learnt. How might this happen? It is obvious that supervised learning from noisy inputs must retain some information about uncertainty to perform optimally (e.g. Orhan and Ma, 2017), but it is unclear whether such a representation will be flexible or generic enough to underpin general probabilistic computation. Here, we analyse the representation of uncertainty that arises through supervised learning in a network tuned to propagate probabilistic messages using the recently proposed distributed distributional coding (DDC) scheme. Following previous work (Sahani and Dayan, 2003; Zemel et al, 1998), the DDC assumes that neurons represent uncertainty through the expectations of pre-specified basis functions under the encoded distribution. DDCs can be used to generate state-of-the-art performance in unsupervised learning of intractable models using a biologically plausible representation of uncertainty. We trained recurrent neural networks (RNNs) to estimate the posterior mean of a non-linear dynamical system without explicitly enforcing a DDC-like representation. Nonetheless, the RNN in which propagation was consistent with DDC message passing performed better than other networks, and its hidden units preserved more information about the posterior variance. Indeed, we found that activities in the hidden layer of this RNN could be interpreted as posterior expectations of functions over the latent variables; these functions did well not only in predicting the hidden activities, but also in reconstructing the posterior distributions. Thus, we conclude that flexible DDC-like codes for uncertainty are learnt naturally within networks of the suitable architecture.

II-119. Central thalamic stimulation restores awake behavior and cortical dynamics in anesthetized macaques

Jacob Donoghue JACOB.A.DONOGHUE@GMAIL.COM
Emery Brown ENBROWN1@MIT.EDU
Earl Miller EKMILLER@MIT.EDU

Massachusetts Institute of Technology

General anesthesia is often considered the greatest discovery in medical history, yet the neural dynamics producing this brain state are remarkably not understood. A convergence of theoretical and electrophysiological findings suggests that a conscious state of wakefulness and awareness may emerge from the integration of information in finely tuned cortical networks coordinated by connections with the thalamus. We developed a non-human primate model of general anesthesia to study how populations of neurons across these networks mediate changes in conscious states. We simultaneously recorded spikes and local field potentials from chronically-implanted multi-electrode arrays in prefrontal, posterior parietal, and auditory cortex and from laminar probes within the central thalamus during the administration of the GABAergic anesthetic propofol. We discovered that cortical networks fragment in the unconscious state in unpredictable ways- hyperconnectivity and functional disconnection simultaneously emerge across distinct brain regions via changes in spiking and oscillatory synchrony. Furthermore, we show that these neural population spike patterns preclude normal sensory stimuli encoding differently across hierarchical networks. We sought to establish a causal role for the thalamus in maintaining this unconscious state by activating the diffusely projecting intralaminar nuclei with high-frequency electrical stimulation. Incredibly, deep brain stimulation (DBS) immediately and continuously reversed general anesthesia. DBS elicited a state of behavioral wakefulness characterized by eye opening, air-puff responses, and restored limb movement, despite continued anesthetic infusion. Thalamic activation produced an awake-like cortical state despite overwhelming GABAergic inhibition, eliminating slow oscillatory activity across brain regions and inducing a shift to higher frequency rhythms with awake-like spiking dynamics. Interestingly, behavioral and cortical “re-awakening” regularly outlasted DBS, with loss-of-consciousness reoccurring sometimes minutes after cessation of stimulation. Together, these results shed light on the network architecture and dynamical systems required to support cortical processing, cognition, and consciousness.

II-120. Connecting feature-based and dynamics-based classifications in V2

Ryan Rowekamp RROWEKAMP@SALK.EDU
Tatyana Sharpee SHARPEE@SALK.EDU

Salk Institute for Biological Studies

Previous studies of visual area V2 have identified multiple sub-populations of neurons based on their response properties. Neural clusters were identified based on the properties of their spatial feature selectivity, such as 'ultra-long' and 'complex-shaped' classes, based on their temporal integration properties, and on the strength of their suppressive subunits. Here we show that these classification schemes actually reflect just two classes because neurons with transient temporal dynamics have more homogeneous spatial integration properties . These results were obtained by fitting quadratic convolutional models to the responses of macaque V2 neurons to natural movie stimuli and analyzed the parameters of these models in order to understand their spatial feature selectivity and temporal weighting. The convolutional model included pooling weights in time, and these weights were analyzed to derive neural dynamical properties. The spatial feature selectivity was characterized by fitting the quadratic kernel of the model using combination of Gabor features. The reconstructed Gabor features formed quadrature pairs (similar to observed properties of V1 complex cells). We also observed that the suppressive features tended to be locally orthogonal to the excitatory features. This organization increased the sparseness of the simulated responses relative to randomly oriented suppressive features. Combined with the invariance provided by the convolutional nature of the model, the increase in selectivity would be useful for object recognition.

II-121. Uncovering the layer-specific role of cortical surround integration during active sensation

Evan Lyall[1] EVAN.LYALL@GMAIL.COM
Scott Pluta[2] SPLUTA@PURDUE.EDU
Hillel Adesnik[1] HADESNIK@BERKELEY.EDU

[1]University of California, Berkeley
[2]Purdue University

Mammals scan their sensors across their local environment to build up an internal representation of that environment in a process known as active sensation. The primary sensory cortex is known to be critical during such sensory processing for generating high-resolution perception. It exhibits a conserved laminar structure and serial hierarchy that integrates information and extracts features of the stimulus. We have previously shown that surround integration is key to generating higher-order features of cortical computations, such as a map of whisker space in cortical layer 2/3 (L2/3) and L5 (Pluta & Lyall et al. 2017). However, it has not yet been observed how the representation of a stimulus changes within a barrel column as it passes from layer to layer, and where the bulk of the surround integration occurs in order to generate such higher-order features. Is surround integration linear in certain layers, and nonlinear in others? Is the map of space in L5 as smooth as it is in L2/3, or is it more discretized? By combining large field of view two-photon calcium imaging (1.2mm by 1.2mm), transgenic animals that restrict the calcium indicator (GCaMP6s) to the thalamus or a single cortical layer, and a novel whisker stimulator that can activate single whiskers during natural whisking, we have begun to observe how a sensory stimulus is represented at various stages of the cortical hierarchy, and how surround information is integrated at each stage. Preliminary analyses have shown that although most whisker combinations summate sublinearly, there is a surprisingly high degree of superlinear summation in certain neurons for specific whisker combinations unique to each neuron. These results provide a constraint on the potential circuit mechanisms that produce the various cortical computations, such as feature extraction.

II-122. First spikes in visual cortex enable perceptual discrimination

Arbora Resulaj[1] ARESULAJ@GMAIL.COM
Sarah Ruediger[1] SARAH.RUEDIGERLEE@UCSF.EDU
Shawn Olsen[2] SHAWNO@ALLENINSTITUTE.ORG
Massimo Scanziani[1] MASSIMO@UCSF.EDU

[1]University of California, San Francisco
[2]Allen Institute for Brain Science

Perceptual decisions involve the sequential activation of several, hierarchically organized cortical areas. Based on the number of areas likely involved in the processing of sensory stimuli it has been hypothesized that in each area a relatively brief time window of activity may be sufficient to enable a perceptual decision. Yet, this time window has never been directly measured for any area. By determining these lower limits and analyzing neuronal activity over this time window within a given area we can reveal how the stimulus is represented within this time frame in that area. Furthermore, this time window defines the time that an area has to be active such that downstream areas can extract sufficient information to enable a perceptual decision. How this time window relates to the time window for an outside observer to extract sufficient information from that area is not clear. Here we developed a visual discrimination task in mice that necessitates visual cortex because both acute cortical silencing and permanent ablation reduces performance of the task to chance. By optogenetically silencing visual cortex at various intervals following the onset of the visual stimulus we show that the lower temporal limits of visually evoked activity for a perceptual discrimination lie within 40-80 ms. The impact on behavioral performance when silencing visual cortex during this time window is particularly sensitive to the difficulty of the task. Using extracellular electrophysiology during behavior we show that during this time window the vast majority of neurons discriminating the stimulus fire one or no spikes and less than 16% fire more than two. Furthermore, the tuning of these neurons to passively viewed stimuli explains how well they discriminate during the task. These results establish that the firing of the first visually evoked spikes in visual cortex is sufficient to enable a perceptual decision.

III-1. Task representation in the macaque posterior parietal cortex during virtual navigation

Kaushik Lakshminarasimhan[1] JKLAKSHM@BCM.EDU
Roozbeh Kiani[2] ROOZBEH@NYU.EDU
Gregory DeAngelis[3] GDEANGELIS@UR.ROCHESTER.EDU
Paul Schrater[4] SCHRATER@UMN.EDU
Xaq Pitkow[1,5] XAQ@RICE.EDU
Eric Avila[1] ERIC.AVILAOROZCO@BCM.EDU
Dora Angelaki[1,5] ANGELAKI@BCM.EDU

[1]Baylor College of Medicine
[2]New York University
[3]University of Rochester
[4]University of Minnesota Twin Cities
[5]Rice University

Much of what we know about how the brain computes comes from highly controlled tasks that use stimuli with stationary statistics and a limited set of actions (usually two). Such tasks may be inadequate to fully reveal the rich structure of neural representations and computations that mediate fluid behaviour. To understand dynamic neural processing underlying natural behaviour, we trained two macaque monkeys on a continuous-time foraging task in which they used a joystick to steer freely and catch targets in a two-dimensional virtual environment devoid of landmarks. Targets appeared briefly at a random location on the ground plane within the field of view. In order to solve the task, monkeys had to dynamically update their position estimates by integrating optic flow generated by self-motion. We implanted multi-electrode arrays to sample the activity of a large number of neurons in the

posterior parietal cortex (PPC). Fitting a generalized additive model to the neural activity revealed that a majority of neurons encoded multiple task-relevant variables ranging from the monkeys' instantaneous linear and angular velocity to more abstract, integrated variables such as distance and direction of heading. We then inferred the structure of neural interactions by extending our model to include coupling between neurons. We found that there was sparse but indiscriminate flow of information between neurons encoding different task variables, and that the coupled model provided a better account of neural responses. To understand how task variables are represented at the population level, we used canonical correlation analysis and found that the dimensionality of task-relevant neural subspace was as high as possible. Similar analyses on uncoupled and coupled model populations showed that coupling between neurons was responsible for broadening the task representation. These results demonstrate that recurrent connections in the primate PPC facilitate processing and integration of sensory inputs in dynamic environments.

III-2. Different neural landscape regulates individual differences in sensory-guided decision making

Tomoki Kurikawa[1] TOMOKI.KURIKAWA@RIKEN.JP
Takashi Handa[2] TAKASHI.HANDA@CAESAR.DE
Tomoki Fukai[1] TFUKAI@RIKEN.JP

[1]RIKEN Brain Science Institute
[2]Research Center Caesar

In the brain, decision-making is instantiated in dedicated neural circuits. However, there is considerable individual variability in decision-making behavior, particularly under uncertainty. The origins of decision variability within these conserved neural circuits are not known. Here, we elucidate underlying mechanism on individual variability in rat decision-making behavior in combination with a network model. In a sensory-guided choice task, rats were trained to associate two cues with two choices. After training, their behavior was exposed for familiar and (untrained) unfamiliar cues. On the unfamiliar stimuli, choice responses varied significantly across individual rats: some rats responded differently to different stimuli (sensitive rats), while others responded similarly to them (insensitive rats). To investigate what kind of mechanism in neural dynamics underlies this individual difference, we recorded neural activities in rat medial frontal cortex (MFC) and constructed a reservoir network model based on the anatomical structure of MFC. After training the network on familiar stimuli, we applied unfamiliar stimuli to the network and analyzed its neural dynamics. The network replicated neural trajectories and individual difference that are observed in the rats. To understand property of neural dynamics, we applied the perturbations to a particular neural state on a neural trajectory and computed the disturbance of the neural trajectories, called susceptibility. We found that the susceptibility predicts behavioral traits: a network with higher susceptibility shows more sensitive behavior, and vice versa. In addition, the susceptibility was correlated with the trial-by-trial variability of neural trajectories and then, the variability explains individual difference in response behavior. We found that trial-by-trial variability in neural trajectories in MFC can predict the difference across individual rats. A trajectory with higher (lower) susceptibility to perturbation implies shallower (deeper) landscape around the trajectory. Our study suggests that the different land scape of neural dynamics in MFC regulates individual differences in responding behavior.

III-3. Corticostriatal circuit for strategic behavior by gain control of action and reward valuation.

Yuval Baumel BAUMELYU@GMAIL.COM
Melissa Warden MRWARDEN@CORNELL.EDU

Cornell University

Optimal behavior is dependent on correct evaluation of the reward structure of the environment and action in accordance with this understanding. Unfortunately, these two processes stand in contrast since information gathering is usually at the cost of immediate earnings. To solve this predicament one could potentially control the gain of valuation in circuits that form action-outcome associations. Neurons in the anterior cingulate cortex (ACC) encode multiple parameters of the decision space as well as the current strategy (Kolling et-al, 2012; Wan et-al. 2015). The striatum participates in the formation of action-outcome associations and exhibits increased activation as behavior settles upon exploitation of the currently preferred option (Daw et-al, 2006). Given the dense projections from the ACC to the dorsomedial striatum (DMS) we sought to investigate what information is conveyed by this corticostriatal circuit during performance of a two alternative forced choice task, as behavioral strategy shifted according to internal state or in response to external changes in the environment.

To investigate the activity of DMS projecting ACC neurons (ACC-DMS) we injected Cre-inducible GCaMP6m into ACC, and a retrogradely transported adenovirus encoding Cre-recombinase (Cav2-Cre) into DMS. We imaged at cellular resolution using a GRIN lens and head mounted mini-microscope. ACC-DMS neurons integrated information about previous and current trials and exhibited action and reward related activity that was contingent on the current behavioral mode. ACC-DMS neurons that were active during exploration also responded to changes in action-outcome contingencies, and maintained stable spatial representation of the action-reward space even between different sessions spaced days apart. ACC-DMS neurons that developed robust activity at the transition to exploitive behavior signaled the preferred option and changed spatial characteristics according to changes in the preferred option. The data indicate two complementary streams of gain control of action and reward valuation in corticostriatal circuitry to enable strategic behavior.

III-4. A neural field model for size tuning and its modulation by locomotion of cell classes in mouse V1

Mario Dipoppa[1] MARIO.DIPOPPA@GMAIL.COM
Matteo Carandini[1] M.CARANDINI@UCL.AC.UK
Kenneth Harris[1] KENNETH.HARRIS@UCL.AC.UK
Adam Ranson[2] RANSON.AD@GOOGLEMAIL.COM
Michael Krumin[1] MICHAEL@CORTEXLAB.NET
Marius Pachitariu[3] MARIUS10P@GMAIL.COM

[1]University College London
[2]Cardiff University
[3]HHMI Janelia Research Campus

An outstanding question in neuroscience is to determine how properties of the cortex emerge from the interactions between different cell classes. For example, size tuning in Pyramidal (Pyr) neurons of the mouse primary visual cortex (V1) is thought to be imposed by somatostatin-expressing (Sst) interneurons (Adesnik et al., 2012). Size tuning, moreover, depends on locomotion (Ayaz et al., 2013), which can profoundly affect sensory responses in multiple neuronal classes (Fu et al., 2014; Pakan et al., 2016; Polack et al., 2013). We sought to determine a canonical circuit underlying the interplay between these phenomena in Pyr, Sst, Parvalbumin (Pvalb), and Vasoactive Intestinal Polypeptide (Vip) cells. We imaged the visual responses using a 2-photon microscope, identifying Pyr, Sst, Vip and Pvalb neurons. We measured for the first time the visual responses of a large number of cells for different combinations of stimulus size, locomotion conditions and receptive field (RF) positions. In contrast with previous literature, Sst neurons with centered RF were size tuned. Additionally, locomotion increased only responses to small stimuli in Vip neurons and only responses to large stimuli in Sst neurons. We devised a neural field model to describe quantitatively the average responses of each cell class. The model provided accurate fits to the data of all cell types, but only if we incorporated the following constraints between cell classes: a broad integration from Sst neurons to Pyr neurons; visual feedforward input that impinges not only on Pyr and Pvalb neurons but also on Sst cells; visual feedforward input that becomes stronger with locomotion. We conclude that a neural field model can capture the complex size- and locomotion-dependent responses of V1 cell classes. The synaptic connections are independent of the stimulus size and therefore the model can be generalized to

explain visual responses to stimuli of any arbitrary size.

III-5. Running reduces firing but improves coding in rodent higher-order visual cortex

Amelia Christensen[1] AMYJC@STANFORD.EDU
Jonathan Pillow[2] PILLOW@PRINCETON.EDU

[1]Stanford University
[2]Princeton University

Running profoundly alters stimulus-response properties in mouse primary visual cortex (V1), but its effects in higher-order visual cortex remain unknown. Here we systematically investigated how locomotion modulates visual responses across six visual areas and three cortical layers using the Allen Brain Institute Brain Observatory dataset. Although running has been shown to increase firing in V1, we found that it suppressed firing in higher-order visual areas. Despite this reduction in firing rate, visual responses during running could be decoded more accurately than visual responses during stationary periods. We show that this effect was not attributable to changes in noise correlations or changes in firing rate, and instead correlated with increased reliability (defined as the variance of each neuron's average response to each stimuli, divided by the overall variance of that neurons responses) of single neuron responses during running. We proposed a biophysical mechanism for this increased single neuron reliability: decreased membrane voltage fluctuations during running. Using leaky integrate and fire (LIF) simulations of individual neurons with differing membrane voltage fluctuation levels, we demonstrated that this mechanism could give rise to neurons whose firing rates decrease during running, yet whose overall response reliability increases.

III-6. Oxytocin modulates neural synchrony between the anterior cingulate cortex and amygdala for social decisions

Olga Dal Monte OLGA.DALMONTE@YALE.EDU
Nicholas Fagan NICHOLAS.FAGAN@YALE.EDU
Steve Chang STEVE.CHANG@YALE.EDU

Yale University

Accumulating evidence suggests that oxytocin (OT) is involved in regulating social behavior. However, it remains unexplored how local OT signaling within a specific brain region modulates such interactions at the neuronal level to guide social decisions. Here, we focally infused OT into the basolateral amygdala (BLA) to examine its direct effects on the neuronal coordination between BLA and the reciprocally-connected anterior cingulate gyrus (ACCg), two regions previously shown to signal social decision outcomes at the single-neuron level. We used a social reward allocation task in which an actor monkey chooses among delivering juice rewards to himself (Self), both himself and an other monkey (Both), an another monkey (Other), or a juice collection bottle (Neither). We recorded local field potentials (LFP) from ACCg and BLA simultaneously, using pairs of axial arrays, to investigate changes in the coordination following either OT or saline infusions into BLA (OT, n = 205 BLA, 205 ACCg sites; saline, n = 159 BLA, 189 ACCg sites). The actors preferred to donate juice to the other monkey (Other) over a bottle (Neither), but also preferred Self over Both, providing the contexts for examining the ACCg-BLA interaction across prosocial (Other over Neither) and antisocial (Self over Both) preferences. OT infusions blocked the natural decline in prosocial choice over time while increasing ACCg-BLA coherence in the gamma band and decreasing coherence in the beta band (relative to saline). OT also enhanced the directional influence of BLA to ACCg (but not ACCg to BLA) in these bands. Furthermore, prosocial versus antisocial choice types across the contexts can be decoded better and earlier after OT from the BLA-ACCg coherence in the two bands using a linear discriminant

analysis. Our results demonstrate that OT within BLA regulates synchrony between ACCg and BLA underlying prosocial and antisocial decisions.

III-7. Sensory to integrative coding of cued reward or punishment learning in thalamus and amygdala

Christopher Leppla CLEPPLA@MIT.EDU
Praneeth Namburi PRANEETH.NAMBURI@GMAIL.COM
Gordon Glober GLOBERG@MIT.EDU
Jake Olson JMOLSON@MIT.EDU
Yu Feng HYFY1995@SINA.COM
Maya Jay MAYAJAY@MIT.EDU
Kay Tye KAYTYE@MIT.EDU

Massachusetts Institute of Technology

The ability to accurately predict the motivational significance of environmental stimuli under changing conditions is critical for everyday life. Valence processing, a term describing this capacity for precise differentiation of stimuli predicting positive or negative outcomes is fundamental for advantageously navigating dynamic environments. Although the medial geniculate nucleus (MGN) of the thalamus and basolateral amygdala (BLA) have been examined in order to better understand the neural circuitry underlying valence encoding and the formation of associative learning, many open questions still exist. MGN is often thought of as a sensory relay station and BLA the first site of associative learning, these assumptions have not been directly tested. In this study, we ask the following questions: Does valence processing occur in MGN? Do amygdala-projecting MGN cells preferentially encode auditory information? How is information transformed from a sensory signal into an integrated associative memory that drives motivated behavior?

Our results confirm the long-held but previously unproven belief that MGN transmits robust auditory conditioned stimulus (CS) information to BLA during learning, with 91% of BLA-projectors encoding CS information during Discrimination. Our results also demonstrate that BLA-projecting MGN cells are more task-responsive than the overall MGN population and transmit a more homogeneous signal about task-relevant tone information than the overall population. Additionally, MGN exhibits more excitatory responding to reward-predictive cue offset than BLA, and a trend for the punishment-predictive cue as well.

In order to characterize this circuit during associative learning, we employed simultaneous multi-site single-unit in-vivo electrophysiology with circuit-specific optogenetic photoidentification during a multi-session Pavlovian task. This allowed for simultaneous characterization of MGN and BLA encoding as well as photoidentification of both BLA-projecting cells in MGN and those downstream BLA cells receiving input. Although the study is ongoing, our results already provide novel data where assumptions have dominated the field of Pavlovian associative learning.

III-8. Deep models of retinal responses to natural scenes generalize to diverse structured stimuli

Niru Maheswaranathan[1] NIRUM@GOOGLE.COM
Lane McIntosh[2] LANEMCINTOSH@GMAIL.COM
David Kastner[3] DBKASTNER@GMAIL.COM
Luke Brezovec[2] BREZOVEC@STANFORD.EDU
Aran Nayebi[2,2] ANAYEBI@STANFORD.EDU
Surya Ganguli[2] SGANGULI@STANFORD.EDU
Stephen Baccus[2] BACCUS@STANFORD.EDU

[1]Google Brain
[2]Stanford University
[3]University of California, San Francisco

Normal retinal function involves encoding information about the natural world. However, most of our understanding of neural computations and circuit mechanisms in the retina derives from artificial stimuli, such as flashing spots, drifting gratings, moving bars, and white noise. Although these stimuli have revealed a number of adaptive computations and nonlinear responses to specific stimulus features, they are extreme caricatures of stimuli likely to occur during ethologically relevant vision. Although the field has invested heavily in the study and analysis of responses to artificial stimuli, we neither understand to what degree natural vision engages the diverse retinal computations elicited by artificial stimuli, nor understand the relationship between these computations and underlying retinal circuitry. Here, we report that convolutional neural network models (CNNs) of salamander retinal ganglion cells (RGCs) trained solely on natural image sequences not only accurately capture RGC responses, but also provide a unified model of a wide range of retinal phenomena derived from responses to artificial stimuli not explicitly present in the natural images used for training. These phenomena include changing temporal bandwidth during contrast adaptation, latency encoding, the omitted stimulus response, polarity (kernel) reversal, motion anticipation, and motion reversal. CNNs trained on white noise stimuli do not exhibit this rich phenomenology. Furthermore, the responses of the model's internal units (cell types), which were never directly trained, are nevertheless highly correlated with recordings from bipolar and amacrine cells. Our results reveal that circuits underlying these previously described phenomena are engaged by and relevant during natural vision, and that natural stimuli but not white noise is sufficient to reveal these computations. Overall, our work uncovers a unified retinal model, with internal units representing interneurons, that can capture retinal responses to both natural and artificial stimuli, thereby elucidating the relationship between natural stimuli and a diverse set of visual computations.

III-9. Synthesizing speech from the human sensorimotor cortex

Gopala Anumanchipalli[1] ANUMANCHIPALLIG@NEUROSURG.UCSF.EDU
Edward Chang[1] EDWARD.CHANG@UCSF.EDU
Josh Chartier[1,2] JOSH.CHARTIER@UCSF.EDU

[1]University of California, San Francisco
[2]University of California, Berkeley

The ventral sensorimotor cortex (vSMC) encodes coordinated, multi-articulator kinematic movements of the vocal tract that accomplish specific articulatory goals needed to produce natural continuous speech. Our goal here was to decode audible speech only from the associated neural activity during speaking. There are three approaches towards this goal - i) direct decoding of audio speech spectrum, ii) decoding discrete word sequences, followed by text-to-speech synthesis and iii) decoding articulatory kinematics followed by articulatory synthesis. Of these choices, kinematics is the closest representational correlate to vSMC neural activity, has the least relative latency, and generalizes well to arbitrary word sequences, without limitations to a fixed vocabulary. Despite these advantages, modeling articulatory kinematics for neural-to-speech decoding has not been previously demonstrated

given methodological constraints on estimating articulatory movements. Here, we developed a model-based approach with two components, i) a neural decoder that converts neural activity into articulator kinematics, and ii) an articulatory synthesizer, that converts articulatory trajectories into audible speech. Both components were computationally implemented using deep recurrent neural networks with LSTM units. Neural data were collected from five human participants (patients with medically refractory epilepsy), implanted with high-density subdural ECoG arrays, as they spoke fluent sentences. An optimal set of 460 sentences (MOCHA-TIMIT) was used as the speaking material. A previously developed statistical approach was employed for acoustic-to-articulatory inversion to estimate vocal tract kinematics from produced speech acoustics. Articulators were represented as 12 dimensional vectors coding displacements in x and y directions of three points on the tongue, jaw, upper and lower lips and the fundamental frequency coding the laryngeal function. Using this approach, we were able to successfully decode neural activity to synthesize intelligible speech. We found high degree of correlation between synthesized and original spectrograms of subjects' produced speech making this a viable path for future speech based brain-computer interfaces.

III-10. Bird-song learning through quenched reinforcement learning

Ran Rubin RR2980@COLUMBIA.EDU
Laurence F. Abbott LFA2103@COLUMBIA.EDU

Columbia University

During their adolescence, Zebra Finches learn to produce a song motif imitating a teacher song sung to them in early childhood. The bird's song learning is a trial-and-error process thought to be an example of reward based reinforcement learning. Indeed, several fundamental components of reinforcement learning, such as reward-prediction error signal and mechanisms for neuronal activity perturbation, have been found experimentally in the bird's anterior-forebrain pathway (AFP). In this work, we use a detailed model of song production, to examine which type of reinforcement learning is capable of learning an actual song, relying solely on comparing the produced song to the memorized teacher song. Our model is based on current knowledge of the bird's motor pathway, modeling the transformation of pre-motor neuronal activity to an actual auditory signal through several strongly nonlinear stages implementing significant dimensionality reduction and expansion. We find that traditional reinforcement algorithms that average random fluctuations to perform a type of local, stochastic gradient descent generally perform poorly and fail to learn the teacher's song. We propose a novel type of reinforcement learning based on learning only from specific quenched realizations of error-reducing fluctuations in pre-motor neuronal activity. This algorithm generally yields better results and learns the teacher's song with relatively high fidelity. We hypothesize that this strategy is preferable because it promotes larger fluctuations in the learned neuronal activity which, in turn, enable a more global exploration of the error function landscape. Interestingly, experiments have identified area X in the AFP as receiving direct input form the time keeping area HVC, a dopamine based reward-prediction error signal and, importantly, an efferent copy of the activity in LMAN, an area that generates fluctuations in the pre-motor neuronal activity. As such area X is a candidate for implementing our proposed Quenched Reinforcement Learning.

III-11. Spatial selection asymmetrically modulates spike count correlations in mouse primary visual cortex

Ethan McBride[1,2] ETHAN.G.MCBRIDE@GMAIL.COM
Su-Yee Lee[1] SULEE@SALK.EDU
Edward Callaway[1] CALLAWAY@SALK.EDU

[1]Salk Institute for Biological Studies
[2]University of California, San Diego

Local decreases in spike count correlations in sensory areas are associated with improved perception of the corresponding location, particularly in primate attention studies. Global decreases are observed in mouse primary visual cortex (V1) with locomotion. To investigate whether decreased correlations are related to improved behavioral performance in mice, we developed a mouse spatial selection task. Head-fixed mice are shown drifting grating stimuli on left and right screens, must detect a contrast change that is more likely to occur in one location (80%) than the other (20%) and lick within a short time window to receive a reward. These mice (80-20) perform well on likely change trials and poorly on unlikely change trials, selectively using information from the likely change side. We also trained mice with equal change probabilities (50-50). 50-50 mice reliably detect changes on both left and right sides. While locomotion improves performance in 50-50 mice, it impairs performance in 80-20 mice, suggesting the two groups are using different brain states to perform their respective tasks. Electrophysiological recordings in both hemispheres of V1 show that on correct vs. miss trials, correlations are locally reduced in 80-20 mice in the hemisphere of V1 that represents the more likely change. In contrast, correlations are globally reduced in 50-50 mice. In addition, we found that locomotion globally increases correlations in 80-20 mice, and globally decreases correlations in 50-50 mice. These results suggest a close link between decreased correlations and improved perception in mice, and imply that the 80-20 mice employ a spatially selective brain state to perform the task, whereas 50-50 mice use global mechanisms compatible with the effect of locomotion. Moving forward, we plan to optogenetically manipulate cell types and investigate their role in altering correlations and producing and supporting this spatially selective brain state.

III-12. Deep sparse coding for invariant multimodal Halle Berry neurons

Edward Kim[1] EDWARD.KIM@VILLANOVA.EDU
Darryl Hannan[1] DHANNAN1@VILLANOVA.EDU
Garrett Kenyon[2] GKENYON@LANL.GOV

[1]Villanova University
[2]Los Alamos National Lab

Deep feed-forward convolutional neural networks (CNNs) have become ubiquitous in machine learning and computer vision; however, advancements in CNNs have arguably reached an engineering saturation point where incremental novelty results in minor performance gains. Although there is evidence that object classification has reached human levels on narrowly defined tasks, for general applications, the biological visual system is far superior to that of any computer. Research reveals there are numerous missing components in feed-forward deep neural networks that are critical in mammalian vision. The brain does not work solely in a feed-forward fashion, but rather all of the neurons are in competition with each other; neurons are integrating information in a bottom up and top down fashion and incorporating expectation and feedback in the modeling process. Furthermore, our visual cortex is working in tandem with our parietal lobe, integrating sensory information from various modalities.

In our work, we sought to improve upon the standard feed-forward neural network by augmenting them with biologically inspired concepts of sparsity, top-down feedback, and lateral inhibition. We define our model as a sparse coding problem using hierarchical layers. We solve the sparse coding problem with an additional top-down feedback error driving the dynamics of the neural network. While building and observing the behavior of our model, we were fascinated that multimodal, invariant neurons naturally emerged that mimicked, “Halle Berry neurons” found in the human brain. These neurons trained in our sparse model learned to respond to high level concepts from multiple modalities, which is not the case with a standard feed-forward autoencoder. Furthermore, our sparse representation of multimodal signals demonstrates qualitative and quantitative superiority to the standard feed-forward joint embedding in common vision and machine learning tasks.

III-13. OnACID: Online Analysis of Calcium Imaging Data in real time

Johannes Friedrich[1] JFRIEDRICH@FLATIRONINSTITUTE.ORG
Anne Churchland[2] CHURCHLAND@CSHL.EDU
Dmitri Chklovskii[1] DCHKLOVSKII@FLATIRONINSTITUTE.ORG
Liam Paninski[3] LIAM@STAT.COLUMBIA.EDU
Andrea Giovannucci[1] AGIOVANNUCCI@FLATIRONINSTITUTE.ORG
Matthew Kaufman[2] MKAUFMAN@CSHL.EDU
Eftychios Pnevmatikakis[1] EPNEVMATIKAKIS@FLATIRONINSTITUTE.ORG

[1]Flatiron Institute
[2]Cold Spring Harbor Laboratory
[3]Columbia University

Optical imaging methods using calcium indicators are critical for monitoring the activity of large neuronal populations in vivo. Imaging experiments typically generate a large amount of data that needs to be preprocessed to extract the activity of the imaged neuronal sources. While deriving algorithms for this purpose is an active area of research, most existing methods require the processing of large amounts of data at a time, rendering them vulnerable to the volume of the recorded data, and preventing real-time experimental interrogation. Here we introduce OnACID, an Online framework for the Analysis of streaming Calcium Imaging Data, including i) motion artifact correction, ii) neuronal source extraction, and iii) activity denoising and deconvolution. Our approach combines and extends previous work on online dictionary learning and calcium imaging data analysis, to deliver an automated pipeline that can discover and track the activity of hundreds of cells in real time, thereby enabling new types of closed-loop experiments. We apply our algorithm on two large scale experimental datasets, benchmark its performance on manually annotated data, and show that it outperforms a popular offline approach.

III-14. Zero-shot neural decoding of object category

Thomas O'Connell[1] THOMAS.OCONNELL@YALE.EDU
Marvin Chun[1] MARVIN.CHUN@YALE.EDU
Gabriel Kreiman[2] GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU

[1]Yale University
[2]Harvard University

The ability to recognize novel objects on the first exposure is known as "zero-shot" learning, in contraposition to traditional supervised learning that requires large numbers of learning examples. Here, we demonstrate a zero-shot neural decoding system that can categorize objects from neural activity despite having had no exposure to neural activity from the test categories during training. Our system employs a convolutional neural network (CNN) trained for object categorization and neural recordings from macaque inferior temporal (IT) cortex in response to natural object stimuli. Instead of learning a direct mapping between IT activity and object category (e.g. Hung et al., 2005), we learn an indirect mapping between IT activity and category labels using CNN unit activity as an intermediate basis representation. First, we build linear readout models to discriminate between pairs of object categories from CNN unit activity sampled from across the model's hierarchy. Next, we learn a linear mapping from IT neuron space to CNN unit space using activity evoked by a subset of the object categories. Finally, we test our model with neural activity from novel categories by transforming test firing patterns from IT-space to CNN-space using the learned transformation and then making predictions regarding which category evoked a given pattern of firing activity using the CNN readout models. Previous work showed that CNN models can capture a significant fraction of the variance in IT recordings (Yamins et al., 2014). This work provides additional evidence that CNNs provide a reasonable basis-space for representations in IT by demonstrating that the mapping between IT neurons and CNN units generalizes across categories. We show the potential of building more flexible, general neural decoding systems and directly translating between neural representations and model representations to apply computational vision models directly to neural activity.

III-15. Gain control in the motor system

Joshua Glaser[1] J-GLASER@U.NORTHWESTERN.EDU
Stephanie Naufel[1] STEPH.NAUFEL@GMAIL.COM
Eric Perreault[1] E-PERREAULT@NORTHWESTERN.EDU
Konrad Kording[2] KOERDING@GMAIL.COM
Lee Miller[1] LM@NORTHWESTERN.EDU

[1]Northwestern University
[2]University of Pennsylvania

Whether placing a contact lens in our eye, or lifting a heavy weight, the motor system regularly facilitates a wide range of dynamic conditions involving various loads. How does the motor cortex, which has a limited bandwidth, function effectively over this wide dynamic range?

Here, we explored the interaction of activity in primary motor cortex (M1) and muscle output (EMGs) during wrist movements made during three tasks requiring a variety of different loading conditions. A single linear model failed to account for the relation between M1 activity and EMG across the broad range of conditions. However, a model with a downstream gain parameter that increased with the target force remained accurate across forces and dynamical conditions. Although cortical output increased with force, the gain component increased even more, accounting for the majority of the EMG increase. We conclude that a mechanism downstream of our recorded M1 neurons amplified their output in order to provide for the full range of forces.

III-16. Learn, forget, relearn, amplify: codependent synaptic plasticity in spiking networks

Everton J Agnes EVERTON.AGNES@CNCB.OX.AC.UK
Tim Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Cortical networks display a rich repertoire of dynamics that depends on brain region and contextual state. The underlying synaptic architecture, feedforward or recurrent, must be carefully constructed to accommodate the appropriate activity. Most attempts to learn structure into network models have utilized independent, fine-tuned plasticity rules for excitation and inhibition with artificial boundaries. Here we present a set of “codependent” synaptic plasticity rules in which each synapse type adjusts its efficacy as a function of all local synaptic activity. The excitatory synaptic plasticity rule has a long-term potentiation (LTP) component that depends on excitatory currents from neighbouring excitatory synapses - a phenomenon generally referred to as cooperativity. Additionally, input from neighbouring inhibitory neurons control the learning rate of both long-term depression (LTD) and LTP, making the synapse ‘depend’ on its inhibitory partner. Codependent inhibitory synaptic plasticity, on the other hand, aims to balance excitatory and inhibitory currents, without any direct homeostatic control over the neuron’s baseline firing-rate. These entwined rules lead to an extraordinarily stable system, because neither synapse type can grow out of bounds. Our synapse models reproduce most experimental results on spike-timing and rate-dependent plasticity in paired recordings. Additionally, their combined action displays some remarkable features. In feedforward architectures, codependent synaptic plasticity maintains stable tuning, but can rapidly remap or erase its receptive field when the system is dis-inhibited, consistent with experiments in mice auditory cortex. In recurrent networks where both excitatory and inhibitory synapses are plastic, the network learns to generate a stable, balanced baseline activity. Reminiscent of recent observations from motor cortex, certain inputs trigger transient amplification with strong, dampened oscillations. We show that intricate learning dynamics emerge naturally when plastic synapses interact. Consistent with recent experimental observations, codependent synaptic plasticity reveals itself as a good candidate for a general theoretical framework of synaptic plasticity.

III-17. Excitatory and inhibitory neural populations coordinate to drive decisions in both novice and expert subjects

Farzaneh Najafi[1] FARZNAJ@GMAIL.COM
Anne Churchland[1] CHURCHLAND@CSHL.EDU
Gamaleldin Elsayed[2] GAMALELDIN.ELSAYED@GMAIL.COM
Eftychios Pnevmatikakis[3] EPNEVMATIKAKIS@SIMONSFOUNDATION.ORG
John Cunningham[2] JPCUNNI@GMAIL.COM

[1]Cold Spring Harbor Laboratory
[2]Columbia University
[3]Simons Foundation

Decisions are driven by the coordinated activity of diverse neural populations in multiple structures. Inhibitory neurons are critical to models of decision-making, but their in vivo role is largely untested. Here, we compare the contributions of excitatory and inhibitory neural populations by imaging their neural activity in the posterior parietal cortex during decision-making. Mice judged a series of multisensory “events” (clicks/flashes), the rate of which fluctuated stochastically over 1000 ms. Mice reported choices with a left/right lick. In each session, ~600 neurons were simultaneously recorded using 2-photon imaging while mice performed ~400 trials. To evaluate the relationship between neural activity and decision-making, we trained classifiers to distinguish single-trial activity preceding left vs. right choices. We first replicated that overall population activity could reliably predict choice. We then evaluated excitatory and inhibitory neurons separately. Surprisingly, decoders trained on inhibitory vs. excitatory neurons were remarkably similar both in accuracy, providing support for models in which excitatory pools target selective inhibitory pools. Interestingly, both populations were also similar in terms of classifier stability: the degree to which a classifier trained at one moment during the trial was effective at other moments in the trial. Additionally, accuracy and stability increased in a coordinated fashion between excitatory and inhibitory neurons as animals transitioned from novice to expert decision-makers. This argues that parallel classification accuracy by the two populations does not require extensive experience. These findings argue in favor of decision-making models in which pools of inhibitory neurons are specifically targeted by populations of excitatory neurons in favor of a particular choice.

III-18. An evolving code for cognitive context during accumulation of evidence

Sue Ann Koay[1] KOAY@PRINCETON.EDU
Stephan Thiberge[1] THIBERGE@PRINCETON.EDU
Carlos Brody[2,1] BRODY@PRINCETON.EDU
David Tank[1] DWTANK@PRINCETON.EDU

[1]Princeton University
[2]Howard Hughes Medical Institute

An evidence accumulation process is thought to underlie decision-making driven by noisy/ambiguous sensory information. Although many brain regions are involved in such tasks, and signatures of top-down feedback have been observed in early sensory regions, many studies conceptualize decision-making as a feedforward process, with stimulus-driven responses in sensory cortices temporally integrated in association areas, then categorized into choice downstream. We trained mice in a virtual T-maze navigation task in which pulsatile visual cues are randomly presented on their left/right as they run up the corridor; the rewarded arm is the side with the greatest total. Behaviorally, mice were found to accumulate multiple cues over many hundreds of milliseconds. We performed cellular resolution imaging in V1, several extrastriate areas, and retrosplenial cortex (RSC). Qualitatively similar neural dynamics were observed in layers 2/3 and 5 of all these areas. The majority of active cells formed, as a population, choice-specific sequences that spanned the duration of the trial. We also observed cells with activity time-locked to the pulsatile cues. The amplitudes of these responses were progressively and strongly modulated

by choice and other cognitive information throughout the accumulation period (c.f. Britten et al. Vis Neurosci 1996). Notably, simultaneously recorded cells of the same cue-side preference sometimes exhibited opposite choice modulations. Although these effects were observed as early as in V1, there is a gradual increase, going towards RSC, in both the persistence of sensory responses and the strengths of cognitive modulations. These data are consistent with a view of parallelized (feedforward) mechanisms for decision formation, and raise the question of whether such modulations arise jointly or independently per region. The strong cognitive modulations we observe in RSC could provide the substrate for choice-specific changes in spatial representations that guide navigation through the correct turn in the T-maze.

III-19. Functional assessment of large-scale cortical networks during multi-sensory decision-making

Simon Musall SIMON.MUSALL@GMAIL.COM
Matthew Kaufman MKAUFMAN@CSHL.EDU
Anne Churchland CHURCHLAND@CSHL.EDU
Steven Gluf SGLUF@CSHL.EDU

Cold Spring Harbor Laboratory

The brain continually integrates sensory inputs to produce appropriate behavior. This transformation from sensory input to motor output involves many cortical areas, but how areas interact in a large-scale decision-making circuit remains elusive. To gain a broad perspective on cortical decision-making, we trained head-fixed mice in a delayed two-alternative forced-choice paradigm for sensory stimulus detection. Mice were presented with visual moving bars or auditory clicks, waited one second and then reported the stimulus side. Using a tandem lens macroscope and GCaMP6f transgenic mice, we measured large-scale cortical activity patterns through the cleared skull during behavior. To attribute activity in specific cortical areas to specific stimuli, behavioral events, and decision variables, we fitted a linear model incorporating numerous task parameters (e.g., stimulus onset, licking and pupil diameter). The model accurately predicted changes in neural activity throughout the dorsal cortex, identifying expected areas as involved with sensory and motor events. We then quantified the predictive contribution of each parameter to the model. Surprisingly, choice and behavioral success did not contribute predictive power; instead, activity was best described as locked to sensory and motor events. This result indicates that our widefield recordings primarily reflect sensory inputs and motor outputs, whereas decision-related variables, though they appeared to modulate neural response, were undetectable at this scale. To better identify distinct regions of interest, we performed clustering using noise correlations between pixels. Despite whitening the correlations with respect to inter-pixel distance, the clusters were spatially compact and always included corresponding regions in both hemispheres. Clusters appeared to correspond to functionally related brain areas and were organized by body part, such as the hindlimb representation, but often spanned both sensory and motor areas. These results change our interpretation of ostensibly decision-related activity in dorsal cortex and open new avenues to better exploring widefield calcium data for targeting cellular resolution recordings and inactivations.

III-20. Attention modulates V1 neuronal responses through a depolarizing mechanism

Duy Tran[1] DUYTRAN1125@G.UCLA.EDU
Peyman Golshani[1] PGOLSHANI@MEDNET.UCLA.EDU
Michael Einstein[1] EINSTEIM@G.UCLA.EDU
Pierre-Olivier Polack[2] POLACK.PO@RUTGERS.EDU

[1]University of California, Los Angeles
[2]Rutgers University

Understanding the cellular and network mechanisms of attention is critical for the development of rational therapies for diseases with attention deficits. Extracellular recordings from non-human primates show that attention increases the responsiveness of visual cortical neurons encoding an attended visual stimulus (Chalk et al., 2010; Fries et al., 2001b; McAdams, 2005; McAdams and Maunsell, 1999; Mehta, 2000a; Motter, 1993; Reynolds et al., 2000b; Treue and Maunsell, 1996b), yet there is little direct evidence to explain this phenomenon due to technical shortcomings in non-human primates. Using a novel multimodal attention task for mice, we performed 2-photon guided whole-cell recordings from layer 2/3 (L2/3) primary visual cortex (V1) neurons as animals attended or ignored visual cues. These data show that L2/3 V1 neurons depolarize while animals attend visual cues relative to when they ignore visual cues. As a result, we show the first direct evidence that a slight depolarization could account for the increased responsiveness of neurons during attention. By depolarizing the membrane potential (Vm) of neurons in the visual cortex, attention would set neurons closer to their action potential threshold, increasing the responsiveness of neurons to excitatory inputs (Polack et al., 2013). Additionally, our data from 128 channel extracellular recordings reveal that V1 neurons significantly increased their responsiveness to visual cues and V1 enters a relatively desynchronized brain state during attention, replicating core findings from experiments in non-human primates. Therefore, these data provide the first direct evidence that V1 neurons are more depolarized during attention and demonstrate the feasibility of this novel animal model for future study.

III-21. Decision-making through evidence integration at long timescales

Michael Waskom MWASKOM@NYU.EDU
Roozbeh Kiani ROOZBEH@NYU.EDU

New York University

Evidence accumulation models of decision-making propose that decisions are formed by integrating multiple samples of evidence, which produces a decision variable that can be compared against a threshold to determine choice. Evidence integration has most often been studied in the context of simple sensory decisions, typically employing a continuous stimulus, and usually limiting the timescale of decision formation to hundreds of milliseconds. In contrast, many naturalistic decisions involve multiple discrete pieces of evidence that are encountered at unpredictable times across longer timescales. Can decision-making in these circumstances be explained by evidence integration? We have developed a novel paradigm to examine this question. In our paradigm, subjects see several brief “pulses” of a contrast patch; each pulse appears with a variable contrast, and pulses are separated by gaps of up to 8 seconds. The task is to determine which of two overlapping contrast distributions generated the pulses on each trial. Optimal performance would be achieved through linear integration of evidence afforded by the pulses in units of log-likelihood ratio. However, the long gaps may introduce noise or leak into the integration process. The discrete samples of evidence and long timescale could also encourage alternative decision-making strategies, such as detection of extreme values or counting the number of pulses that favor each alternative. We used computational modeling and model-free analyses of behavior to evaluate these competing accounts. We found that human behavior approximates the optimal strategy: subjects weighted pulses equally regardless of the history of pulses or the duration of the gaps between them. This indicates that there is no biophysical limitation that prevents generalization of integration to longer time scales and discontinuous streams of evidence. We suggest that variability of decision strategy in naturalistic conditions, when present, stems from a lack of emphasis on accuracy and inadequate understanding of the task.

III-22. Invariances in a combinatorial olfactory receptor code

Guangwei Si[1] GSI@FAS.HARVARD.EDU
Aravinthan Samuel[1] SAMUEL@PHYSICS.HARVARD.EDU
Jessleen Kanwal[1] JKANWAL@FAS.HARVARD.EDU
Yu Hu[2,3] YU.HU@MAIL.HUJI.AC.IL
Christopher Tabone[1] CTABONE@FAS.HARVARD.EDU
Jacob Baron[1] JBARON@PHYSICS.HARVARD.EDU
Matthew Berck[1] BERCKM@JANELIA.HHMI.ORG
Gaetan Vignoud[1] GAETAN.VIGNOUD@HOTMAIL.FR

[1]Harvard University
[2]Hebrew University of Jerusalem
[3]ELSC

Animals can identify an odorant type across a wide range of concentrations, as well as detect changes in concentration for variant odorant types. How primary olfactory representations are structured in a complete sensory system to support these functions remains poorly understood. We have uncovered a set of quantitative invariances in the representation of odorant identity and intensity in the Drosophila larva which could simplify this fundamental problem. We find that dose-response relationships across odorants and olfactory receptor neuron (ORN) types follow the Hill function with shared cooperativity but different activation thresholds. Thus, the only free variable in an interaction between any odorant molecule and receptor is the activation threshold. Furthermore, all activation thresholds across odorants and receptors are drawn from the same power law statistical distribution. A fixed activation function and power law distribution of activation threshold predict the scaling relationship between odor intensity and overall receptor activity across olfactory space. We find that similar temporal response filters of ORNs across odorant types and concentrations extend these invariances in olfactory representation to fluctuating or turbulent environments. Common patterns in ligand-receptor binding and sensory transduction across olfactory receptors may give rise to these invariant quantitative structures in the olfactory code. Invariant patterns in the activity responses of individual ORNs and the ORN ensemble may simplify decoding by downstream circuits.

III-23. Premotor network exploration during practice of a stereotyped, learned motor action

William Liberti III[1] WLIBERTI@BERKELEY.EDU
Jun Shen[2] JUNSHEN@BU.EDU
Nathan Perkins[2] LNP@BU.EDU
Daniel Leman[2] DPLEMAN@BU.EDU
Timothy Gardner[2] TIMOTHYG@BU.EDU

[1]University of California, Berkeley
[2]Boston University

With practice, learned motor actions can be maintained for decades, but the biological basis of skill acquisition and persistence remains largely unknown. We do know that precise motor skills are learned and maintained through practice; a process of exploratory trial-and-error learning where actions are evaluated and ‘good’ performances are reinforced. However, the neural underpinnings of how practice may evaluate, maintain and sharpen the stereotypy of a precise motor action are not well understood. The adult male zebra finch sings a highly stereotyped song as a courtship ritual to attract females. These finches will also sing in isolation, otherwise known as undirected song, and these songs have slightly lower stereotypy. Undirected song could serve as an opportunity for premotor networks to explore potentially better configurations to optimize motor output, since the performance outcome is less critical. We deploy head mounted miniature microscopes and cell-type specific genetic tools in the pre-motor nucleus HVC to investigate if this area undergoes state-dependent changes in variability. We find that on average, individual neurons have similar firing times in both conditions. However during undirected

practice, some neurons are active more probabilistically, and less robustly- and trial to trial variability is not highly correlated across cells. In contrast, when song is directed to a female, this exploration is largely diminished. These state-dependent variations could contribute to the long-term stability of the motor network, or to trial and error learning to minimize vocal errors.

III-24. Perceptual straightening of natural videos

Olivier Henaff[1] OJH221@NYU.EDU
Robbe Goris[2] ROBBE.GORIS@UTEXAS.EDU
Eero Simoncelli[1] EERO.SIMONCELLI@NYU.EDU

[1]New York University
[2]The University of Texas at Austin

Much of behavior relies on making predictions about future outcomes given recent observations. Yet visual input evolves according to complex, nonlinear dynamics that are difficult to extrapolate. We hypothesize that the brain transforms the incoming stream of images so as to make them more predictable. Specifically, we propose that internal representations are structured so as to straighten the trajectories of natural videos, thereby enabling prediction through linear extrapolation. In contrast, image sequences that are unlikely to occur in the real world need not be straightened and will most likely be distorted by the visual system. To test this 'temporal straightening' hypothesis, we developed a novel procedure for estimating the curvature of the human perceptual representation of a sequence of images. Specifically, we first measured the discriminability of pairs of briefly presented video frames. Next, we formulated an observer model in which each frame is represented as a Gaussian distribution in a fixed-dimensional perceptual space. The likelihood of a trajectory can be evaluated by computing the probability that the measured human responses would have arisen from the overlap of the corresponding distributions. Finally, we derived a data-efficient and largely unbiased estimator of perceptual curvature by searching for the curvature that best describes all plausible perceptual trajectories. By comparing this value to the curvature calculated from the pixel intensities of the image sequence, we tested three distinct predictions of our hypothesis. First, natural videos that are curved in the intensity domain should be straighter perceptually. Conversely, unnatural videos that are straight in the intensity domain should be curved perceptually. Finally, natural videos that are straight in the intensity domain should remain straight perceptually. Our results are consistent with all three predictions, demonstrating that the visual system deploys nonlinear transformations targeted to straighten natural videos, in support of tasks that rely on prediction.

III-25. Enhanced capacity and dynamic gating in a model of context-dependent associative memory

William F Podlaski WILLIAM.PODLASKI@CNCB.OX.AC.UK
Everton J Agnes EVERTON.AGNES@CNCB.OX.AC.UK
Tim Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

An increasing amount of evidence suggests that memory formation and retrieval are modulated by contextual signals, such as behavioral or emotional state. However, typical models of associative memory do not incorporate this dependency. Here we propose an extension to the Hopfield network which takes into account contextual modulation. The network is divided into a set of overlapping subnetworks, each representing a different context with a separate set of memory patterns. Only one subnetwork is active at any given time, thereby reducing interference from memories found in other contexts, which remain dormant through inhibitory control. Using theoretical and numerical methods, we show that these context-modular Hopfield networks have substantially increased memory capacity, as well as robustness to noise and to memory overloading. Their performance

depends on two parameters—the number of subnetworks, and their relative size—and when chosen optimally, the capacity is up to ten times greater than the standard Hopfield model. Improved performance comes at the cost of limited retrieval, because only memories stored in the active subnetwork can be recalled. To address this, we propose a system in which a controller network dynamically switches the memory network to a desired contextual state before storage or retrieval. Through simulations, we successfully show that this system is able to bias memory retrieval based on context. Overall, our work illustrates the benefits of context-dependent memory, and may have implications for our understanding of cortical memories and their interaction with contextual signals in the prefrontal cortex and hippocampus.

III-26. Oscillatory encoding of visual stimulus familiarity

Samuel Kissinger[1] ST.KISSINGER@GMAIL.COM
Sotiris Masmanidis[2] SMASMANIDIS@UCLA.EDU
Alexandr Pak[1] PAK6@PURDUE.EDU
Alexander Chubykin[1] CHUBYKIN@PURDUE.EDU

[1]Purdue University
[2]University of California, Los Angeles

The primary visual cortex (V1) is well known as a feature detector, but has also been shown to exhibit significant experience dependent neural activity. Repetitive presentation of visual stimuli in awake mice results in significant potentiation of visually evoked potentials (VEPs) in V1, a phenomenon which is sensitive to the familiarity of the animal with the stimulus. Potentiation of VEP amplitude has also been shown to report learned sequences of visual stimuli. Interestingly, neural activity in the form of persistent spiking or theta (4-8Hz) oscillations in V1 has even been shown to report the timing of reward delivery after presentation of a visual cue. We sought to determine if visually evoked theta oscillations are sensitive to stimulus novelty, and explored potential mechanisms underlying their generation in V1. Using 64 channel silicon probes acutely implanted in binocular V1, we recorded VEP's and units in awake mice after several days of visual training to sinusoidal grating stimuli. Oscillations across all cortical layers, most prominently expressed in the theta range, emerged in both VEPs and individual units in an experience dependent manner. Intriguingly, only stimulus locked responses were observed when a novel visual stimulus was presented, suggesting that these oscillations are evoked by familiarity to visual stimuli. Both the acquisition and expression of these oscillations were blocked by the muscarinic acetylcholine receptor (mAChR) antagonist scopolamine, a drug previously shown to block other experience dependent phenomena in V1. Surprisingly, k-means clustering of oscillatory units revealed unique neuronal ensembles that fire preferentially to particular oscillation cycles. These subpopulations are biased towards particular cortical layers, providing insight into potential circuit mechanisms underlying the oscillation. Finally, the oscillations were not observed in the visual thalamus, suggesting that they are generated either locally within V1 or via feedback from higher order visual areas.

III-27. Matrix-normal models for fMRI analysis

Michael Shvartsman[1] MS44@PRINCETON.EDU
Adam Charles[1] ADAMSC@PRINCETON.EDU
Jonathan Cohen[1] JDC@PRINCETON.EDU
Mikio Aoi[1] MAOI@PRINCETON.EDU
Narayanan Sundaram[2] NARAYANAN.SUNDARAM@INTEL.COM
Ted Wilke[2] TED.WILKE@INTEL.COM

[1]Princeton University
[2]Intel Parallel Computing Lab

Functional magnetic resonance imaging (fMRI) is a challenging analysis problem in neuroscience: signal-to-noise ratio is low, data dimensionality is high, and noise creates spatiotemporal correlations that can mask signal and magnify false alarms. Early methods approached fMRI analysis with univariate regression, but the fundamentally spatial nature of fMRI limited their success, driving development of multivariate analysis methods organized under the umbrella of Multi- Voxel Pattern Analysis (MVPA; Norman et al., 2006). Unlike univariate fMRI methods, which are typically instances of the general linear model (GLM), MVPA methods have not been organized in a unified theoretical framework, with methods developed independently for different problems, including dimension reduction (e.g. Chen et al., 2015; Manning et al., 2014), correlation estimation (e.g. Cai et al., 2016; Simony et al., 2016), multivariate regression (Allefeld and Haynes, 2014), and classifier-based decoding (e.g. Polyn et al., 2005). While these methods have had considerable success, they created new challenges for interpretation and hypothesis testing (e.g. Allefeld et al., 2015; Cai et al., 2016; Schreiber and Krekelberg, 2013). Furthermore, the lack of unified theoretical perspective has led to a lack of consistency even in addressing identical problems. We propose matrix-variate normal (MN) models as a unifying framework for multivariate fMRI analysis. This framework combines explicit spatio-temporal modeling of fMRI with the formality and interpretability of probabilistic generative modeling, and includes as special cases many of the above methods. We use the formalism to develop a number of novel method variants: MN-RSA, MN-SRM and MN-ISFC. MN-RSA outperforms the previous best RSA method in both speed and accuracy, MN-SRM improves on SRM in reconstruction accuracy, and MN-ISFC achieves comparable accuracy to ISFC but is guaranteed to return valid correlation matrices. The shared mathematical structure additionally enables the creation of an MN modeling toolkit that admits flexible prototyping of spatiotemporal analysis methods.

III-28. Simple integration of fast excitation and offset, delayed inhibition computes directional selectivity in Drosophila

Eyal Gruntman GRUNTMANE@JANELIA.HHMI.ORG
Michael Reiser REISERM@JANELIA.HHMI.ORG
Sandro Romani SANDRO.ROMANI@GMAIL.COM

HHMI Janelia Research Campus

The detection of visual motion is a fundamental neuronal computation that subserves many critical behavioral roles. A neuron that extracts directionally selective motion information from signals that are not motion selective must compare the visual responses of spatially offset inputs. This computation has been studied for decades in both vertebrate and invertebrate visual systems and given rise to competing algorithmic models: either using a synergistic combination of offset excitatory inputs or using an inhibitory input to “veto” an offset excitatory input. In the Drosophila visual system, stimuli are processed through ON and OFF pathways. The 4th-order neuron, T4, is the first cell type in the ON pathway to exhibit directionally selective signals. Attempts to identify the mechanism responsible for directional selectivity in the T4 circuit have resulted in several contradictory claims. Due to the small size of these neurons, recordings have been restricted to calcium imaging, limiting temporal resolution and preventing the direct measurement of inhibition. These limitations preclude a clear demonstration of the neuronal computation underlying directional selectivity. Here we use in-vivo whole-cell recordings of T4 to show that directional selectivity originates from simple integration of spatially offset fast excitatory and slow inhibitory inputs, resulting in a suppression of responses to the non-preferred motion direction. We constructed a passive, conductance-based model of a T4 cell that accurately predicts the neuron’s response to moving stimuli. These results connect the known circuit anatomy of the motion pathway to the algorithmic mechanism by which the direction of motion is computed.

III-29. A novel deep recurrent network for predicting large scale population responses to natural video

Fabian Sinz[1] SINZ@BCM.EDU
Dimitri Yatsenko[1] DVYATSEN@BCM.EDU
Alexander Ecker[2] ALEXANDER.ECKER@UNI-TUEBINGEN.DE
Paul Fahey[1] PAUL.FAHEY@BCM.EDU
Edgar Walker[1] EYWALKER@BCM.EDU
Erick Cobos[1] ERICK.COBOS@BCM.EDU
Emmanouil Froudarakis[1] FROUMAN@GMAIL.COM
Jacob Reimer[1] REIMER@BCM.EDU
Andreas Tolias[1] ASTOLIAS@BCM.EDU

[1]Baylor College of Medicine
[2]University of Tuebingen

To understand the representations in visual cortex, we need to be able to faithfully predict neural activity in response to its natural input: a continuous video stream. Since cortical activity is highly variable and context-dependent, this prediction is already difficult for integrated neural activity to static natural images, and even more difficult for dynamic responses to movies. In awake animals under free-viewing conditions, eye movements and brain states add to this response variability, making the prediction problem even harder. While deep convolutional networks have recently been shown to improve prediction performance over linear-nonlinear type models and are currently considered state-of-the-art, they make suboptimal use of the data, because they cannot account for stimulus-independent variability.

Here, we developed a new deep recurrent network architecture that predicts the deconvolved Ca++ activity of thousands of simultaneously recorded neurons in mouse V1 to natural videos, recorded at 7Hz and 30Hz, respectively, while simultaneously estimating dynamic gaze position and brain state changes related to running state and pupil dilation. In addition to the natural movie input, the network uses pupil position and dilation extracted from a video of the animal's eye, as well as treadmill velocity. The unknown relation between pupil position and gaze position on the monitor is learned by the network during training based solely on predicting neural activity. We find that incorporating all these elements (nonlinear recurrent network, running speed, pupil position, and pupil dilation) significantly increases the prediction performance of the network. Our network achieves between 40% and 60% of a leave-one-out estimate of single-trial correlation with the mean response over repeated presentations. To the best of our knowledge, this makes our model the state-of-the-art on single trial prediction of dynamic responses to natural movies on large neuronal populations.

III-30. Low frequency local field potentials in the parietal reach region encode the pattern of bimanual reaching

Eric Mooshagian ERICMOOSHAGIAN@EMAIL.WUSTL.EDU
Chuck D. Holmes CHUCK@EYE-HAND.WUSTL.EDU
Lawrence H. Snyder LARRY@EYE-HAND.WUSTL.EDU

Washington University in St. Louis

Primates are able to use their two arms in complex ways. We reach with both arms to grasp a heavy object on a high shelf; we hold a mug with one hand while the other hand pours coffee into it from a pot; we steer our car with our left hand and shift gears with our right hand. Our understanding of the cortical basis for coordinating our limbs is limited. Previous work has shown that single units in the parietal reach region (PRR) primarily encode the intended movement of the contralateral limb, with small non-systematic encodings of the bimanual pattern and very little representation of the ipsilateral limb movement on its own. This leaves us with a puzzle: If PRR primarily encodes the contralateral arm, how are its bimanual coordination signals produced? We hypothesized

that single units in PRR in each hemisphere might exchange information about what the individual arms are doing. This would allow each hemisphere to form representations related to bimanual coordination without either side explicitly encoding the movement of the ipsilateral arm. We tested this hypothesis by recording the local field potential (LFP) in PRR. We asked what role PRR might play in unimanual and bimanual reach tasks by comparing LFP and single unit responses. Low frequency LFP power contains a rich representation of intended reach movements, including movement of the ipsilateral arm. The temporal characteristics of LFP across the two hemispheres (coherence) also indicate that information about the two arms is shared across hemispheres. We interpret these data to mean that each PRR codes information about the contralateral arm, and sends that information to the opposite PRR so that each side can generate contralateral arm control signals that result in coordination of the two arms.

III-31. Capturing monosynaptic dynamics from pairwise spike times in non-stationary conditions

Jonathan Platkiewicz ZSACCOMANO@GRADCENTER.CUNY.EDU
Zach Sacccomano ZSACCOMANO4@GMAIL.COM
Asohan Amarasingham ASOHAN.AMARASINGHAM@GMAIL.COM

City University of New York

A major challenge in systems neuroscience is the difficulty of studying brain connectivity in behavioral conditions. One approach is to attempt to infer microcircuit diagrams indirectly from large-scale spike population recordings (Fujisawa et al., 2008). In particular, fine-timescale pairwise spike interactions, exemplified by the sharp peak in the cross-correlogram of Panel B, can often indicate cells are monosynaptically-coupled. It is important that such an inference be agnostic regarding extreme nonstationarities in the presynaptic and network environment, which occur independently of monosynaptic dynamics. One such inference develops a rigorous nonparametric statistical method, based on the statistical approach of conditional inference, in which the underlying probability models explicitly separate fine- from coarse-time- scale parameters through the appropriate decomposition of the data distribution (Amarasingham et al., 2012). The timescale separation renders the resulting inference robust to coarse timescale nonstationarities in a striking way (see Additional Detail). There is however no validation or calibration yet of timescale separation, which is a modeling approach motivated on theoretical grounds, due to the lack of ground truth data. The recently-reported experimental work of (English et al., 2017) is a step toward this goal. Biophysical models of neural firing, closer to experimental physiology than statistical models, offer a framework for assessing the relevance of conditional models. In this study, we propose first an integrate-and- fire-based model of ultra-precise monosynaptic spike transmission, incorporating genuine fast nonstationary input as the source of spike-time unreliability (see Panel A and Additional Detail). Multiple biophysical studies examined the monosynaptic origin of fine-temporal spike correlation (e.g., Ostojic et al., 2009), but none considered the ultra-brief peak observed under in vivo conditions (English et al., 2017). Two key ingredients of our model are a slowly autocorrelated input noise (i.e., synaptic filtering) and an adaptive threshold dynamics (i.e., coincidence detection). Second, using a multi-scaled version of the conditional model, we derived an unbiased closed-form estimator for the monosynaptic contribution to inferred fine- temporal spike correlation based on the observed spike train data, and which is robust to (coarse) nonstationarities. Such a formulation speeds up significantly the spike data computational analysis. Third, we validated the performance of this formula on spike trains generated using our biophysical model. In particular, we were able to correctly track the change of the monosynaptic peak conductance on long and shorter timescale (see Panel C). We discuss implications for a global analysis of population data.

III-32. Saccade kinematics communicate covert decision-related computations during urgent choices

Joshua Seideman JSEIDEMA@WAKEHEALTH.EDU
Terrence Stanford STANFORD@WAKEHEALTH.EDU
Emilio Salinas ESALINAS@WAKEHEALTH.EDU

Wake Forest School of Medicine

The empirical study of perceptual decision making hinges on the ability to make inferences about covert cognitive states based on overt behaviors. And yet, while saccadic choice paradigms have been instrumental in advancing our understanding of decision making in general, few studies have directly linked saccade metrics (e.g., peak velocity) themselves to underlying decision-related processes. Thus, it is currently unknown whether and how the formation and development of a perceptual decision influence the metrics of a saccadic choice upon execution, or complementarily, whether the metrics of saccades are a reliable tell of perceptual decision-making dynamics. Here, we show that various saccade metrics are systematically modulated as functions of sensory cue viewing time during performance of an urgent-decision task. Using a race-to-threshold model previously proven to replicate both standard performance measures (e.g., choice accuracy, response time distributions) and cortical oculomotor neuronal activity (i.e., the frontal eye field; FEF) in the task, we present a plausible physiological mechanism by which sensory evidence influences saccade peak velocity. We discover that the simulated and behavioral data both correlate with the probability that a choice will be correct given the sensory evidence (i.e., with the statistical definition of confidence)—indicating that cortical neuronal activity and saccade kinematics encode/communicate the degree of certainty associated with the urgent perceptual decision-making process. Ultimately, our empirical and theoretical lines of evidence converge to support a unified, mechanistic framework whereby sensory evidence informs not only what saccadic choices we make, but when and how we make them. These findings not only (1) change the way we think about saccade metrics and their communicative capacity, but also (2) provide mechanistic insight into the cortical control of visually-guided eye movements, and (3) help bridge conceptual gaps between oculomotor and signal detection-theory (SDT) based models of decision making.

III-33. Does macaque anterior cingulate cortex represent the valence of social decision outcomes?

Sharika K. M. SHARIKA@MAIL.MED.UPENN.EDU
Michael Platt MPLATT@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Human fMRI studies have for long implicated anterior cingulate cortex (ACC) and anterior insula in the empathy for pain, fear or disgust (Wicker et al., 2003; Singer et al. 2004; Botvinick et al., 2005; Jackson et al., 2005; Saarela et al., 2007; Fan et al., 2011, for meta-analysis; Lamm et al., 2015). However, so far, neuronal recordings in macaque ACC gyrus have been shown to correlate with positive reward outcomes for self and others (Chang et al., 2013). Whether ACC neurons encode valence specificity of vicarious reward outcomes is not known as yet. To investigate this, we recorded the activity of ACC neurons while monkeys performed a 'willingness to pay' task in which they chose between two differently colored targets associated with varying magnitudes of juice across trials. The juice on offer was cued to be either sweet tasting or bitter tasting for the actor monkey (the one making the decisions) or a recipient monkey (sitting across the room and facing the actor monkey) in different trials. Preliminary results showed that not only did decisions made by the actor monkey indicate an awareness of vicarious reward outcomes learnt by observation alone, the activity of ACCg neurons distinctly represented positive and negative outcome for self and other as well. This is consistent with the idea of ACC's role in making value based decisions in a social context.

III-34. How strong are correlations in strongly recurrent neuronal networks?

Ran Darshan[1] DARSHANR@JANELIA.HHMI.ORG
David Hansel[2] DAVID.HANSEL@UNIV-PARIS5.FR
Carl van Vreeswijk[3] CARL.VAN-VREESWIJK@BIOMEDICALE.UNIV-PARIS5.FR

[1]HHMI Janelia Research Campus
[2]CNPP CNRS UMR8119 & Universite Paris Descartes
[3]Center for Neurophysics Physiology and Pathology; CNRS-UMR8119

Cross-correlations in the activity in neural networks are commonly used to characterize their dynamical states and their anatomical and functional organizations. Yet, how these latter network features affect the spatiotemporal structure of the correlations in recurrent networks is not fully understood. Here, we develop a general theory for the emergence of correlated neuronal activity from the dynamics in strongly recurrent networks consisting of several populations of neurons. We apply this theory to the case in which the connectivity depends on the anatomical or functional distance between the neurons. We establish the architectural conditions under which the system settles into a dynamical state where correlations are strong, highly robust and spatially modulated. We show that such strong correlations arise if the network exhibits an effective feedforward structure which can be inferred from the Jordan form of the interaction matrix. In networks lacking such a structure correlations are extremely small and only weakly depend on the number of connections per neuron. We demonstrate how the feedforward structure determines the way correlations scale with the network size and the degree of the connectivity. Our work shows how strong correlations can be consistent with highly irregular activity in recurrent networks, two key features of neuronal dynamics in the central nervous system.

III-35. An encoding model reveals spatiotemporal specificity of eye-related effects during natural vision

Hiroto Yamaguchi[1] HYFORMAL@GMAIL.COM
Satoshi Nishida[2] S-NISHIDA@NICT.GO.JP
Shinji Nishimoto[2] NISHIMOTO@NICT.GO.JP

[1]Osaka University
[2]National Institute of Information and Communications Technology

We incessantly make eye movements and blinks, but our visual experiences remain largely steady and uninterrupted. This phenomenon is referred to as visual consistency, and the functional characteristics underlying it are still unclear. Here we propose a new approach to study functional properties underlying visual consistency by modeling cortical representation and its effective modifications around eye-movements and blinks. We recorded blood-oxygen-level-dependent (BOLD) signals using functional MRI while human subjects watched naturalistic movies. Subjects were allowed to make eye movements, and we recorded the gaze positions and blink events. We then adopted the motion energy model as a voxel-wise encoding model. Specifically, retinal images were reconstructed by combining the movie scenes and the gaze positions, and motion energy features were calculated for the reconstructed images. Then we modeled BOLD signals by temporally convolving and summing up over the energy features. Parameters of the convolution and summation were estimated by ridge regression. By using the estimated models, we could reconstruct a retinotopic map, which indicates the feasibility of our approach. We further studied spatiotemporal specificities of saccadic eye movements and blinks. Considering spatiotemporal characteristics of these movements, it is conceivable that the effects are also specific in spatiotemporal domain. We thus incorporated eye-related suppressive effects into the model by systematically cutting off energy features in spatiotemporally specific manners. Their functional relevance was quantified by the modeling accuracy. Our result shows 1) we can predict the BOLD signals in the early and mid-level visual cortex under the free viewing condition, 2) models with simulated saccadic and blink suppressions provide more accurate predictions than a model without them, and 3) the effective spatiotemporal components of the simulated suppressions are different between saccade and blink. These results demonstrate the effectiveness of our approach in studying the

functions that do underlie visual consistency under naturalistic condition.

III-36. Reverse engineering transient computations in nonlinear recurrent neural networks through model reduction

Christopher Stock CHSTOCK@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Stanford University

Increasingly, complex recurrent neural networks are being employed to model biological neural circuits mediating behaviors involving sequential, context-dependent temporal processing. Often, such networks are successfully used as a black box to capture variance in dynamic neural data, but a conceptual understanding of how the connectivity and dynamics of such networks conspire to generate behavior remains lacking. Previous attempts at extracting such an understanding of dynamics alone involved studying fixed points of the dynamics. This yielded insight into the operation of nonlinear networks whose dynamics is indeed dominated by the fixed point structure. However, in many temporal signal processing tasks in biology, including sequence memory, sequence classification, and natural language processing, neural networks operate in a transient regime far from fixed points.

Here we develop a general method to attain conceptual insight into transient computations in recurrent networks by dramatically reducing the complexity of networks trained to solve transient processing tasks. Our method, called dynamically reweighted singular value decomposition (DR-SVD), performs a reweighted dimensionality reduction to obtain a much lower rank connectivity matrix that preserves the dynamics of the original neural network.

We apply DR-SVD to a network trained on a context-dependent temporal computation task, and we find low-rank recurrent network reductions that can solve the entire task. By contrast, we demonstrate that previous approaches to model reduction, based on fixed points or low-rank approximations to the connectivity itself, fail to replicate transient network computations. Because DR-SVD uncovers simplified, reduced and interpretable models of low dimensional network connectivity, without sacrificing accuracy in capturing dynamical network computations, we expect that it will be be widely applicable to conceptually understanding how transient computations unfold in recurrent neural networks.

III-37. Input correlations impede suppression of chaos and learning in balanced rate networks

Ramin Khajeh RK2899@COLUMBIA.EDU
Alessandro Ingrosso AI2367@COLUMBIA.EDU
Rainer Engelken RE2365@COLUMBIA.EDU
Laurence F. Abbott LFABBOTT@COLUMBIA.EDU
Columbia University

Cortical circuits exhibit complex activity patterns both spontaneously and evoked by external stimuli. Information encoding and learning in neural circuits depend on how well time-varying input can control network activity. Previous work showed that input correlations can be detrimental to learning in balanced networks, but the reasons for this were not clear. We show that in firing-rate networks in the balanced state (Vreeswijk, Sompolinsky, 1988), external control of recurrent dynamics strongly depends on the correlations of the input: one might expect that driving all neurons with a common input helps to dominate network dynamics. Surprisingly, we find that the network is far easier to control with independent inputs into each neuron. We discover that this discrepancy can be explained by the dynamic cancellation of external input by recurrent feedback - a phenomenon previously described in dense binary networks in the 'balanced state' (Renart, et al., 2010). This cancellation impedes the suppression of chaos in the network, since only residual fluctuations of order O(1/sqrt(K)) - where K scales both excitatory and inhibitory currents - contribute to the control of the recurrent dynamics. In contrast, for inputs independent across neurons, no cancellation occurs, and a weaker external input is sufficient for control. We present a scaling analysis of the critical external modulation strength necessary to suppress chaotic activity. As predicted by mean field theory (Renart, et al., 2010), the critical input modulation scales as sqrt(K) for correlated input, whereas it is independent of network size for independent external input. In summary, we identified a novel link between excitatory-inhibitory circuit dynamics and chaos. This can help to harness the computational capabilities of recurrent circuits for plasticity and learning of stable trajectories. Specifically, we predict that learning in balanced networks is facilitated by uncorrelated inputs.

III-38. Back-propagating errors with burst coding

Richard Naud[1] RNAUD@UOTTAWA.CA
Friedemann Zenke[2] FRIEDEMANN.ZENKE@CNCB.OX.AC.UK
Jean-Claude Beique[1] JBEIQUE@UOTTAWA.CA

[1]University of Ottawa
[2]University of Oxford

The remarkable properties of learning with a back-propagation of errors compels the question How can this simple but effective algorithm be used in the brain? Much recent interests has focused on biological approximations the symmetry of connections, but often using unrealistic learning protocols based on the separation of a transmit and a learning phase. Using computer simulations constrained by data, we have explored the hypothesis that these two phases are taking place simultaneously when active dendrites give rise to spike-burst multiplexing. In our simulations, tonic spikes are travelling up the hierarchy through synapses with short-term depression. This propagation is unaffected by concomitant top-down information. Simultaneously, a prediction error artificially generated at the top of the hierarchy is propagated down through synapses with short-term facilitation onto dendritic compartments. We find that the correspondance between the spiking simulations and the backprop algorithm is improved by including feedback inhibition from somatostatin- and parvalbumin-positive cells cells. Furthermore, the activity of VIP-positive cells enacts an instantaneous control of the learning rate. Back-propagating errors with burst coding gives rise to three important predictions. Firstly, it predicts the triplet and frequency dependence of spike-timing dependent plasticity (STDP). Secondly, it predicts a dependence of intrinsic plasticity on the relative proportion of bursts. Lastly, it ascribes a strong and weak burst rate to false negative error and false positive errors, respectively.

III-39. Circuit mechanisms of persistent activity in the primate oculomotor system

Eric Hart EHART004@GMAIL.COM
Alex Huk HUK@UTEXAS.EDU

The University of Texas at Austin

Persistent neural activity in cortical neurons is classically viewed as a neural substrate of working memory. Neurons in both the parietal cortex (area LIP) and the prefrontal cortex (FEF) of monkeys show sustained responses during delayed saccade tasks. These regions share bidirectional connections in anatomy, but detailed functional assessments of how these neurons and areas come to exhibit persistent activity have yet to inform quantitative and theoretical models. We therefore conducted simultaneous multi-neuron recordings in both areas in the brain of a macaque monkey performing a standard memory-guided saccade task. We analyzed 176 single-unit and multi-unit clusters in both areas (99 in LIP, 77 in FEF) using a generalized linear model (GLM). In addition to characterizing the impact of task variables (such as the visual target and the saccade) on neural responses, the

GLM supports tests of whether single-trial fluctuations in a particular neuron's spike train impacted responses in other neurons, both within (LIP-LIP, FEF-FEF) and across (LIP-to-FEF, FEF-to-LIP) areas. This revealed strong coupling within both areas. Such pairs usually interacted over time scales spanning ~5-30 msec. Clear cross-area coupling was also observed on similar time scales, but with lesser magnitudes than within-area coupling. Both FEF-to-LIP and LIP-to-FEF directional interactions were detected, with a small but notable asymmetry whereby FEF activity exerted a larger influence on LIP spiking and acted on longer time scales than the complementary direction. Thus, although both areas FEF and LIP exhibit sustained responses and have functional connections, the slow temporal integration evident in their persistent activity is mediated by fast interneuronal interactions both within and across areas. The time scale and directionality of these coupling results place firm constraints on realistic models of persistent activity.

III-40. A Bayesian boolean decoder for robust odor recognition

Lijun Zhang LIJUNZHANG@WUSTL.EDU
Srinath Nizampatnam SRINATH@WUSTL.EDU
Debajit Saha DSAHA@WUSTL.EDU
Rishabh Chandak RISHABH.CHANDAK@WUSTL.EDU
Baranidharan Raman BARANI@WUSTL.EDU

Washington University in St. Louis

Sensory stimuli evoke spatially and temporally patterned complex neural responses that are hypothesized to encode the stimulus identity. However, our experimental results show that both spatial and temporal features of odor-evoked responses can vary significantly across different encounters of the same stimuli. This observation raises an important question: how are sensory stimuli robustly recognized. We examined this issue in this study. We envisioned a flexible decoder that performed a simple logical/Boolean operation: OR-of-ANDs or disjunction of conjunctions. The flexible decoder summed and thresholded the contributions of all neurons activated by solitary introductions of the target stimulus and disregarded the contributions of all other neurons. By appropriately choosing the threshold of the decoding neuron, information distributed in variable subsets of neurons can be used to achieve robust recognition. This simple neural decoding solution provides an adequate tradeoff between the representational stability needed for robust recognition and flexibility needed for adaptive sensory computations. We next wondered if our approach can be generalized into a simple Boolean Neural Network to achieve robust stimulus recognition for general machine learning systems. The Boolean neural network was built on a Bayesian treatment of linear logistic regression with Bernoulli priors on weights. We validated the generality of this approach by applying the proposed technique for pattern recognition on our olfaction dataset as well as other standard machine learning datasets such as MNIST.

III-41. Visual and linguistic semantic representations are aligned at the boundary of human visual cortex

Sara Popham[1] SPOPHAM@BERKELEY.EDU
Alexander Huth[2] ALEX.HUTH@GMAIL.COM
Natalia Bilenko[1] NBILENKO@GMAIL.COM
Jack Gallant[1] GALLANT@BERKELEY.EDU

[1]University of California, Berkeley
[2]The University of Texas at Austin

Humans can learn about the world both from direct sensory experience and indirectly through language. We wanted to know how the brain systems that represent semantic knowledge gained through both of these means relate to each other in their cortical locations. Here, we look specifically at visual and linguistic representations

around anterior occipital cortex. To address this issue we used fMRI to record brain activity while six subjects watched natural movies (120 minutes) and while they listened to narrative stories (129 minutes). Object and action categories were labeled in the movies (one second resolution), and the stories were transcribed and aligned (single word resolution). Movie labels and story words were projected into a common 985-dimensional semantic word embedding space constructed using a separate text corpus. Separate voxel-wise encoding models were then estimated for each modality using regularized linear regression, and the fit models were projected onto individual cortical flat maps. Evaluation of the cortical maps reveals a clear boundary around anterior occipital cortex. Voxels posterior to this boundary respond to movies but not stories, voxels anterior to the boundary respond to stories but not movies, and voxels along the boundary respond to both. The semantic categories represented along the boundary are diverse. However, for nearly any visual category that is represented on the posterior side of the boundary, the same linguistic category is represented on the anterior side of the boundary. That is, there is a close spatial correspondence between the semantic categories represented in movies and those represented in stories. These findings suggest that visual cortex is only part of a contiguous multimodal cortical map, and its boundary is the functional and anatomical interface between vision and language.

III-42. Decoding neural population activity within a latent variable framework

Matthew Whiteway[1] WHIT8022@UMD.EDU
Bruno Averbeck[2] AVERBECKBB@MAIL.NIH.GOV
Ramon Bartolo[2] RAMON.BARTOLOOROZCO@NIH.GOV
Daniel Butts[1] DAB@UMD.EDU

[1]University of Maryland
[2]National Institute of Mental Health

Neural populations produce complex patterns of activity, which play an important role in the neural computations that underlie perception, cognition and behavior. Decoding these patterns of activity to extract task-relevant information can be hindered by the presence of correlations between pairs of neurons on nominally identical trials (“noise correlations”). Here we present a perspective on this decoding problem that employs concepts from latent variable models of population activity, whereby we assume that noise correlations (or at least some portion of them) arise from a small number of latent variables that drive the activity of a larger number of neurons. These latent variables could, for example, capture the effect of internally generated signals that are not experimentally controlled. The simple structure of the noise correlations induced by this assumption motivates the development of a new class of decoding algorithms. We apply both linear and nonlinear versions of this latent variable approach to large populations of hundreds of simultaneously recorded neurons in prefrontal cortex during a visually-guided saccade task. Our approach can dramatically outperform standard decoding algorithms such as logistic regression, and the nonlinear version can in principle outperform optimal linear decoding. This work offers a practical means by which large populations of neurons can be efficiently decoded, and suggests the importance of considering latent variable methods in understanding the effects of structured variability on neural population coding.

III-43. Neural signature of Bayesian interval timing in dorsomedial frontal cortex

Hansem Sohn HANSEM@MIT.EDU
Devika Narain DNARAIN@MIT.EDU
Mehrdad Jazayeri MJAZ@MIT.EDU

Massachusetts Institute of Technology

Bayesian inference has emerged as a unified framework for understanding the ubiquitous role of prior statistics and beliefs in perception, cognition and action. However, how the brain encodes prior information and uses it

for optimal inference is not known. To tackle this question, we analyzed neural activity in the dorsomedial frontal cortex of monkeys during a Bayesian timing task. Animals measured a sample interval between two flashes (Ready and Set), and reproduced it afterwards by making a saccade (Go) to a visual target. Sample intervals were drawn from one of two overlapping uniform prior distributions. The two distributions alternated in a blocked fashion and were cued on every trial. Animals learned to perform the task and were able to switch rapidly between the two prior conditions. Produced intervals increased monotonically with the sample interval in both conditions and exhibited prior-dependent regression to the mean that was accurately captured by a Bayesian model. During the measurement epoch, neural trajectories associated with the passage of time featured a distinct rotational segment that was temporally matched to the prior statistics and was therefore displaced across the two prior conditions. This rotation afforded a warped time axis along a dimension that the system used to represent a Bayesian estimate of elapsed time. Activity along this warped time axis served as the initial condition during reproduction, and adjusted the speed of the ensuing neural trajectories in accordance with Bayes-optimal behavior. Analysis of a recurrent neural network model trained to emulate monkeys' Bayesian behavior during the task revealed that the observed prior-dependent rotations were likely governed by attractor dynamics. These results indicate that prior statistics exert their influence on behavior by shaping cortical latent dynamics.

III-44. Hierarchical organization of neural states in freely behaving rodents

Dyuti Bhattacharya[1] DYUTIBHATTACHARYA@FAS.HARVARD.EDU
Daniel Lee[2] DDLEE@SEAS.UPENN.EDU
Haim Sompolinsky[1,3] HAIM@FIZ.HUJI.AC.IL
Ashesh Dhawale[1] DHAWALE@FAS.HARVARD.EDU
Bence Olveczky[1] OLVECZKY@FAS.HARVARD.EDU

[1]Harvard University
[2]University of Pennsylvania
[3]vard University/ Hebrew University of Jerusalem

The ultimate goal of systems neuroscience is to understand how neural activity generates behavior. Behavior evolves over a range of timescales, from elementary movements, to complex sequences of movements and behavioral states such as grooming, eating and exploration. To understand the neural mechanisms that give rise to such hierarchical behavioral dynamics, we recorded neural activity from large numbers of neurons over weeks and months in freely behaving rats, while also carefully monitoring and classifying the animals' behavior (Dhawale et al. 2017). Though we could often follow the activity of single units over days and weeks, the identity of the recorded neurons also changed significantly over the course of the experiments. To reliably track changes in neural states over longer time periods, we developed a 'stitching algorithm' that aligns the neural states over different days. This allowed us to describe neural dynamics as transitions between recurring states (defined by a Hidden Markov Model), a description that can extend over weeks of continuous recordings. We found the mapping between neural and behavioral states to be robust across days. We also found a strong correlation between the similarity of two neural states and the probability of transitioning between them. This was true for both quiescent and active periods, suggesting that neural dynamics switches preferentially between similar states. Further analysis of the neural state sequences revealed long time-scale dependencies, hinting at the presence of temporal hierarchies. To explore this further, we modeled the neural dynamics as a hierarchical HMM. We found that the multi-scale temporal structure in our data was well captured by the presence of longer 'super-states' that modulate the transition dynamics of shorter neural states. Overall, these results suggest that patterns of behavioral state transitions are strongly constrained by the structure and hierarchical nature of neural state space.

III-45. Temporal structure of hippocampal activity during offline periods

Kourosh Maboudi Ashmankamachali[1] KOUROSH@UWM.EDU
Etienne Ackermann[2] ETIENNE.ACKERMANN@RICE.EDU
Caleb Kemere[2] CALEB.KEMERE@RICE.EDU
Kamran Diba[3,1] KDIBA@UMICH.EDU

[1]University of Wisconsin-Milwaukee
[2]Rice University
[3]University of Michigan

Sharp wave/ripple complexes and associated population bursts form a major mode of hippocampal activity. Evidence indicates that the sequential activation of neurons during the population bursts facilitates plasticity that is needed for formation of spatial and episodic memory. Detection of these sequences is therefore beneficial to understanding the formation of hippocampus-dependent memories. Common methods of sequence detection use templates derived from the ordered activity of neurons during behavior. However, the utility of these templates is limited by the necessity of recording population neuronal activity through multiple repetitive trials of the same spatial task. Under many mnemonic paradigms, learning occurs quickly over very few trials and activity templates are consequently not reliable or limited to a small set of specific sequences. Other vital memories or hard-wired components of the network may pass undetected. To tackle this problem, we recently trained hidden Markov models (HMMs) to capture sequential hippocampal neuronal patterns during quiet waking periods, independent of behavioral templates. States, transitions and correspondence to neuronal firing were learned directly from data. We found that states represent distinct zones in the environment, similar to place fields. The population vectors for states and the transition between them were both highly sparse. Graphs corresponding to state transitions were dominated by long paths containing sequentially connected states, representing positions in the environment. Finally, HMMs were able to recognize many of the same sequences detected using the more common Bayesian decoding method of replay detection. However, the HMM-based method appeared to be less vulnerable to false positives compared to the Bayesian decoding method. In the future, we aim to extend the application to other data, such as during sleep or anesthesia when there may be less direct correspondences between neuronal sequences and animals' overt behavior.

III-46. Complex adaptive internal model subserves perceptual sequential decision making

Adam Koblinger KOBLINGER_ADAM@PHD.CEU.EDU
Jozsef Arato JOZSEF.ARATO@GMAIL.COM
Jozsef Fiser FISERJ@CEU.EDU

Central European University

Despite recent findings of sequential effects in perceptual serial decision making (SDM) (Chopin & Mamassian 2012; Fischer & Whitney 2014), SDM is typically investigated under the assumption that the decisions in the sequence are independent or at most, are influenced by a few previous trials. We set out to identify the true underlying internal model of event statistics that drives decision in SDM by investigating and modeling a set of novel sequential 2AFC visual discrimination tasks by humans and rats. Participants solved the same decision task across trials, but experienced one shift in baseline appearance probabilities of noisy stimuli during the experiment. We found non-trivial interactions between short- and equally strong long-term effects guiding evidence accumulation and decisions in such SDM. These interactions could elicit paradoxical and long-lasting net serial effects, for example, a counterintuitive negative decision bias towards the recently less frequent element. Our findings cannot be explained by previous models of SDM that either assume a sequential integration of prior evidence, or presume an implicit compensation of discrepancies between recent and long-term summary statistics, or adjust learning rates of those statistics at change points. To provide a normative explanation for the empirical data, we developed a hierarchical Bayesian model that could simultaneously represent the priors over the ap-

pearance frequencies and a potentially non-uniform noise model over the different stimulus identities. The results of simulations with the model suggest that humans are more disposed to readjust their noise model instead of updating their priors on appearance probabilities when they observe sudden shifts in the input statistics of stimuli. In general, regardless of the simplicity of the decision task, humans automatically utilize a complex internal model during SDM and adaptively alter various components of this model when detecting sudden shifts in the conditions of the task.

III-47. Modeling the sensorimotor computations that direct orientation behavior through active sampling

Ajinkya Deogade[1] AJINKYA.DEOGADE@LIFESCI.UCSB.EDU
Elena Knoche[2] ELENA.KNOCHE@GMAIL.COM
Daniel Malagarriga[2] DANIEL.MALAGARRIGA@GMAIL.COM
Eric T. Trautman[3] TRAUTMANE@JANELIA.HHMI.ORG
Vivek Jayaraman[3] VIVEK@JANELIA.HHMI.ORG
Matthieu Louis[1] MATTHIEU.LOUIS@LIFESCI.UCSB.EDU

[1]University of California, Santa Barbara
[2]Centre for Genomic Regulation
[3]HHMI Janelia Research Campus

Chemotaxis is a powerful paradigm to study how dynamical sensory input is converted into directed motor output. Drosophila larvae navigate odor gradients by alternating phases of forward locomotion (runs) with reorientation maneuvers (turns). The sensory information encoded by a single functional olfactory sensory neuron (OSN) is sufficient to control larval chemotaxis. Recently, we showed that negative olfactory gradients inhibit OSN activity, which in turn promotes stopping behavior. By contrast, positive gradients promote strong OSN activity and sustained running behavior. Here, we combine optogenetics and high-resolution closed-loop tracking to decipher the sensorimotor mechanisms controlling navigation in virtual olfactory realities. Reorientation is directed by an active sampling routine relying on lateral head sweeps (head casts) analogous to sniffing in vertebrates. During runs, low-amplitude head casts (“run- casts”) are locked with the rhythmic pattern of peristalsis. Upon interruption of peristalsis, larvae engage in head casts of wider amplitudes (“stop-casts”). Replaying the odor experience of the animal during electrophysiological recordings in immobilized larvae, we show that run-casts likely mediate changes in OSN activity that are sufficient to bias the direction of the run towards positive gradients. During stop-casts, the amplitude of head casts is modulated by OSN activity: the detection of positive gradients maintains casting while negative gradients interrupt it. The larval olfactory system is capable of perceiving and responding to transient changes in OSN activity on a timescale shorter than 50 ms. By integrating a quantitative model of the dynamical encoding of odors with the sensorimotor control of head-casting behavior and forward locomotion, we reproduce essential aspects of larval chemotaxis in agent-based simulations. Our work presents a fine-grained multi-scale model of the sensorimotor loop that underlies the acquisition of sensory information through active movements of the body.

III-48. Population voltage imaging in behaving animals reveals dynamics of sensorimotor decisions

Takashi Kawashima[1] KAWASHIMAT@JANELIA.HHMI.ORG
Ahmed Abdelfattah[1] ABDELFATTAHA@JANELIA.HHMI.ORG
Jonathan Grimm[1] GRIMMJ@JANELIA.HHMI.ORG
Johannes Friedrich[2] JFRIEDRICH@FLATIRONINSTITUTE.ORG
Liam Paninski[3] LIAM@STAT.COLUMBIA.EDU
Luke Lavis[1] LAVISL@JANELIA.HHMI.ORG
Eric Schreiter[1] SCHREITERE@JANELIA.HHMI.ORG
Misha Ahrens[1] AHRENSM@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]Flatiron Institute
[3]Columbia University

Brains integrate sensory information and make decisions about which behavior to execute through brainwide communication at the millisecond timescale. To fully track transformations from sensory input to behavior, it is necessary to record spikes from large numbers of neurons across multiple brain areas. Here we introduce technology that makes this possible, through light-sheet imaging in behaving larval zebrafish with a novel chemigenetic voltage indicator, Voltron. Voltron has a hybrid design of protein voltage sensing domain and a photostable chemical dye, which tolerates photobleaching over extended periods of time during high-speed imaging. We imaged spike responses in dozens of glutamatergic cells in a motor-related midbrain nucleus in zebrafish swimming in a virtual environment. We identified different types of spiking patterns within this nucleus. In one population of neurons, spike rates increased about one second prior to swimming, suggesting the presence of preparatory populations for motor execution. A second, non-overlapping population of neurons showed spiking synchronized to motor output. Further analysis showed that the spike timing of this second population is delayed relative to spinal output by several milliseconds, suggesting that they receive efference copy signals from motor control circuits; but remarkably, those same cells show subthreshold depolarization prior to swim onset. This co-existence of pre- and post-motor spike signals in a common population in a single brain area suggests that both are integrated to achieve stable generation of behavior. This study demonstrates the analysis of millisecond-timescale neural computations across the brain during behavior using scalable optical technology.

III-49. Revealing multiple timescales of structure in larval zebrafish behavior

Robert Johnson[1] ROBERTEVANJOHNSON@FAS.HARVARD.EDU
Scott Linderman[2] SCOTT.LINDERMAN@GMAIL.COM
Thomas [3] THOMAS.PANIER@GMAIL.COM
Caroline Wee[3] CAROLINEWEE@GMAIL.COM
Erin Song[3] SONG6@FAS.HARVARD.EDU
Kristian Herrera[3] KHERRERA01@FAS.HARVARD.EDU
Andrew Miller[3] ACM@SEAS.HARVARD.EDU
Florian Engert[3] FLORIAN@MCB.HARVARD.EDU

[1]Duke University
[2]Columbia University
[3]Harvard University

Predictive models of an animal's behavior quantify and characterize the output space of its nervous system. This space is huge for wild animals and is often constrained in laboratory environments to limit its complexity and the volume of data needed to evaluate its structure. The study of natural behavior is further complicated by the challenge of temporally segmenting behavioral sequences into constituent elements. Zebrafish larvae are well suited for study because their behavior is naturally discrete, i.e. they swim in punctuated bouts. We reveal patterns

in larval zebrafish behavior by monitoring individual larvae swimming freely in a large arena containing abundant unicellular food resources. We collected 40 hours of high-resolution video with a tracking camera programmed to move above the fish, recording its postural dynamics and interaction with prey objects. We categorized swim action types and constructed a probabilistic model—a marked point process—to predict how these actions are deployed given the internal hunger-state of the fish, its behavioral history, and the current environmental input. The distribution of swim actions is left-right symmetric at the population level and we exploit this symmetry to simplify behavioral classification and construct more robust models. We find that hunger increases the likelihood of transitioning from exploratory to hunting behavior and promotes shorter intervals between consecutive exploratory actions. We construct symmetric neural networks to interpret how statistics of these natural scenes influence action selection. Our probabilistic modeling approach has enabled us to compare the predictive power of different features (e.g. external factors like prey locations to internal factors like previous action) and sample from multi-timescale models to generate realistic behavior of a virtual fish in simulated environments. Our work reveals previously unknown patterns in larval zebrafish behavior and can serve to generate hypotheses about possible neural implementations of these behavioral algorithms.

III-50. Unsupervised learning of sequential activity with temporally asymmetric Hebbian learning rules

Maxwell Gillett[1] MAX.GILLETT@GMAIL.COM
Nicolas Brunel[1] NICOLAS.BRUNEL@DUKE.EDU
Ulises Pereira[2] ULISES@UCHICAGO.EDU

[1]Duke University
[2]University of Chicago

Sequential activity is observed in neuronal circuits across species, neural structures, and cognitive tasks. It is believed that sequences can arise from unsupervised learning processes, but to what extent biologically plausible unsupervised learning rules can organize this activity remains unknown. Moreover, there exists a gap in the theoretical understanding of the conditions required for this learning, as well as the storage capacity of these rules. In this work we investigate unsupervised temporally asymmetric Hebbian rules in recurrent rate networks, and develop a theory of the transient sequential activity observed after learning. These rules transform a sequence of random input patterns into synaptic weight updates. After learning we find that sequential activity is reflected in the transient overlap of network activity with each of the stored input patterns, and that multiple sequences can be stored. Each sequence can be activated independently by providing brief input corresponding to the first pattern in the sequence, and activity is robust to perturbations. Using mean-field theory, we find a low-dimensional description of the network dynamics in which N degrees of freedom are reduced to an effective delay line system with only O(1) variables. We use this reduced system to derive a statistical description of network activity and to compute the sequential capacity of these networks, determining the maximal number of fixed-length sequences that can be stored as a function of network size. We show that the capacity grows linearly with network size, comparable to that found in networks storing fixed-point attractors. Our work yields insight into the interplay between learning rules and network parameters in determining the presence and robustness of this activity.

III-51. Automated high-throughput cellular resolution neural circuit mapping with online experimental design

Ben Shababo[1] SHABABO@BERKELEY.EDU
Karl Kilborn[2] KARL@INTELLIGENT-IMAGING.COM
Xinyi Deng[3] XD2188@COLUMBIA.EDU
Johannes Friedrich[4] JFRIEDRICH@FLATIRONINSTITUTE.ORG
Hillel Adesnik[1] HADESNIK@BERKELEY.EDU
Liam Paninski[3] LIAM@STAT.COLUMBIA.EDU
Shizhe Chen[3] SHIZHE.CHEN@GMAIL.COM

[1]University of California, Berkeley
[2]Intelligent Imaging Innovations (3i)
[3]Columbia University
[4]Flatiron Institute

Circuit-mapping experiments combining whole-cell electrophysiology with two-photon optical stimulation of potentially presynaptic neurons have produced rich data on cell type-specific monosynaptic connectivity of neural circuits. However, mapping large volumes of densely-packed neurons (e.g. cortical excitatory neurons) at cellular resolution with two-photon optogenetics has proven challenging because of two main problems: 1) stimulation sensitivity and resolution varies across cells due to differential opsin expression and intrinsic excitability, making the precise localization of connected neurons difficult, and 2) experimental time is severely limited compared to the number of potential connections to map. We present a method which overcomes these limitations using statistical modeling, online experimental design, and the combination of a fast, high-potency soma-targeted opsin (st-ChroME) with computer generated multi-target holography. To infer posterior distributions for the probability of synaptic transmission and opsin expression level of potentially connected neurons, we fit a generative model with three main components: a neural response model which predicts presynaptic spike rates given the power and location of stimulation targets, a connectivity model which filters presynaptic spike rates into a postsynaptic event rate, and a postsynaptic model which converts the postsynaptic event rate into a voltage-clamp observation. To improve mapping efficiency, we implement a closed-loop parallel computational system which designs batches of stimulation targets online based on fast stochastic variational inference of these posteriors. Specifically, stimulation design for each neuron automatically switches between coarser (multi-target) and finer (single-target) protocols depending on the posteriors. This approach allocates high-resolution single-target stimuli only at locations of evoked connections and ambiguous responses. Furthermore, our experimental system allows us to collect data at 20 trials per second for a large portion of experimental time while analyzing data in the background. We demonstrate the efficacy of our method in vitro and in vivo by mapping connectivity (including short-term plasticity) in mouse cortex.

III-52. Inferring pre- and post-synaptic activity from dendritic calcium imaging

Jinyao Yan[1] YANJ11@JANELIA.HHMI.ORG
Aaron Kerlin[1] KERLINA@JANELIA.HHMI.ORG
Laurence Aitchison[2] LAURENCE.AITCHISON@GMAIL.COM
Karel Svoboda[1] SVOBODAK@JANELIA.HHMI.ORG
Srini Turaga[1] TURAGAS@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]University of Cambridge

In vivo calcium imaging can be used to probe the functional organization of synaptic activity across the dendritic arbor. Synaptic input onto spines can produce calcium transients that are largely isolated to the spine head. However, in otherwise unperturbed cells, the back-propagating action potential (bAP) also contributes strongly to the change in fluorescence measured at individual spines. To address this problem, we propose a statistical model

to separate these sources and infer the probability of both pre- and post-synaptic spiking activity. This model is a simplified nonlinear approximation of the biophysical processes by which synaptic input and the bAP contribute to the fluorescent measurements at different sites. We use the framework of variational autoencoders (VAE) – a recent advance in machine learning – training a deep neural network (DNN) as part of the VAE to efficiently infer an approximate posterior distribution over spike trains from fluorescence traces. Our training procedure jointly optimizes parameters of the generative model and the recognition network in order to maximize a lower bound on the log likelihood of the observed data. In developing these methods for dendritic inference, we have also extended and improved our previous work using VAEs for somatic spike inference. These methods are a crucial step towards understanding the transformation of dendritic inputs to somatic spike output in vivo.

III-53. Structured features in connectivity between olfactory regions

Hamza Giaffar[1] HGIAFFAR@CSHL.EDU
Sergey Shuvaev[1] SSHUVAEV@CSHL.EDU
Daniel Kepple[1] DKEPPLE@CSHL.EDU
Alexei Koulakov[1] KOULAKOV@CSHL.EDU
Dmitry Rinberg[2] DMITRY.RINBERG@NYUMC.ORG

[1]Cold Spring Harbor Laboratory
[2]New York University

The computation of invariant features in stimuli is a key function of sensory systems. The olfactory system, for instance, is faced with the task of generating a representation of odor identity that is invariant with respect to substantial variations in concentration, timing and background. The primacy code has recently been proposed as a simple and robust solution to this complex computational task (Wilson 2017). According to this scheme, the identity of an odorant is encoded by the activity of a small (primacy) set of highest affinity olfactory receptor types (ORs). The primacy set of ORs respond with the shortest latency out of all OR responses to a given odourant and their activation is invariant with respect to concentration. A variety of network architectures can process a primacy code, from simple feedforward (Wilson 2017) to more complex recurrent dual networks (Kepple 2017). Recent theoretical work reveals some of the subtle implications of structure vs. randomness in the connectivity of similar feedforward models of sensory systems (Babadi 2014), suggesting possible benefits of structured synapses that encode the high order correlations implied by the primacy model and found in OR-ligand response data. For simple feedforward networks, we find that (sparse binary) structured connectivity that reflects the primacy based associations between ORs improves performance in a classification task relative to random synapses, when stimuli are drawn from a low dimensional manifold. We examine the sensitivity of PCA for detecting primacy like connectivity in current datasets, finding the technique to be insensitive to high levels of even simple correlation structure when connectivity data is heavily subsampled. This result reopens questions regarding the extent of order in olfactory connectivity. Inspired by the visual system, we propose a simple Hebbian mechanism to account for the formation of this type of structure during organismal development. Simulation studies provide strong support for the effectiveness of this mechanism in generating the appropriate structure for efficiently processing a primacy code.

III-54. Pyramidal cell-interneuron circuit architecture and dynamics in hippocampal networks

Daniel English[1] ENGLISHDANIEL@GMAIL.COM
Talfan Evans[2] TALFAN.EVANS.13@UCL.AC.UK
Sam McKenzie[1] SAM.A.MCKENZIE@GMAIL.COM
Kanghwan Kim[3] KHANKIM@UMICH.EDU
Euisik Yoon[3] ESYOON@UMICH.EDU
Gyorgy Buzsaki[1] GYORGY.BUZSAKI@NYUMC.ORG

[1]New York University
[2]University College London
[3]University of Michigan

Hippocampal networks exhibit internally generated representations of the external world, including those of space and time. The organization of these representations at the synaptic and circuit level is believed to involve the transient binding of principal neurons into assemblies, whose sequential evolution involves competitive interactions between them. In the CA1 region, principal neurons primarily interact with one another via GABAergic interneuron-mediated feedback inhibition. Thus, the recruitment of interneurons by pyramidal neurons is central to hippocampal function. However, the excitatory control by principal neurons of interneurons is poorly understood due to the difficulty of studying synaptic connectivity in vivo. Here, we inferred such connectivity through analysis of spike timing of pre- and postsynaptic neurons, and validated this inference using juxtacellular and optogenetic control of presynaptic spikes in awake-behaving mice. We observed that neighboring CA1 neurons had stronger connections, and that superficial pyramidal cells projected more to deep interneurons. Connection probability and strength were skewed, with a minority of highly connected hubs. Divergent presynaptic connections led to synchrony between interneurons. Synchrony of convergent presynaptic inputs boosted postsynaptic drive. Presynaptic firing frequency was read out by postsynaptic neurons through short-term depression and facilitation, with individual pyramidal cell-interneuron pairs displaying a diversity of spike transmission filters. Additionally, spike transmission was strongly modulated by prior spike timing of the postsynaptic cell. These results bridge anatomical structure with physiological function, and provide a circuit-level framework for investigating cell assemblies.

III-55. Attentional modulation of neural covariability in a distributed circuit-based population model

Matthew Getz MATTPGETZ@GMAIL.COM
Jeff Dunworth JBD20@PITT.EDU
Marlene Cohen MARLENERCOHEN@GMAIL.COM
Brent Doiron BDOIRON@PITT.EDU
Chengcheng Huang CHENGCHENGHUANG11@GMAIL.COM

University of Pittsburgh

Relating the structure of cortical circuits to their processing of input stimuli is a long-standing problem at the interface between circuits and systems neuroscience. One important observation in building circuit models is that cortical response often depends upon cognitive state, thereby providing two conditions from which to constrain a single circuit model [Doiron et al. 2016]. In primate visual cortex, directed attention amplifies the magnitude and sensitivity of trial averaged firing rates, as well as reduces the trial-to-trial shared variability between neurons [Cohen and Maunsell 2009]. These constraints recently allowed us to show how top-down modulation of inhibition in a recurrent coupled excitatory and inhibitory model of cortex is a reasonable model of visual attention [Kanashiro et al. 2017]. However, our model treated cortex as a bulk network, ignoring any distributed stimulus tuning across neurons. Since the experiments used to probe attention involved both spatial and feature based attention [Cohen and Maunsell 2011] this shortcoming prevents a deep analysis of how attention affects population-wide

responses. To rectify this we extend our model to include stimulus tuning and incorporate recurrent wiring and attentional modulation that is distributed over tuning space. We show that if attentional modulation targets the network with a precision that is above the scale of stimulus tuning, the shared variability between similarly tuned neurons increases with attention, in disagreement with experiment. However, if attentional modulation is broadly distributed, then the model shows an overall reduction of population-wide shared variability with attention. Our work exposes how the relative distribution of top-down and bottom-up signals is a critical feature, often unstudied, that determines how modulation affects the response of recurrent circuits.

III-56. Complementary direct and indirect pathway activity encodes sub-second 3D pose dynamics in striatum

Jeffrey Markowitz[1] JEFFREY_MARKOWITZ@HMS.HARVARD.EDU
Winthrop Gillis[1] WGILLIS@G.HARVARD.EDU
Celia Beron[1] CELIA_BERON@G.HARVARD.EDU
Shay Neufeld[1] SHAYNEUFELD@FAS.HARVARD.EDU
Keira Robertson[1] KEIRA_ROBERTSON@HMS.HARVARD.EDU
Minsuk Hyun[1] MINSUKHYUN@FAS.HARVARD.EDU
Emalee Peterson[1] EMALEE_PETERSON@HMS.HARVARD.EDU
Ralph Peterson[1] RALPH_PETERSON@HMS.HARVARD.EDU
Neha Bhagat[1] NEHA_BHAGAT@HMS.HARVARD.EDU
Scott Linderman[2] SCOTT.LINDERMAN@GMAIL.COM
Bernardo Sabatini[1] BERNARDO_SABATINI@HMS.HARVARD.EDU
Sandeep Datta[1] SRDATTA@HMS.HARVARD.EDU

[1]Harvard University
[2]Columbia University

Many naturalistic behaviors are built from modular components that are expressed sequentially. Although corticostriatal circuits have been implicated in action selection and implementation, the neural mechanisms that compose behavior in unrestrained animals are not well understood. Here we simultaneously record neural activity in the direct and indirect pathways of dorsolateral striatum (DLS) as mice spontaneously express action sequences. These experiments reveal that average activity in these pathways is largely correlated, but fast-timescale decorrelations allow both pathways to collaboratively and systematically encode information about sub-second 3D pose dynamics. Action-associated neural activity is modulated in a sequence-dependent manner, and DLS lesions altered the assembly of behavioral sequences. By characterizing naturalistic behavior at neural timescales, these experiments identify a code for elemental 3D pose dynamics built from complementary pathway dynamics, support a role for the DLS in elaborating the grammatical structure of behavior, and suggest models for how actions are sculpted over time.

III-57. Biophysically interpretable, model-free identification of neuronal dynamics

Vignesh Narayanan VIGNESH.NARAYANAN@WUSTL.EDU
Jr-Shin Li JSLI@WUSTL.EDU
ShiNung Ching SHINUNG@WUSTL.EDU

Washington University in St. Louis

Neurons produce a rich repertoire of activity, mediated by ionic channel dynamics that span multiple time-scales. A basic problem in experimental neuroscience involves the data-driven identification of key ionic channel contributors to the overall dynamics of a given neuron. Such a problem is challenging since even in the best cases,

identification relies on (noisy) recordings of membrane potential only. Thus, strict inversion to the constituent channel dynamics is mathematically ill-posed. In this work, we develop a biophysically interpretable, learning-based strategy for performing data-driven neuronal dynamical identification. In particular, we propose two variants of a partially model-free framework to learn and approximate neural dynamics from an observed voltage trajectory in order to identify and recover internal gating variables at different time-scales. Both variants rely on the availability of a library of known ionic gating variables with prescribed kinetics of channel opening/closing. Then, the dynamics of the trans-membrane potential is represented in a parametric form with the conductances and membrane capacitances forming a vector of unknown parameters and the gating variables forming a vector of nonlinear regression functions. By forcing the unknown parameters to move in a direction which guides the estimated voltage trajectory towards the observed trajectory, the identification error is pushed towards zero. Since the gating variables are characterized by stable internal dynamics, by contraction theory [2], convergence of the estimated voltage trajectory implies that the internal variables converge to their true values. By using a biophysical gating variable library, the resultant learned parameters indicate those channels that most contribute to the observed voltage trace. A variant of this approach is also proposed that further relaxes assumptions by using generic basis functions within a multi-layer learning network to represent the known channel library. The efficacy of this approach is demonstrated on a variety of neuronal activity patterns spanning several time-scales.

III-58. Three-factor embedded learning on neuromorphic systems

Georgios Detorakis[1] GDETORAK@UCI.EDU
Travis Bartley[1] TBARTLEY@UCI.EDU
Roman Parise[1] PARISEROMAN@GMAIL.COM
Sadique Sheik[2] SSHEIK@UCSD.EDU
Charles Augustine[3] CHARLES.AUGUSTINE@INTEL.COM
Somnath Paul[3] SOMNATH.PAUL@INTEL.COM
Bruno Pedroni[2] BPEDRONI@ENG.UCSD.EDU
Nikil Dutt[1] DUTT@ICS.UCI.EDU
Jeffrey Krichmar[1] JKRICHMA@UCI.EDU
Gert Cauwenberghs[2] GERT@UCSD.EDU
Emre Neftci[1] ENEFTCI@UCI.EDU

[1]University of California, Irvine
[2]University of California, San Diego
[3]Intel Corporation

Neuromorphic devices mimic the architecture, the dynamics and the computations of biological brains. Significant progress in the field has led to remarkable capabilities in simulating large-scale neural networks or accelerating computations and undertaking massively parallel and distributed emulations. However, embedded learning on neuromorphic devices is limited or non-existent, and most often based on Hebbian learning. The latter, including pair-wise Spike-Time-Dependent Plasticity take into account the activity of pre- and post-synaptic neurons and changes their synaptic strengths accordingly, but is limited to basic unsupervised settings. Recently, several researchers argued for three-factor learning rules that take into account neuromodulatory signals reflecting extrinsic, task-related factors at the neural or synaptic level. We extend this insight into neuromorphic architecture design for energy-efficient learning in dedicated hardware. We describe the Neural and Synaptic Array Transceiver (NSAT), a neuromorphic architecture and framework facilitating flexible and efficient embedded learning in hardware using a three-factor learning rule that matches the requirements of gradient descent on the neural dynamics. Thanks to this matching, the NSAT supports event-driven supervised, unsupervised and reinforcement learning algorithms, including deep learning, as demonstrated in three different learning tasks: First, we show how a feed-forward spiking neural network learns to classify hand-written digits using a random backpropagation algorithm (supervised learning). Second, we demonstrate how NSAT and its three-factor learning rule can be used in training a spiking restricted Boltzmann machine in a classification task (unsupervised learning). Finally, we illustrate a sequence learning example, where a spiking neural network is trained to distinguish a template pattern from noise using a voltage-based learning rule. We anticipate that this contribution will establish the foundation for a new generation

of multi-purpose neuromorphic devices supporting on-line, learning.

III-59. The cortical clock

Philippa Karoly PKAROLY@STUDENT.UNIMELB.EDU.AU
Ewan S. Nurse E.NURSE@STUDENT.UNIMELB.EDU.AU
David B. Grayden GRAYDEN@UNIMELB.EDU.AU
Mark J. Cook MARKCOOK@UNIMELB.EDU.AU
Dean R. Freestone DEANRF@UNIMELB.EDU.AU

The University of Melbourne

Regardless of the outside environment or any human construct of time, the brain imposes a periodic order on almost every regulatory function of the body. The activity of hormone levels, sleep, body temperature, and metabolism all follow cyclic patterns (Bass, 2012). The molecular machinery that governs these patterns has long been understood (Bargiello et al., 1984). The central timekeeper of the body's clock has been discovered in the suprachiasmatic nucleus of the brain (Moore & Eichler, 1972, Stephan & Zucker, 1972), and the neurobiological function of this timekeeper has been widely studied both in vitro and in vivo. However, directly studying the cortical mechanisms of the body clock in humans is challenging.

We investigated the cortical representation of time using electrocorticography (ECoG) in eight subjects with long-term intracranial recording electrodes (mean recording duration of 2.1 years). Specifically, we evaluated the accuracy of a "cortical clock" including how stable timekeeping is over long periods of time and whether accuracy is affected by environmental factors (such as weekdays vs. weekends when sleep cycles may be different). To evaluate clock accuracy, we trained convolutional neural networks (CNNs) to classify the hour of the day (24 classes) using 30s segments of ECoG. CNNs were trained with 100 days of consecutive data, starting from day 100 of the recording (to allow the signal to settle), then evaluated continuously on all remaining data (from 6 months to 2 years). We found that the brain's clock is astonishingly accurate—keeping track of the hour with high accuracy (4 to 6 times greater than chance). This precise representation of time was not degraded on weekends. Our results provide evidence of a clock mechanism that is widespread across cortex and stable over many years, demonstrating that there is much still to learn about cerebral circadian cycles.

III-60. The virtual rat: Building a computational model of the rodent whisker trigeminal system

Chengxu Zhuang[1] CHENGXUZ@STANFORD.EDU
Nadina Zweifel[2] NADINAZWEIFEL2022@U.NORTHWESTERN.EDU
Jonas Kubilius[3] QBILIUS@MIT.EDU
Mitra Hartmann[2] HARTMANN@NORTHWESTERN.EDU
Daniel Yamins[1] YAMINS@STANFORD.EDU

[1]Stanford University
[2]Northwestern University
[3]Massachusetts Institute of Technology

Rodents "see" the environment mostly through their whiskers. The exquisitely sensitive, actively-controllable whisker array receives raw sensory data in the form of mechanical signals. These signals are carried along the trigeminal pathway through a sequence of increasingly complex processing stages, from brainstem to somatosensory ("barrel") cortex. Although many aspects of these processing stages have been characterized by a long history of experimental studies, the computational operations of the whisker-trigeminal system are poorly understood. In the present work, these computations are modelled through a goal-driven deep neural network (DNN) approach. Using a biophysically-realistic whisker array model (see companion poster, Zweifel et al., 2018), we sweep the array across a wide variety of 3D objects in highly-varying poses, angles, and speeds to generate a large dataset. Next, DNNs from several distinct architectural families are trained on this dataset to solve a shape recognition task. Each architectural family is based on a structurally-distinct hypothesis for processing in the whisker-trigeminal system. These hypotheses correspond to different ways in which spatial and temporal information can be integrated. After training, we find that reasonable performance levels on the challenging shape recognition task are only achieved by specific architectures from several families, while most networks perform poorly. Finally, we show that Representational Dissimilarity Matrices (RDMs), a tool for comparing population codes between neural systems, can separate these higher performing networks. And the data for computing RDMs is in a type that could plausibly be collected in an imaging or neurophysiological experiment. Our results are a proof-of-concept that DNN models of the whisker-trigeminal system are potentially within reach.

III-61. Contextual modulation of response variability in primary visual cortex

Dylan Festa DYLAN.FESTA@EINSTEIN.YU.EDU
Selina Solomon SELINA.SOLOMON@GMAIL.COM
Amir Aschner AASCHNER@MAIL.EINSTEIN.YU.EDU
Adam Kohn ADAM.KOHN@EINSTEIN.YU.EDU
Ruben Coen-Cagli RUBEN.COEN-CAGLI@EINSTEIN.YU.EDU

Albert Einstein College of Medicine

Response variability to repeated presentations of the same sensory input is a prominent feature of cortical activity, which can affect the quality of perceptual experience. The structure and mechanisms of such variability have been studied extensively, yet its functional role remains debated. Theory suggests that primary visual cortex (V1) performs approximate probabilistic inference about latent features of the visual input, and variability represents uncertainty about those features (Orban et al, Neuron 2016). However, it is unclear what source of uncertainty could account for supra-Poisson increase in variance as a function of the mean response, a widely observed property of cortical variability (Goris et al, Nat. Neurosci. 2014)

Here we first derive new analytical results identifying the source of uncertainty that could explain that observation. We assume that the inference relies on a generative model of natural images (Coen-Cagli et al, Nat. Neurosci.2015), in which visual inputs are generated as linear combinations of local latent features (oriented edges) collectively multiplied by a global modulator (contrast). We show that inference of the local latents, via marginalization of the modulator, often leads to supra-Poisson variability, consistent with data. Our analysis further shows how inference relies on divisive normalization, clarifying the effects of normalization on response mean and variance. This establishes a precise link between a canonical neural mechanism (normalization), a widespread property of cortical variability, and probabilistic inference.

Building on these results, we extend the framework to generate detailed predictions about how variability depends on the visual input, and test them with recordings in macaque V1. Specifically, the model predicts that stimulation of the receptive field surround affects response mean and variability differently, producing richer effects than the uniform response stabilization reported previously with surround stimulation. Our data provide initial support for these predictions.

Our results indicate a precise, functional link between normalization and supra-Poisson variability, and therefore could explain the structure of variability well beyond V1.

III-62. Constrained matrix factorization methods for denoising and demixing voltage imaging data

E. Kelly Buchanan[1] EKELLBUCHANAN@GMAIL.COM
Johannes Friedrich[2] JFRIEDRICH@FLATIRONINSTITUTE.ORG
Ian Kinsella[1] IAK2119@COLUMBIA.EDU
Patrick Stinson[1] PWS2108@CUMC.COLUMBIA.EDU
Pengcheng Zhou[1] PZ2230@COLUMBIA.EDU
Felipe Gerhard[3] FELIPE.GERHARD@QSTATEBIO.COM
John Ferrante[3] JOHN.FERRANTE@QSTATEBIO.COM
Graham Dempsey[3] GRAHAM.DEMPSEY@QSTATEBIO.COM
Liam Paninski[1] LIAM@STAT.COLUMBIA.EDU

[1]Columbia University
[2]Flatiron Institute
[3]Q-State Biosciences, Inc.

Voltage imaging (VI) at single-neuron resolution promises to be a transformative neurotechnology, and better voltage indicators are rapidly being developed. Specialized methods for analyzing the resulting video data will be critical for fully exploiting this new technology. Constrained matrix factorization (CMF) methods have been successfully applied to a variety of calcium imaging (CI) datasets, but VI data has lower signal-to-noise ratio (SNR) and radically different timescales and signal characteristics. Thus, here we develop new CMF pipelines specialized for VI data. We begin by denoising the raw video data, using a combination of purely spatial, purely temporal, and mixed spatiotemporal denoising approaches. All of these denoisers are highly scalable: by working in small local spatial patches, they scale linearly both in T (number of movie frames) and d (number of pixels); we additionally utilize the fact that shot noise can be spectrally separated from the voltage signals of interest, enabling automated selection of all smoothing parameters. We apply a form of anisotropic spatial Wiener filtering to each movie, then apply a local adaptive form of principal component denoising, and finally apply L1 trend filtering (a form of nonparametric spline denoising based on convex optimization) to the resulting temporal signals. We also experimented with neural network methods for denoising the extracted voltage traces. This combination of approaches denoises the video data dramatically while leaving the signals of interest largely unperturbed. Next we apply CMF methods to the denoised data to demix the spatiotemporal video signal into individual contributions from each neuron visible in the field of view (FOV). We demonstrate good performance on a challenging in vitro dataset with highly synchronous activity driven by optogenetic stimulation. We further find that similar methods are effective on simulated and in vivo CI and VI data.

III-63. Discrete attractors underlie preparatory activity in rodent frontal cortex

Lorenzo Fontolan LORENZO.FONTOLAN@GMAIL.COM
Hidehiko Inagaki INAGAKIH@JANELIA.HHMI.ORG
Karel Svoboda SVOBODAK@JANELIA.HHMI.ORG
Sandro Romani SANDRO.ROMANI@GMAIL.COM

HHMI Janelia Research Campus

Short-term memory (STM) is the ability to temporarily hold stimulus-related information in order to carry out cognitive tasks. At the core of STM maintenance is the question of how to sustain a representation after the stimulus is removed. Persistent, stimulus-dependent neuronal activity has been observed in numerous brain areas during tasks requiring the temporary maintenance of information (1,2,3), either related to external stimuli or motor planning. Given that these persistent states can last seconds, proposed mechanisms to explain the experimental results either rely on single-cell multistability or require circuit-level interactions. In this work, we explored the ability of these mechanisms to account for electrophysiological data and perturbation experiments in the anterior lateral motor cortex (ALM) during a delayed-response task in mice (4,5). Analysis of these data excluded

single-cell multistability and neuronal sequences, where information is passed along each node in a chain of connected neurons. Two classes of network models could potentially account for the observed activity: i) integrator models with a continuum of fixed points; ii) attractor models featuring a discrete set of stable firing patterns. To discriminate between these classes, we generated model predictions, such as expected changes in across-trial fluctuations, and the phase portrait during non-selective optogenetic manipulations. These predictions were tested experimentally. Models in (i) could be ruled out, since their dynamics reflected the integration of the perturbation, which was not observed in the data. Instead, models in (ii), implemented as a bistable attractor network, could accurately reproduce the observed neuronal dynamics, either recovering to the unperturbed trajectory or switching to the other fixed point. Our combined computational and experimental approach reveals the presence of discrete attractors during the preparation of motor actions, which may be a general mechanism underlying STM in cortical circuits.

III-64. Nonrandom sampling of olfactory input to Drosophila mushroom body

Zhihao Zheng ZHENGZ11@JANELIA.HHMI.ORG
Feng Li LIF@JANELIA.HHMI.ORG
J. Scott Lauritzen LAURITZENS@JANELIA.HHMI.ORG
Matthew Nichols NICHOLSM@JANELIA.HHMI.ORG
Corey Fisher FISHERC@JANELIA.HHMI.ORG
Nadiya Sharifi SHARIFIN@JANELIA.HHMI.ORG
Steven Calle-Schuler CALLESCHULERS@JANELIA.HHMI.ORG
Lucia Kmecova LKMECOVA@WINDOWSLIVE.COM
Jawaid Ali JAWAIDALII@JANELIA.HHMI.ORG
Davi Bock BOCKD@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

The mushroom body (MB) of the Drosophila brain is critical for olfactory learning and memory. The MB contains ~2,000 Kenyon cells (KCs) that receive olfactory input in the MB calyx from ~150 projection neurons (PNs). Light microscopy (LM) data pooled across many animals (Murthy et al. 2008, Caron et al. 2013), as well as theoretical arguments (Litwin-Kumar et al. 2017), have suggested that the PN-to-KC synaptic network is completely random. However, the PN-to-KC network has never been mapped within a single animal. We used a whole-brain EM dataset (Zheng et al. 2017) to map the connections of ~125 olfactory PNs to ~15% of the KCs in the MB (~350 total). The mapped KCs included a randomly selected set (~170 total) and a tightly cofasciculated set (~180 total; the 'bundle' KCs). We compared actual connectivity to that predicted by two null models: one in which KC claws select PN boutons completely at random, and one in which claws select randomly from nearest neighbor boutons. We find a community of 12 olfactory PN types provides highly enriched input to bundle KCs in the MB calyx compared to a completely random model. However, this effect persists in the nearest-neighbor permutation, suggesting that a precise and developmentally determined geometry of arbor overlap creates network communities within the PN-to-KC graph. Additional reconstruction and analyses are underway to determine whether these communities reflect a structured olfactory space (Koulakov et al. 2011) used in associative memory formation and recall in the fly MB.

III-65. Theory and physiology of spatial frequency tuning in cortical area MT

Ambarish Pawar[1] AMBARISH@SALK.EDU
Sergei Gepshtein[1] SERGEI@SALK.EDU
Sergey Savel'ev[2] S.SAVELIEV@LBORO.AC.UK
Thomas Albright[1] TOM@SALK.EDU

[1]Salk Institute for Biological Studies
[2]Loughborough University

Background: Sensory neurons in cerebral cortex are characterized by their selectivity to stimulation. This selectivity was originally viewed as a stable property of individual neurons, later challenged by the evidence that selectivity depends on stimulus context and history of stimulation. Here we use theoretical and physiological methods to investigate how neuronal selectivity arises in the cortex and whether it depend on stimulus parameters. Theory: We used the Wilson-Cowan circuit that consists of excitatory and inhibitory cells connected reciprocally and recurrently. We generalized the circuit to form a distributed system: a neural chain. Using mathematical analysis and numerical simulations, we found that the generalized system is intrinsically tuned to certain spatial and temporal frequencies (SF, TF) of stimulation. The intrinsic frequency of the circuit is determined by the synaptic weights that provide network stabilization. Notably, we found that the intrinsic frequency depends on stimulus luminance contrast. Increasing the contrast shifts the intrinsic SF up or down depending on whether the circuit is dominated by excitation or inhibition. Physiology: We tested model predictions by measuring firing rate responses of neurons in the cortical area MT of two alert macaque monkeys. The stimuli were sinusoidal luminance gratings at multiple contrasts and spatiotemporal frequencies. In both monkeys, peak responses depended on contrast in most cases (90%), similar to model predictions. On average, increasing the contrast caused the preferred SF to increase for an amount that depended on TF. The change of preferred SF was large at low TF and small at high TF. Conclusions: Physiological results confirm the model prediction that cortical selectivity for SF depends on stimulus contrast, supporting the view that selectivity is determined by the balance of excitation and inhibition in local cortical circuitry.

III-66. Pitch-trained neural networks replicate properties of human pitch perception

Ray Gonzalez RAYGON@MIT.EDU
Josh H McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

Ideal observer models play an important role in perception research, revealing optimal performance characteristics on psychophysical tasks. However, ideal observers are difficult to explicitly specify for many real-world perceptual problems because knowledge of how the relevant sensory signals are generated is typically incomplete. Here we propose to use artificial neural networks to empirically derive classifiers optimized for real-world problems, the performance characteristics of which can be compared to those of biological systems. We applied this approach to pitch perception, which has a long experimental tradition but remains poorly understood in normative terms.

We trained a convolutional neural network to estimate the fundamental frequency (F0) of vowel-like sounds in noise. The tones were generated from a source-filter model that replicated basic spectral properties of natural sounds. We then ran the network in simulations of classic psychoacoustic experiments using synthetic tones. The tones varied in harmonic composition and phase, variables that yield well-established effects on performance in humans. The network qualitatively replicated these effects, performing better for tones containing low-numbered harmonics and exhibiting effects of phase only for stimuli containing only high-numbered harmonics. This pattern did not occur for a network optimized for the psychoacoustic stimuli themselves, demonstrating that the observed performance characteristic does not merely reflect the intrinsic difficulty of the underlying conditions. The results are consistent with the notion that human pitch perception is near-optimal for the task of estimating F0 from peripheral auditory representations of natural sounds, in that optimizing for performance of this task on naturalistic stimuli is sufficient to reproduce human perceptual characteristics. The results provide a unifying normative explanation for the key psychophysical features of human pitch perception, and illustrate the use of deep neural networks as optimized observer models for real-world tasks.

III-67. Opposing effects of summary statistics on peripheral discrimination

Corey Ziemba CZIEMBA@NYU.EDU
Eero Simoncelli EERO.SIMONCELLI@NYU.EDU

New York University

A converging view of peripheral vision holds that the brain represents statistical summaries of image content in local regions of the visual field. The resulting loss of information can explain the phenomenon of “crowding”, in which recognition of peripheral objects is impaired by surrounding distractors (Balas et.al., 2009; Freeman & Simoncelli, 2011). Here, we show that such statistical representation can either help or hinder visual discrimination, depending on the observer’s task. We created samples of synthetic texture constrained to match a set of higher-order statistics obtained from natural photographs, and measured the ability of observers to discriminate these stimuli within apertures of different sizes. Observers performed two tasks: category discrimination between images with different statistics, and sample discrimination between different images with matching statistics. For both tasks, performance of an ideal observer should improve as stimulus size (and thus stimulus information) increases. In contrast, humans became better at category discrimination but worse at sample discrimination. This occurred regardless of whether stimuli were presented simultaneously at different spatial locations, or in separate temporal intervals at the same location. We found that these opposing effects are predicted by a decision model operating on noisy higher-order statistics that are computed within localized regions that grow with eccentricity at roughly the scale of V2 receptive fields. This model predicts discriminability should scale in a lawful way with eccentricity, similar to the scaling of perceptual crowding, even though the contents of our stimuli are entirely task-relevant, and do not include distractors or clutter. We experimentally verified this prediction, further solidifying the relationship to crowding. These results are consistent with analogous effects in discrimination of auditory textures as a function of temporal window duration (McDermott et.al. 2013), suggesting they may be a general consequence of cascaded sensory transformations.

III-68. Hierarchy as a principle for microcircuit specialization and large-scale dynamics of human cortex

John D. Murray[1] JOHN.MURRAY@YALE.EDU
Murat Demirtas[1] MURAT.DEMIRTAS@YALE.EDU
Joshua B. Burt[1] JOSHUA.BURT@YALE.EDU
Markus Helmer[1] MARKUS.HELMER@YALE.EDU
Jie Lisa Ji[1] JIELISA.JI@YALE.EDU
William Eckner[1] WILLIAM.ECKNER@YALE.EDU
Natasha Navejar[2] NNAVEJAR@TULANE.EDU
William J. Martin[3] BILL@BLACKTHORNRX.COM
Alberto Bernacchia[4] A.BERNACCHIA@GMAIL.COM
Alan Anticevic[1] ALAN.ANTICEVIC@YALE.EDU

[1]Yale University
[2]Tulane University
[3]BlackThorn Therapeutics
[4]University of Cambridge

The large-scale organization of dynamical neural activity across cortex emerges through the interplay between structured long-range interactions and local circuit processes. However, the role of heterogeneity of intrinsic dynamical properties across cortical areas is not well understood. Hierarchy provides a unifying principle for the macroscale organization of many anatomical and functional properties across primate cortex. Functional specialization across the cortical hierarchy may involve interareal heterogeneity of local microcircuitry. Cortical hierarchy is conventionally informed by invasive measurements of long-range projections, creating the need for a principled proxy measure of hierarchy in humans. We found that a noninvasive neuroimaging measure, the MRI-derived myelin map, reliably indexes hierarchy in primate cortex. We hypothesized that functional specialization of human cortical microcircuitry involves hierarchical gradients of gene expression, and analyzed the topography of gene expression using the Allen Human Brain Atlas. We found that the myelin map closely resembles the dominant areal pattern of transcriptomic variation across human cortex. Human cortex exhibits strong hierarchical gradients in expression profiles of genes related to microcircuit function, including synaptic receptors, inhibitory interneuron subtypes, and laminar differentiation, which we validated with anatomical measurements in monkeys. Next, we developed a large-scale dynamical circuit model of human cortex that incorporates heterogeneity of local recurrent synaptic strengths, following the hierarchical myelin map topography, and fit the model using multimodal neuroimaging data from the Human Connectome Project. We found that hierarchical heterogeneity substantially improves the fit to empirical resting-state functional connectivity and captures the hierarchical organization of multiple functional connectivity measures. The heterogeneous model predicts hierarchical topography of higher-frequency spectral power, which we found to be consistent with resting-state magnetoencephalography recordings. Our findings suggest that hierarchy defines an axis shared by transcriptomic, anatomical, and dynamical architectures of human cortex, and that hierarchical gradients of microscale properties contribute to macroscale specialization of cortical function.

III-69. Sensorimotor gain control as a novel mechanism to prevent preparatory activity from causing movement

Timothy Darlington TRD12@DUKE.EDU
Stephen Lisberger LISBERGER@NEURO.DUKE.EDU

Duke University

In the smooth pursuit eye movement region of the frontal eye fields (FEFsem), preparatory activity evolves during fixation and signals expectations about an impending target motion (Figure 3A). In the context of recent analysis of arm-movement preparatory activity in terms of “output-potent” versus “output-null” dimensions, we find that FEFSEM preparatory activity is not confined to output-null dimensions; it has large projections into output-potent dimensions (Figures 1C & D). We understand the presence of output-potent activity in the absence of movement in terms of the previous finding that the output of FEFsem regulates the strength, or “gain”, of visual motion access to the motor system without driving movement itself. To test this theory, we presented brief pulses of visual motion at various times during fixation while monkeys were performing a pursuit task (Figure 2). Behavioral responses to identical pulses of motion were larger when the pulses were given later in a fixation epoch, at times when preparatory activity within FEFsem has a larger projection into output-potent dimensions. Further, behavioral responses are larger and modulated more strongly by preparation when they provide motion in the direction of the upcoming pursuit direction. We conclude that FEFsem preparatory activity increases visual motor gain in a directionally specific manner in anticipation of behaviorally relevant visual motion. Our combined neural and behavioral findings propose a novel mechanism that allows preparatory activity in sensorimotor cortex to evolve without producing inappropriate movement: preparatory signals may act downstream as a gain controller in anticipation of an imminent signal. Even when the gain signal is high, there can be not movement without the sensory signal. Preparatory modulation of gain may be a general mechanism that allows the brain to use expectation to deal with sensory uncertainty and sensory processing delays.

III-70. The virtual rat: from whisker mechanics to the representation of natural tactile scenes

Nadina Zweifel[1] NADINAZWEIFEL2022@U.NORTHWESTERN.EDU
Chengxu Zhuang[2] CHENGXUZ@STANFORD.EDU
Ian Abraham[1] IANABRAHAM2015@U.NORTHWESTERN.EDU
Todd Murphey[1] T-MURPHEY@NORTHWESTERN.EDU
Daniel Yamins[2] YAMINS@STANFORD.EDU
Mitra Hartmann[1] HARTMANN@NORTHWESTERN.EDU

[1]Northwestern University
[2]Stanford University

The nervous system of an animal co-evolves with the morphology of its body, and both brain and body evolve within a particular ethological niche. Neural computations are thus constrained by the material properties and mechanics of an animal's sensorimotor system; understanding these computations requires examining neural circuits in the context of naturally-occurring environmental and organism complexity. Our laboratory uses the rat vibrissal (whisker) system as model to study the neural computations that underlie active tactile perception. Specifically, we aim to characterize the complete mechanical input to the whisker array during exploratory behaviors in naturalistic environments, and then to simulate the neural circuits involved in these behaviors. In conjunction with the companion poster (Zhuang et al., 2018), the present work takes four steps towards this goal. First, we developed and validated a model to accurately simulate the three-dimensional (3D) dynamics of individual whiskers. Second, we constructed two versions of a model of the complete vibrissal array; the versions have increasing biological realism. Third, we collected 3D scans of natural rat habitats and simulated the tactile signals for all possible head positions and orientations within that environment. We quantify environmental statistics (the tactile prior) both in terms of mathematical variables (distance, slope, curvature) as well as in terms of variables accessible to the rat (the mechanical signals at the base of each whisker). Finally, we used the mechanical model in an object classification deep neural network (DNN) to model the computations involved in tactile object classification. Together, these results offer one of the first models to capture the complete input to an active sensing system during natural tactile exploration and to generate some preliminary predictions for the associated neural computations.

III-71. Recording neural activity in unrestrained animals with 3D tracking two-photon microscopy

Mirna Mihovilovic Skanata MIRNA@NYU.EDU
Doycho Karagyozov DPK257@NYU.EDU
Amanda Lesar AAL380@NYU.EDU
Marc Gershow MHG4@NYU.EDU

New York University

Optical recordings of neural activity in behaving animals can reveal the neural correlates of decision making, but such recordings are compromised by brain motion that often accompanies behavior. Two-photon point scanning microscopy is especially sensitive to motion artifacts, and to date, two-photon recording of activity has required rigid mechanical coupling between the brain and microscope. To overcome these difficulties, we developed a two-photon tracking microscope with extremely low latency (360 microsecond) feedback implemented in hardware. Using conventional galvonometric mirrors and a resonant ultrasonic lens, we move the microscope's focal spot rapidly in a cylinder about the center of a single targeted neuron. We count the photons emitted from each point on the scan, use these to form an updated estimate of the neuron's location, and execute the next scan around the updated center. We maintained continuous focus on neurons moving with velocities of 3 mm/s and accelerations of 1 m/s^2 both in-plane and axially, allowing high-bandwidth measurements with modest excitation power. We recorded from motor- and inter- neurons in unrestrained freely behaving fruit fly larvae, including visually responding interneurons and pairs of neighboring neurons, and correlated neural activity with stimulus presentation and

behavioral outputs. This work presents the first measurements of neural activity in behaving larvae and the first two photon measurements of activity in any animal not rigidly coupled to a microscope objective. Our technique can be extended to stabilize recordings in a variety of moving substrates.

III-72. The generation of multiple output channels from the olfactory bulb.

Shelly Jones SHELLY.JONES@UCDENVER.EDU
Nathan Schoppa NATHAN.SCHOPPA@UCDENVER.EDU
Joel Zylberberg JOEL.ZYLBERBERG@GMAIL.COM

University of Colorado

Odor molecules bind to olfactory sensory neurons (OSNs) in the nasal epithelium, which project to two classes of excitatory cells in the olfactory bulb, mitral cells (MCs) and tufted cells (TCs), which project to olfactory cortex. MCs and TCs respond differently to odorant stimulation: TCs have broader tuning profiles, faster response times, and are more sensitive to low odor concentrations than MCs. This effectively creates two output channels from the olfactory bulb, that carry different information to the cortex. Herein, we study the circuit mechanism that generates these two distinct channels for olfactory information. Previous work hypothesized that the difference in response properties of MCs and TCs comes from different numbers of OSN synapses onto those cells, however, recently published electron microscopy studies showed that MCs and TCs have similar densities of OSN inputs, indicating that another mechanism underlies this difference in response properties. Based on known differences in connexin-36 (Cx36) mediated gap junction (GJ) coupling between MCs and TCs, as well as electrophysiology recordings in Cx36 knock out (KO) animals, we hypothesized that selective shunting of synaptic inputs through gap junctions is responsible for the difference in response properties between MCs and TCs. We investigated this GJ hypothesis with a combination of slice electrophysiology in mouse olfactory bulb, and computational modeling in NEURON. We found that GJ coupling can lead to this type of difference in response, and is likely a key factor in generating the difference between MC and TC responses, and thus in shaping the communication of olfactory information from the bulb to the cortex. Given the ubiquity of gap junctions in the brain, this mechanism may also apply to other areas and/or sensory systems.

III-73. A pulvino-cortical circuit model for attention, memory, and decision-making

Jorge Jaramillo JJ99@NYU.EDU
Jorge F. Mejias JORGE.F.MEJIAS@GMAIL.COM
Xiao-Jing Wang XJWANG@NYU.EDU

New York University

The thalamus serves as a bridge between the environment and the cortex, but a comprehensive account of its contribution to cognition remains elusive. Once thought to be a passive relay of sensory information, different thalamic nuclei are now known to play an active role in many of the cognitive functions typically attributed to the cortex. To relate thalamic to cortical computation, we propose a framework that distinguishes two anatomical pathways that reciprocally connect the thalamus to the cortex: (1) a “feed-forward* pathway” that emanates from one cortical area, targets the thalamus, and is relayed to a second cortical area and (2) a “feed-back* pathway” that emanates from one cortical area, targets the reticular-thalamic nucleus and thalamus, and is reciprocated to the same cortical area. What is the link between these anatomical pathways and the neural computations underlying higher-order cognitive functions? In this computational study, we focus on the pulvinar, a prominent subdivision of the visual thalamus that is involved in visual attention and confidence during decision-making. We constructed a thalamocortical circuit model composed of two interconnected cortical areas and the pulvinar. Our model suggests that the pulvinar is a key node that, subject to top-down modulation, gates the effective connectivity between the

two cortical modules via the feed-forward pathway. We show how this pulvinar-mediated gain-control can be used resolve conflicting responses during a combined attention and memory task. Next, we propose that the feed-back pathway, endowed with plastic cortico-thalamic synapses, accounts for the transformation of an implicit representation of decision-making confidence in the parietal cortex to an explicit one in the dorsal pulvinar. Finally, we simulate lesions to the pulvinar and show how they impact behavior and cortical physiology as observed in recent experiments in primates. Overall, our results suggest that the pulvinar, through the feed-forward and feed-back cortico-thalamic pathways, provides crucial contextual modulation to cortical computations associated with attention, working memory, and decision-making.

III-74. Rats optimize reward rate and learning speed in a 2-AFC task

Javier Masis JMASIS@FAS.HARVARD.EDU
Andrew Saxe ASAXE@FAS.HARVARD.EDU
David Cox DAVIDCOX@FAS.HARVARD.EDU

Harvard University

At the heart of perceptual decision-making research is the question of the speed-accuracy trade-off: the more time an actor spends accumulating evidence, the more likely they are to be correct, but the slower they make a choice. When speed and accuracy are balanced to optimize reward rate in free response two-alternative forced choice (2-AFC) tasks, behavior converges to an optimal performance curve (OPC). We examined whether rats approach optimal behavior after learning a visual object recognition task. We found that, like humans, a subset of animals approach the OPC while many respond too slowly and remain above the OPC. Like humans, rats are also more likely to approach optimal behavior for lower error rates. We tracked the rats' development throughout learning and found that, like humans, rats tended to respond too slowly early in learning, and improved error rate before lowering reaction times (RTs). One potential explanation for this behavior, suggested in prior work but not previously formalized, is that rats optimize learning speed in addition to reward rate, in order to maximize future rather than present rewards. We develop a tractable theory of learning in free response 2-AFC tasks based on error corrective learning in deep linear neural networks. Our theory predicts the entire learning trajectory in speed-accuracy space, and shows a decisive advantage to slowing early trials in order to learn faster. To our knowledge, our study is the first analysis of optimal behavior in this context in rats, and our theory is the first to directly incorporate the learning process into free response 2-AFC models. More generally, learning is a critical capacity in higher animals, and the objective to learn rapidly from limited experience may be a fundamental principle which can explain frequently observed suboptimal behavior in humans.

III-75. Learning nonlinearities for identifying regular structure in the brain's inference algorithm

KiJung Yoon[1,2] KIJUNG.YOON@GMAIL.COM
Xaq Pitkow[1,2] XAQ@RICE.EDU

[1]Baylor College of Medicine
[2]Rice University

Neurons in the brain are not simply linear filters followed by a half-wave rectification, and exhibit properties like divisive normalization, coincidence detection, and history dependence. Instead of fixed canonical nonlinear activation functions such as sigmoid, tanh, and relu, other nonlinearities may be both more realistic and more useful. We are particularly interested in multivariate (e.g. two- or three-argument) nonlinearities like f(w1.x, w2.x, ...) which could allow inputs that arise from multiple distinct pathways such as feedforward, lateral, or feedback connections, or different dendritic compartments. Such multi-argument nonlinearities could allow one feature to modulate the processing of others. Many single-argument nonlinearities permit universal computation, but the right nonlinearity

could allow faster generalization during inference, both for the brain and for artificial networks. To address this, we parameterize the nonlinear input-output transformation flexibly by an “inner” neural network, which becomes a ‘subroutine’ called from the conventional “outer” network. These parameters are shared across all layers and all nodes of a given cell type. We evaluate fully-connected feedforward networks on image classification tasks given a diverse set of random initial conditions. We focus especially on the two-argument nonlinearities learned from MNIST and CIFAR-10 datasets. We demonstrate that learned two-argument nonlinearities are reliably shaped roughly like quadratic functions, possibly with a linear transformation on the inputs such as a shift and/or rotation. We therefore separate the training and testing phases by a phase in which we fit an algebraic functional form to the learned inner-network nonlinearities. The algebraic nonlinearity does indeed perform as well as the more richly parameterized nonlinearity in the tasks. In general, these nonlinearities are particularly well-suited for contextual gating of information, and an integral part of the message-passing inference in the brain, because they allow us flexible methods to learn canonical messages as they transform parameters of a population code.

III-76. CELLMax: Maximum likelihood based cell sorting of large-scale neural calcium imaging data

Biafra Ahanonu[1] BAHANONU@STANFORD.EDU
Lacey Kitch[1] LJKITCH@GMAIL.COM
Tony Kim[1] KIMTH@STANFORD.EDU
Margaret Larkin[2] MAGGIE.C.LARKIN@GMAIL.COM
Elizabeth Otto Hamel[1] ELIZABETHJOTTO@GMAIL.COM
Jerome Lecoq[3] JLECOQ@GMAIL.COM
Mark Schnitzer[1] MSCHNITZER@GMAIL.COM

[1]Stanford University
[2]University of California, San Francisco
[3]Allen Institute for Brain Science

Recent advances in large-scale neural calcium imaging allow neuroscientists to visualize the concurrent dynamics of hundreds to thousands of individual neurons in live animals, but analysis of these datasets remains a bottleneck. Existing methods for extracting individual cells and their activity traces from calcium imaging datasets have several limitations, including inherent tradeoffs between signal detection fidelity and crosstalk between nearby cells. To date, no algorithm has demonstrated the requisite speed, scalability, accuracy, and versatility to provide a general solution. Here we present CELLMax (Cell Extraction by Log-Likelihood Maximization), a high-fidelity cell extraction method that makes no assumptions about the temporal structure of cells’ activity traces. The algorithm is highly parallelizable, and under suitable conditions its runtime scales sub-linearly with dataset size. In validation studies with simulated datasets, we found that CELLMax generally yielded superior estimates of cells’ fluorescence activity traces as compared to the most commonly used cell sorting algorithms. In studies with real data acquired in the hippocampus of freely behaving mice, neural activity traces extracted by CELLMax allowed superior positional estimates of the mouse’s running trajectory. Overall, our results show that CELLMax is a versatile, reliable, and scalable approach for extracting cellular signals from a wide variety of calcium imaging datasets. Thus, we expect its usage should help improve the pace and accuracy of experiments relying on large-scale neural calcium imaging.

III-77. Recovering human reward expectation in a bandit setting using Bayesian models

Dalin Guo[1] DAG082@UCSD.EDU
Florent Meyniel[2] FLORENT.MEYNIEL@CEA.FR
Angela Yu[1] AJYU@UCSD.EDU

[1]University of California, San Diego
[2]Commissariat a l'Energie Atomique et aux Energies Alternatives

Making repeated choices among options with uncertainty, such as in the multi-armed bandit task, is suitable for studying how the brain represents and learns about reward statistics based on prior knowledge and sequential observations. Bayesian statistical models have been widely used to explain human behavior as they elegantly quantify inferential uncertainty and identify individual differences. We consider two Bayesian learning models, Dynamic Belief Model (DBM) and Fixed Belief Model (FBM) [1,2], coupled with a decision policy that approximates Thompson Sampling [3], which balances between exploration and exploitation. In stationary bandit games, FBM is the "correct" generative model, assuming the reward rates to be fixed and updating beliefs via Bayes' rule; DBM is an extension that assumes the reward rate undergoes unsignaled, discrete changes. Using FBM, we have shown subjects have substantially different, idiosyncratic prior expectations of mean reward [3]; separately, we have observed DBM to better predict human choice in the bandit task than FBM [2]. Here, we fit DBM and FBM on two datasets: (a) 57 subjects playing two-armed bandit task [3], and (b) 107 subjects playing four-armed bandit task in high abundance (Beta(4, 2)) and low abundance (Beta(2, 4)) environments, where some of the subjects also report expected reward rate of the arms never chosen by the end of each game. We obtained the following key results: DBM predicts trial-by-trial choice better than FBM; individual differences in prior reward expectation are substantial, and consistent whether estimated by FBM or DBM; humans under-estimate reward rates as evidenced in self-reported reward expectations, a finding more accurately recovered by DBM than FBM. Altogether, the results add to prior evidence that humans assume nonstationary reward statistics even in stationary environments, and modeling this component explicitly allows better recovery of both context-dependent and individual-specific parameters related to prior reward expectations.

III-78. Non-Arrhenius dynamics in a balanced cortical network with plastic synapses

Jeff Dunworth JBD20@PITT.EDU
Brent Doiron BDOIRON@PITT.EDU
Bard Ermentrout BARD@PITT.EDU

University of Pittsburgh

Cortical neuron spiking activity is broadly classified as temporally irregular and asynchronous. Model networks with a balance between large recurrent excitation and inhibition capture these two key features, and are a popular framework relating circuit structure and network dynamics. Balanced networks stabilize the asynchronous state through reciprocal tracking by the inhibitory and excitatory population activity, leading to a cancellation of total current correlations driving cells within the network, a phenomenon seen in in vitro preparations (Graupner and Reyes, 2013). While asynchronous network dynamics are often a good approximation of neural activity, in many cortical datasets there are nevertheless brief epochs wherein the network dynamics are transiently synchronized (Buzsaki and Mizuseki, 2014, Tan et al., 2014), or where cortex alternates between intervals of low and high firing rates (Jercog et al., 2017). Traditional balanced networks with linear firing rate dynamics have a single attractor, and fail to exhibit macroscopic synchronous events. Mongillo et al. (2012) showed that balanced networks with short-term synaptic plasticity can depart from strict linear dynamics through the emergence of multiple attractors. We extend this model to incorporate finite network size, introducing quenched noise into the system, and allowing finite size effects to generate transitions between multiple attractors. Principled models of balanced networks require a specific scaling relationship between the system size, N, and the synaptic connection strength, with j

~N^1/2. Recent experimental results provide some evidence for this scaling (Barral and Reyes, 2016). Finite size effects and strong synaptic scaling work in concert to generate an optimal system size for maximizing transitions between two stable states. This is in contrast to the strictly monotonic relationship between transition rates and system size predicted by classical Arrhenius escape. This optimal system size is a novel example of an unexpected result that can come from an intrinsic linking of noise and dynamics.

III-79. A computational model of speech production with internal error detection and correction mechanism

Meropi Topalidou MTOPALID@UCI.EDU
Emre Neftci ENEFTCI@UCI.EDU
Gregory Hickok GREG.HICKOK@UCI.EDU

University of California, Irvine

Sensorimotor interaction is critical for motor control including speech [2, 3, 4]. Recently, it has been argued that speech planning at the phonological level involves an internal feedback control mechanism that can detect and correct phonological selection errors prior to overt speech output [1]. Evidence for this claim includes (i) speech error repair in healthy talkers, which occur more quickly than would be possible using overt feedback alone and (ii) an increase in phonological error rate following brain damage to auditory-motor integration areas. Our present goal was to implement a mechanism that can achieve internal speech error detection and correction. We used the architecture proposed as one level in the Hierarchical State Feedback Control (HSFC) model as described in [1]. The network comprises four structures corresponding to functional-anatomic regions: lexical (pMTG), auditory-phonological (pSTS), motor-phonological (pIFG), and auditory-motor intermediary (Spt) levels. The lexical level is bidirectionally connected to both the auditory and motor levels, which themselves are connected to each other via the Spt auditory-motor interface level. Internal error correction is hypothesized to occur via auditory-motor interaction in cases where the motor plan does not match the lexical and auditory targets [1]. Analysis of network behavior showed that motor errors can be corrected by Spt driving the correction. Activation of just the auditory level also produced activation of the corresponding motor nodes providing a mechanism to support speech repetition. Finally, during correct motor output, activity in Spt suppressed activation of the auditory target, consistent with motor-induced suppression effects, which provides a mechanisms for external error detection [5]. This network behavior may be important for the production of speech sequences where activation of past items must be suppressed to avoid interference in producing the next item [1]. Thus, this simple neuro-psycholinguistically constrained network shows several empirically documented properties that characterize speech production.

III-80. Spectral EEG features track an integrated recognition signal

Christoph Weidemann[1] CTW@COGSCI.INFO
Michael Kahana[2] KAHANA@PSYCH.UPENN.EDU

[1]Swansea University
[2]University of Pennsylvania

A repeated encounter with a person or object frequently elicits feelings of familiarity and recollections of previous interactions. The question of whether these constitute independent signals for recognition has been an issue of contention for decades. Dual-process models of recognition memory typically assume that independent familiarity and recollection signals with distinct temporal profiles can each lead to recognition (enabling two routes to recognition), whereas single-process models posit a unitary "memory strength" signal. In a large-scale study encompassing 132 participants, who each returned for 20 experimental sessions, we used multivariate classifiers trained on spectral EEG features to quantify neural evidence for recognition decisions as a function of time. These classifiers corresponded closely with overt responses with trial-by-trial classifier output correlating strongly with subsequent confidence ratings. Classifiers trained on a small portion of the decision period performed similarly to those also incorporating information from previous time points indicating that neural activity reflects an integrated evidence signal. We propose a single-route account of recognition memory that is compatible with contributions from familiarity and recollection signals, but relies on a unitary evidence signal that integrates all available evidence. Our novel applications of machine learning techniques to answer basic questions about neural evidence during recognition decisions links recognition memory to perceptual and other types of decisions which are commonly thought to rely on a unitary evidence signal. Whereas recordings of brain activity in humans and other animals have supported the view that perceptual decisions are based on a unitary evidence signal, the near consensus among neuroscientists has been that recollective processes add a fundamentally different quality to recognition memory decisions. Instead, our results suggest that any evidence from recollective processes is integrated yielding a unitary memory strength signal that drives recognition decisions.

III-81. Maintained avalanche dynamics during changing pair-wise correlations in nonhuman primates

Stephanie Miller[1,2] STEPHANIE.MILLER@NIH.GOV
Dietmar Plenz[1] DIETMAR.PLENZ@NIH.GOV

[1]National Institute of Mental Health
[2]National Institutes of Health

Neuronal avalanches characterize spontaneous and evoked neuronal activity clusters in cortex across many species and experimental conditions. The hallmark of avalanche dynamics is a power law distribution in size s with exponent close to -3/2. The power law states that the formation of avalanches obeys a fixed ratio in size over all spatiotemporal scales. Such invariance seems contradictive to the flexible and adaptive role of cortex functions, especially given that the brain's functional networks change constantly. So far there is no empirical evidence nor theoretical understanding of how invariable aspects of avalanche dynamics can be achieved by variable functional networks of the cortex. We therefore followed neuronal avalanche dynamics over many days and weeks in premotor and prefrontal cortices of non-human primates (n = 3). Ongoing local field potentials (LFPs) were recorded with chronically implanted multielectrode arrays (10 x 10 electrodes; 400 um interelectrode distance) for 20-60 min each day over successive weeks. We found that the power law in avalanche size was stable in each monkey and cortex area demonstrating the invariance of avalanche dynamics. Similarly, the mean of the pairwise correlation between local cortical sites on the array remained stable. In contrast, individual pairwise correlations changed significantly over time, with a faster change in premotor compared to prefrontal cortex. By using a parametric model, we demonstrated that correlation matrices, i.e. functional connectivities, reconstructed based on the same mean but different individual pairwise entries yield power-law distributions in cluster sizes similar to experimentally observed distributions. Our simulations suggest that a proper and constant level of average pairwise correlation across the cortical network is sufficient to maintain stable avalanche dynamics. We conclude that changes in functional connectivity during ongoing activity do not affect information transmission properties associated with avalanche dynamics.

III-82. Drift correction for electrophysiology and two-photon calcium imaging

Marius Pachitariu[1] MARIUS10P@GMAIL.COM
Carsen Stringer[1] CARSEN.STRINGER@GMAIL.COM
Matteo Carandini[2] MATTEO@CORTEXLAB.NET
Kenneth Harris[2] KENNETH.HARRIS@UCL.AC.UK

[1]HHMI Janelia Research Campus
[2]University College London

Vertical drift (Z-drift) is a major confound for neural recordings during behavior. The brain floats in liquid, and movements of tens of microns can easily occur, even in head-fixed animals. In the mouse, for instance, postural changes such as locomotion can cause vertical brain movements of up to 20 microns. This displacement creates an apparent change in the activity of neurons recorded with either electrode arrays or two-photon calcium imaging. Here, we present methods to estimate and correct the drift in both optical and electrical recordings. We demonstrate three methods to recover Z-drift in 2-photon calcium imaging. (1) Alignment to a densely-scanned reference volume (z-stack). (2) Estimation from a non-functional channel- such as tdTomato expressed in a neuronal subclass. (3) Estimation from changes in the shape of identified cells in functional recordings. We validate methods 2 and 3 by comparing to method 1, which provides ground truth. We then develop correction methods that remove the effects of Z-drift, and show that correlations of neuronal activity with running are significantly decreased. Finally, we develop a convenient online module for drift correction that eliminates Z-drift at sub-micron resolution. Z-drift also affects electrophysiological recordings. The amplitude and shape of extracellular action potentials changes when the electrode moves relative to the brain, and neurons may even disappear altogether from the set of recorded channels. Fortunately, new electrodes such as Neuropixels have dense arrays of channels, with inter-site spacings as low as 20um. We found that we could estimate the drift in extracellular recordings with linear electrodes by tracking neuronal waveform shifts, and corrected for it by spatially interpolating the raw data prior to spike sorting. In summary, the algorithms presented here provide effective methods to remove Z-drift, a major confound for neural recordings during behavioral experiments. We provide the code as part of the Suite2p and Kilosort pipelines.

III-83. Convolutional recurrent neural network models of dynamics in higher visual cortex

Aran Nayebi[1] ANAYEBI@STANFORD.EDU
Jonas Kubilius[2] QBILIUS@MIT.EDU
Daniel Bear[1] DBEAR@STANFORD.EDU
Surya Ganguli[1] SGANGULI@STANFORD.EDU
James DiCarlo[2] DICARLO@MIT.EDU
Daniel Yamins[1] YAMINS@STANFORD.EDU

[1]Stanford University
[2]Massachusetts Institute of Technology

Neurons in the ventral visual pathway exhibit behaviorally relevant temporal dynamics during image viewing. However, the most accurate existing computational models of this system are feedforward hierarchical convolutional neural networks (HCNNs), which capture neurons' time-averaged responses, but do not account well for their complex temporal trajectories. Here we show that HCNNs augmented with both local and global recurrent connections are quantitatively accurate models of dynamics in higher visual cortex.

We began with a five-layer HCNN that achieved state-of-the-art predictions of temporally-averaged visual responses in macaque V4 and IT neurons. To model within-area dynamics, we replaced units in each layer with one of several local recurrent circuit motifs, including simple Recurrent Neural Networks (RNNs), Gated Recurrent Units (GRUs), and Long Short-Term Memory (LSTM) units. We also included combinations of global feedback connections, in which outputs of later convolutional layers were added to inputs of earlier layers. Using back-propagation through time, these new parameters were optimized to predict V4 and IT neural response patterns. Finally, we tested these networks' ability to predict responses on held-out images and neurons not used for model optimization.

We found that the best network structure led to substantial improvements over the feedforward baseline, explaining close to 100% of the explainable variance in V4 neurons and above 75% in IT neurons on average across time points. This network made use of gated local recurrence, with LSTMs and GRUs proving superior to simple RNNs. Furthermore, the presence of specific global feedback connections in this network was critical for best predicting V4 neuron dynamics. In summary, we have developed a deep recurrent neural network architecture that accurately captures temporal dynamics in several ventral cortical areas, opening the door to more detailed computational study of the circuit structures underlying complex visual behaviors.

III-84. Identifying circuit parameters using datasets of diverse neural tuning curves

Gregory Barello GBARELLO@UOREGON.EDU
Takafumi Arakaki TARAKAKI@UOREGON.EDU
Yashar Ahmadian YASHAR@UOREGON.EDU
University of Oregon

Tuning curves characterizing the response selectivities of biological neurons often exhibit large degrees of irregular diversity across cells. Theoretical network models featuring heterogeneous connectivity or single-cell properties can also generate diverse tuning curves. This correspondence can potentially be used to infer or constrain circuit parameters (e.g., statistics of connectivity between different cell-types) solely based on the observed statistics of recorded tuning curves. To this end, we propose to view mechanistic network models as generative models whose parameters can be optimized so the distribution of their generated tuning curves match that of experimentally measured tuning curves. Traditional likelihood-based fitting procedures fail at this task, however, as the likelihood function for most theoretically-motivated models is unavailable or intractable. Recently developed frameworks in machine learning, such as Generative Adversarial Networks (GANs), provide tools for systematically fitting such generative models with implicitly defined likelihoods. Here, we extend and apply GANs to models of cortical network previously developed in theoretical neuroscience, and show that this framework can be used effectively to infer network parameters from datasets of recorded tuning curves. This method avoids the computationally expensive step of inferring latent variables, such as the biophysical parameters of individual recorded cells or the particular realization of the connectivity matrix between them, and directly learns the parameters characterizing the statistics of connectivity or single-cell properties. Moreover, this approach fits the entire joint distribution of tuning curves, instead of matching a few summary statistics picked a priori by the user, as done, e.g., in moment matching. More generally, this framework opens the door to fitting theoretically-grounded dynamical network models directly to simultaneously or non-simultaneously recorded neural data.

III-85. A high resolution data-driven model of the mouse connectome

Joseph Knox[1] JOSEPHK@ALLENINSTITUTE.ORG
Hongkui Zeng[1] HONGKUIZ@ALLENINSTITUTE.ORG
Eric Shea-Brown[2] ETSB@UW.EDU
Stefan Mihalas[1] STEFANM@ALLENINSTITUTE.ORG
Kameron Decker Harris[2] KAMDH@UW.EDU
Nile Graddis[1] NILEG@ALLENINSTITUTE.ORG
Jennifer Whitesell[1] JENNIFERWH@ALLENINSTITUTE.ORG
Julie Harris[1] JULIEHA@ALLENINSTITUTE.ORG
[1]Allen Institute for Brain Science
[2]University of Washington

Knowledge of mesoscopic brain connectivity is important in understanding inter-region communication and information processing. Models of structural connectivity have been used to investigate the relationship with functional connectivity, to compare brain structures across species, and more.

Previous models were constructed with the assumption that regions are homogeneous. While these have proven useful, they depend on predefined regional parcellations and describe connectivity at this same region-limited level of resolution. Here, we go beyond the regional approach and construct a model of the whole brain connectivity at the scale of 100 micron voxels. We use the Allen Mouse Brain Connectivity Atlas, a large scale dataset

measuring with two photon tomography the projections of sets of neurons infected with a viral tracer to approximately 5x10^5 target voxels. While this dataset is enormous, the 429 sources we have in wild type C57Bl/6 mice pale in comparison to the 2.5x10^5 source voxels (as all experiments were conducted in one hemisphere).

To meet this challenge we propose a new model, which relaxes the assumption of homogeneity of connections within a region, and instead uses an assumption of smoothness within a region, with sharp boundaries between regions. We model the connectivity at each source voxel as the kernel weighted average of the projection patterns of nearby injections. We fit the meta-parameters of the model using complete nested cross-validation. The voxel-scale model strongly outperforms the previous regional modeling both in relative mean squared error in cross-validation, and when compared to a human-curated dataset.

This new voxel-scale model been used in several applications including analysis of the modularity of the mouse cortical network. We believe that this is the tip of the iceberg, and that this new voxel-scale model of the mouse connectome will permit researchers to extend their previous analyses of structural connectivity with unprecedented levels of resolution.

III-86. Inferring what you believe from what you do

Xaq Pitkow[1,2] XAQ@RICE.EDU
Paul Schrater[3] SCHRATER@UMN.EDU
Zhengwei Wu[1] ZHENGWEI.WU@BCM.EDU

[1]Baylor College of Medicine
[2]Rice University
[3]University of Minnesota Twin Cities

An animal's behavior is driven by its internal model of the world, which allows the animal to integrate memories, motivation, beliefs, and sensory information. Learning and using this internal model are core functions of the brain, yet these processes are not directly observable in experiments. Here we provide a statistical method for inferring an agent's internal model parameters and the dynamic latent beliefs that the model determines. These properties can provide specific predictors for latent dynamics that can be used to explain measurable nonlinear neural dynamics. We develop this method within a normative model for foraging behavior. To capture the action strategies in an uncertain sensory environment, we use a Partially Observed Markov Decision Process (POMDP) to model the behavior and beliefs of an agent. We assume the animal knows the task structure and dynamics but may significantly misestimate its parameters, such as overweighting sensory evidence. We assume the policy of an agent is optimal under these possibly mistaken parameters, and measure which parameters best account for its observable actions. This can be framed as a maximum likelihood estimation of parameters under a latent variable model for the animal's unobservable, uncertain beliefs about the world. Since the POMDP is Markovian, this estimation problem can be solved using a hidden Markov model (HMM) where the belief is the latent variable. Given the animals' observable behaviors and sensory information, the transition probabilities between states and the policy of animals can be iteratively estimated by a constrained Expectation-Maximization algorithm. We validated this method on a simplified model for a foraging task, and analyzed which parameter combinations are well-constrained by behavioral data. Furthermore, we show how to generalize the method to encompass all the complex task properties in our ongoing foraging experiments.

III-87. Saliency computation in spiking circuit simulations of superior colliculus

Richard Veale[1] RICHARD.E.VEALE@GMAIL.COM
Tadashi Isa[1] ISA.TADASHI.7U@KYOTO-U.AC.JP
Masatoshi Yoshida[2] MYOSHI@NIPS.AC.JP

[1]Kyoto University
[2]National Institute for Physiological Sciences

The superior colliculus (SC) is a layered midbrain structure that receives retinotopically organized visual input from the retina and cortex to its superficial layers. The deeper layers are in spatial register to the superficial layers, and send output to brainstem motor regions which control orienting movements. Despite extensive anatomy and physiology studies examining individual cells in superior colliculus, we still do not understand its intrinsic circuit dynamics, nor how these dynamics contribute to behavior. Towards this goal, we constructed full-scale spiking neural circuit models of horizontal slices of the superficial and deeper layers of superior colliculus, and estimated anatomical and synaptic parameters of these simulations based on physiological data collected from acute slices of the superior colliculus of mouse. To estimate parameters in a statistically rigorous manner, we novelly applied Differential Evolution Markov Chain Monte Carlo (DE-MCMC), and applied approximate Bayesian computation (ABC) to explicitly define the class of acceptable models. After parameter estimation, performant models captured the temporal and spatial dynamics of the physiological data under all experimental conditions. Furthermore, marginal likelihood for anatomical parameters closely match the results from independent anatomical studies, despite no such constraints imposed on the model. Finally, we tested the response of parameterized circuits to realistic visual input derived from a model of mouse retina. The superficial layer circuits had responses to visual input that highly correlated (0.62) with the luminance channel output of a computational saliency map model, suggesting that the retino-tectal circuit implements a computation similar to part of the saliency map model.

III-88. Spike adaptation as the optimal neural code

Hui-An Shen[1,2] HUI-AN.SHEN@UZH.CH
Simone Carlo Surace[1,2] SURACE@INI.UZH.CH
Jean-Pascal Pfister[1,2] JPFISTER@INI.UZH.CH

[1]University of Zurich
[2]ETH Zurich

Spike-based communication is crucial for information transfer between neurons, as digital signals do not deteriorate over long distances. How does a neuron optimally encode its analog somatic membrane potential trajectory into a digital spike train, from an information-theoretic point of view? Here we show how optimal encoding could be implemented in a presynaptic neuron with adaptation, when the synapse is assumed as a near-optimal decoder (filter) for the incoming spikes.

III-89. A probabilistic population code based on neural sampling

Ankani Chattoraj ACHATTOR@UR.ROCHESTER.EDU
Richard Lange RLANGE@UR.ROCHESTER.EDU
Ralf Haefner RHAEFNE2@UR.ROCHESTER.EDU
Shu Chen Wu SWU32@U.ROCHESTER.EDU

University of Rochester

Visual processing is often characterized as implementing probabilistic inference: networks of neurons compute posterior beliefs over unobserved causes given the sensory inputs. How these beliefs are computed and represented by neural responses is much-debated (Fiser et al. 2010, Pouget et al. 2013), often focusing on whether the code is “explicit,” with neurons representing concrete features of the outside world (e.g. Gabor-shaped patches for V1 neurons), or whether it is “implicit”, with responses representing beliefs about abstract variables such as orientation (Pitkow & Angelaki 2017). A second debate concerns whether neural responses represent samples of latent variables (Hoyer & Hyvarinnen 2003) or parameters of their distributions (Ma et al. 2006). We present a model that exhibits key characteristics of both sampling and parametric codes, and that agrees with classic empirical data about neural responses in V1. We propose that V1 spikes represent binary samples from a linear model of the image. The spike rate in such a code is proportional to the marginal posterior probability over the variable represented by the neuron, making this at once a sampling-based as well as a parametric code. Previous work has shown that learning natural images in such a model yields localized, oriented, bandpass, receptive fields in agreement with empirical data (Bornschein et al. 2013). Correspondingly, neural responses in this model show tuning to external stimulus parameters like orientation and spatial frequency. Surprisingly, those tuning curves are approximately contrast-invariant with spike counts that are approximately Poisson-distributed – compatible with a linear probabilistic population code (PPC) for orientation. Finally, extending work by Buesing et al. (2011), we translate the sampling equations into a network of integrate-and-fire neurons that compute conditional probabilities in dendritic nonlinearities and generate spikes according to the correct posterior distribution given an input image.

III-90. EXTRACT: automated cell sorting for large-scale neural calcium imaging based on a framework of robust statistics

Hakan Inan[1] INANH@STANFORD.EDU
Murat Erdogdu[2] ERDOGDU@CS.TORONTO.EDU
Biafra Ahanonu[1] BAHANONU@STANFORD.EDU
Mark Schnitzer[1] MSCHNITZ@STANFORD.EDU

[1]Stanford University
[2]Microsoft Research

Fluorescence calcium imaging is a widely used methodology for tracking the simultaneous activity patterns of large numbers of neurons in awake behaving animals. High-fidelity computational extraction of individual neurons and their activity time traces from the raw calcium video datasets is important for attaining reproducible, high-quality biological results. However, most calcium imaging datasets acquired in behaving animals contain substantial sources of non-stationary noise and signal contamination, such as from brain motion, neuropil activation, or clustered patterns of neural activation. Previous research has led to improved computational cell sorting algorithms, but problems associated with non-stationary noise and contaminants persist, perhaps due to the basic limitations of the statistical frameworks that have been applied. Here we propose a new cell sorting approach based on the statistical framework of robust estimation. Using the theory of M-estimation, we derive a minimax optimal robust loss, and find a simple and practical optimization routine for this loss with provably fast convergence. Using simulated datasets, we showcase the advantages of our robust estimation approach over existing methods commonly used in the field. Finally, we report a fast and efficient computational implementation of our robust approach that has allowed us to extract hundreds of neurons and their dynamics from calcium imaging datasets acquired in freely behaving animals.

III-91. Online convolutional compressed sensing for sparse signal recovery from neuronal spiking activity

Sebastian Weingartner[1] WEINGARTNER@STANFORD.EDU
Mehmet Akcakaya[2] AKCAKAYA@UMN.EDU
Tirin Moore[1] TIRIN@STANFORD.EDU

[1]Stanford University
[2]Department of Electrical and Computer Engineering

The recent introduction of high-density microelectrode arrays with ultra-high channel counts enables the investigation of networks of thousands of neurons. As the recording technology improves, time-efficient and automated processing of neurophysiological data becomes imperative. In particular, online processing has received increasing interest, since operating in real-time on incoming data is a prerequisite for closed-loop experiments or brain-machine-interface applications. Neuronal spiking appears as sparse events on the pre-processed signal, with a single characteristic waveform per neuron. Compressed sensing (CS) is a theoretical framework that has been used for sparse signal recovery from noisy data in a wide variety of applications. Recently convolutional CS with translationally invariant basis functions has also been introduced. In this work, we aim to develop an online convolutional CS reconstruction, extending on the literature of sparse neuronal signal recovery. To this end, we exploit the circular-shifted structure of the encoding matrix, representing the temporally shifted waveforms. This allows decompoisition of the global compressed sensing problem to a number of weakly coupled local regularization problems. We introduce a novel algorithm based on sparse reconstruction by separable approximation (SpaRSA) to perform signal recovery by alternating local thresholding and global aggregation steps. This is further augmented by a localized iterative regularization relaxation scheme for accelerated convergence, which allows online integration of new data in the joint optimization. A localized criterion was used for convergence of early incoming data, which decouples this from recent gradient steps. Finally, due to to the interleaved structure of the convolutional encoding matrix the Newton-step length was adaptively approximated on a per-waveform basis. Numerical simulations were performed with simulated recordings using temporally shifted neuronal waveform templates from actual recordings and additive noise. The proposed algorithm for online computation was compared to batch processing, yielding comparable performance with negligible error for SNRs bigger than 10.

III-92. Empirical vine copula modeling to study multivariate neural representations during complex behaviors

Houman Safaai[1,2] HOUMAN_SAFAAI@HMS.HARVARD.EDU
Selmaan N. Chettih[1] CHETTIH@FAS.HARVARD.EDU
Arno Onken[3] AONKEN@INF.ED.AC.UK
Stefano Panzeri[2] STEFANO.PANZERI@IIT.IT
Christopher Harvey[1] HARVEY@HMS.HARVARD.EDU

[1]Harvard University
[2]Istituto Italiano di Tecnologia
[3]University of Edinburgh

An important goal is to understand how neural activity relates to specific parameters in the external world and behavioral outputs. However, in many cases, it is difficult to isolate the contributions of single behavioral, external, or internal variates to neural activity, especially as the behaviors studied increase in complexity and focus on internal, cognitive variables. Behavior and experiment variables often have complex statistical dependencies, and generating a correct model of the relationships between neural activity and these variables requires understanding the dependency structure between all these variables. State-of-the-art approaches to model neural activity in terms of individual task or stimulus components, such as GLMs or deep neural networks, either make strong assumptions about the dependency structure or do not expose the dependencies explicitly, which makes the analysis and interpretation of how neural activity relates to behavior and experiment variables challenging. We propose a

new approach for the modeling of neural activity that accounts for any general statistical dependency structure, regardless of its complexity and without making any parametric assumptions. We developed an analytical kernel vine copula to estimate the joint density function between neural activity patterns and all the dimensions of interest for the behavioral task or sensory stimulus. The copula structure captures the dependency structure of the multivariate density function and gives an empirical representation of the full statistical dependencies between all the variables, including neural activity. Using the vine decomposition allows us to estimate the density function regardless of the dimensionality. The density function can be used generatively to produce realistic samples, with similar dependency structure to the data, and can utilize mutual information tools to investigate neural coding. We used vine copula modeling to study the mouse posterior parietal cortex during a navigation-based choice task and to distinguish representations of movement-related actions and task-related choice signals.

III-93. Cortical mechanisms for robust sensory coding in the olfactory system

Kevin Bolding BOLDING@NEURO.DUKE.EDU
Kevin Franks FRANKS@NEURO.DUKE.EDU

Duke University

Early sensory responses reflect varying stimulus conditions and behavioral state, and can be noisy and unreliable. Cortical representations of these stimuli must be stabilized across such conditions to extract meaningful information and support adaptive behaviors. How this stabilization is implemented in cortical circuits remains unclear. Here we identify a temporal filtering mechanism, implemented by the interplay of recurrent excitation and feedback inhibition, that decorrelates and stabilizes odor representations in primary olfactory cortex (piriform, PCx) using noisy, correlated and state-dependent input from olfactory bulb (OB). We obtained simultaneous recordings from OB-PCx populations in head-fixed mice before and during anesthesia. Both odor-evoked cell-odor-pair responses and population responses changed substantially in OB but much less so in PCx, suggesting cortical odor representations are recovered from partially degraded OB input. Robust (i.e. cross-state) and state-specific OB responses had similar amplitudes, indicating robust responses are not simply stronger. However, robust OB responses had markedly shorter latencies, indicating that early OB output conveys reliable odor information. In PCx, robust responses had shorter latencies and larger amplitudes than state-specific responses, suggesting a temporal filter that amplifies early inputs. Robust cells were preferentially in deep layer II, where excitatory cells are densely interconnected through recurrent collaterals that also drive strong feedback inhibition. Does this recurrent plexus implement the temporal filter? We used a viral tetanus toxin strategy to selectively eliminate recurrent excitation and recruitment of feedback inhibition. Normally, OB responses occur throughout the sniff while, in PCx, a sparser subset of cells respond shortly after inhalation. After eliminating recurrent collaterals, PCx responses were much larger and sustained, indicating that recurrent circuits abruptly truncate cortical odor responses. Moreover, odor identity coding was impaired, especially at longer latencies. Thus, recurrent cortical circuits implement a temporal filter that selectively transmits robust stimulus information to cortex and truncates later, less specific information.

III-94. Efficient coding in V1: Oriented filters vs. orientation selectivity

Sophia Sanborn SANBORN@BERKELEY.EDU
Dylan Paiton DPAITON@BERKELEY.EDU
Bruno Olshausen BAOLSHAUSEN@BERKELEY.EDU

University of California, Berkeley

A longstanding hypothesis in visual neuroscience is that sensory neurons are adapted to natural image statistics to form an efficient code [Barlow 1961]. Independent Component Analysis (ICA) has been proposed as a normative model for simple cells in V1 based on its ability to reduce higher-order redundancy in natural images. When

applied to natural images, the filters that emerge resemble the localized, oriented, and band-pass receptive fields of V1 neurons. However, Eichhorn, et al. showed that the neural code produced by ICA fails to provide any appreciable gain in redundancy reduction beyond second order methods such as PCA, thus challenging the higher-order redundancy reduction account of V1 function [Eichhorn 2009]. Here, we show that this failure is not due to the learned filters, but rather to ICA's linear encoding scheme. We show that sparse coding [Olshausen 1996] - a related model with a nonlinear encoding scheme - is able to achieve a more efficient code than both ICA and PCA when evaluated in the rate-distortion framework. Our findings suggest that sparse coding's ability to encode orientation, a higher-order statistical property of natural images, enables it to provide a more efficient representation of images, and a better model of V1 response properties, than either ICA or PCA.

III-95. A 3rd factor w/o a 2nd: Dopamine and pre alone rule Drosophila's kenyon cell synapses

Bayar Menzat BAYAR.MENZAT@GMAIL.COM
Scott Waddell SCOTT.WADDELL@CNCB.OX.AC.UK
Tim Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Associative learning depends on the relative timing of stimulus and reinforcement. In Drosophila melanogaster, previous experiences of different odours, and their valences — i.e. whether the smell was associated with reward or punishment — are learned in the mushroom body. During aversive learning, the timing of events will determine the outcome: if an odour immediately precedes a shock, a negative association is formed and flies will try to avoid the odour. Conversely, if the shock is encountered prior to the odour, flies will learn that the odour predicts relief and will then be attracted to that odour. Interestingly, fruit flies are able to form parallel aversive and appetite memories for mixed experiences, because positive and negative valences are learned in synapses onto different mushroom body output neurons (MBONs). Here, we introduce a framework for dopamine modulated learning in a spiking model of the mushroom body that is inspired by experimental observations of dopamine modulated depression (Owald et al. 2014) and potentiation (Cohn et al. 2015). Unlike traditional Hebbian and three factor rules (which fail) we implemented a rule that depends on dopamine spikes and presynaptic spiking alone, independent of postsynaptic activity. We show that our model can reproduce a novel extinction learning experiment in which an odour is first paired with reward but in a later trial coincides with punishment. Furthermore, we find that when odour and reward are co-delivered instantaneously, the firing rates of MBONs decrease monotonically with increasing dopaminergic activity, and we show three possible relationships that can be experimentally tested. Additionally, by linearly combining our aversive and appetitive learning rules, we arrive at a mechanism for integrating mixed valence memories that is in line with current results and inspires a new perspective on valence learning in flies.

III-96. Neuron dendrograms uncover asymmetrical motifs

Roozbeh Farhoodi[1] ROOZBEHFARHOUDI@GMAIL.COM
David Rolnick[2] DROLNICK@MIT.EDU
Konrad Kording[3] KOERDING@GMAIL.COM

[1]University of Philadelphia
[2]Massachusetts Institute of Technology
[3]University of Pennsylvania

While growing neuron databases are beginning to enable rigorous data-driven analysis, the geometric structure of neurons remains poorly understood. We present statistical models that explain much of the variance in branching morphologies demonstrated by neurons, incorporating distance from the soma and asymmetric structure using

hidden variables. By analyzing the models, we observe that many databases of neuron morphologies follow a similar pattern; symmetric branching close to the soma and asymmetric branching for distant parts. This is witnessed by the fact that the frequency of asymmetrical motifs, as captured by graph theoretical trees, is higher than symmetrical motifs. Statistical models of branching in neuron morphologies promise to make their analysis more meaningful and suggest directions for investigating underlying biological processes.

III-97. A biologically inspired neural network model of integration and arbitration of decision making

Kai Krueger[1] KAI.KRUEGER@COLORADO.EDU
Ananta Nair[1] ANANTA.NAIR@COLORADO.EDU
Jessica Mollick[2] JMOLLICK@GMAIL.COM
Seth Herd[1] SETH.HERD@GMAIL.COM
Randy O'Reilly[2] RANDY.OREILLY@COLORADO.EDU

[1]eCortex Inc
[2]University of Colorado, Boulder

Many different and often contradictory constraints go into optimal decision making. Particularly, the range and complexity of problems that need decisions, e.g. from deciding the direction of random dot motion to whether to go to grad school, vary greatly. Model-free (MF) and model-based (MB) decision making have been pervasive in the literature, as a solution to dealing with these differing constraints (multiple controller theory). This raises the question of how these multiple systems interact to make the final decision and initiate action. On the normative, computational side, several hypothesis have been proposed as candidates of arbitration. The more biological, mechanistic side however is still less well understood. We propose a decision making model that combines a model-free and model-based component into a single biologically inspired neural network in order to study its ability to appropriately switch between MF and MB decisions. The model incorporates cortical model-based processing, the subcortical dopamine system and the basal ganglia acting as a model-free controller. The model thereby postulates that arbitration happens through the model-free controller of the basal ganglia. Specifically, the model proposes that cortical layers re-encode the complexities of model-based processing into simplified representations feeding into the basal ganglia that are then similarly available as direct stimuli for reinforcement based decision making. We show that the basal ganglia can learn through the dopamine signal to act (through the cortico-thalamic gating network) based either on the cortical model based output signal, or in a model-free fashion directly from sensory input depending on which signal predicts reward outcome more reliably. Thus we show that the basal ganglia through its gating mechanism provides a natural locus of arbitration and explore the properties of such interaction.

III-98. A Bayesian psychophysics model of sense of agency

Roberto Legaspi ROBERTO.LEGASPI@RIKEN.JP
Taro Toyoizumi TARO.TOYOIZUMI@BRAIN.RIKEN.JP

RIKEN Brain Science Institute

Despite the increasing significance of sense of agency (SoA) research, the literature lacks its formal model: What computational principles underlie SoA—the registration that oneself initiated an action that caused something to happen? We theorize SoA as optimal Bayesian cue integration with mutually involved principles: correlation (consistency), causality (unity) and optimal estimation of agentic attributes. Given that human perception is inherently noisy, the brain resolves ambiguity by drawing on prior expectation of action-outcome spatiotemporal consistency. This correlation can be cognitively modulated by prior experience of cause-effect pairing. Lastly, these joint priors inform SoA only when precise estimates of action-outcome attributes are shifted towards each other. To formal-

ize our theoretical model upon these principles, we drew parallels from the ventriloquism effect—an excellent demonstration of multisensory cue integration.

We used our model to explain what could underlie intentional binding, a robust measure that reveals SoA as temporal binding between volitional action and outcome. Our results show that with our model, albeit simple, we qualitatively reproduced the outcomes of two well-known experiments: the seminal experiment on action awareness (Haggard et al., 2002) where binding occurred only with volitional actions, and one that showed action-outcome binding even in increased uncertainty of outcome (Wolpe et al., 2013). More importantly, however, our results explain how these observed phenomena occurred. We also show how SoA is correlated with volitional actions, but decreased with increased outcome uncertainty. Furthermore, we show that temporal binding is perceived in both self-intended and observed (e.g., TMS-induced) actions suggesting that intention is not strictly necessary as suggested previously as long as subjects have prior experience of the causal relation and its temporal behavior. Temporal binding is therefore a matter of causal perception, i.e., it is causal rather than intentional binding. Hence, SoA is causality that binds action and outcome.

III-99. Using multiple optimization tasks to improve deep neural network models of higher ventral cortex

Chengxu Zhuang CHENGXUZ@STANFORD.EDU
Daniel Yamins YAMINS@STANFORD.EDU

Stanford University

Recent goal-driven deep neural network (DNN) models of higher ventral visual cortex have leveraged the rich behavioral task of object recognition to impose powerful top-down constraints on network parameters. DNNs optimized to solve the multi-way object categorization in challenging real-world images have been shown to provide state-of-the-art predictions of neural responses in visual areas throughout the primate ventral pathway. Here, we show that such models can be improved by using a combination of multiple behaviorally realistic tasks as network optimization targets. Specifically, we optimized a DNN to simultaneously solve high level tasks including object categorization and scene classification, as well as intermediate visual tasks including depth estimation, normal map estimation and semantic segmentation. Task optimization was synergistic, in that performance levels for each task in the combined training were higher at a given number of training examples than for models trained on each task separately. Moreover, the model trained on the combined tasks provided improved ability to fit response patterns in neurons from both cortical areas V4 and IT. These results suggest that identifying a richer and more ecologically relevant variety of visual behaviors as network “goals” may lead to substantially improved understanding of the neural computations in the visual system.

III-100. In the footsteps of learning: changes in network dynamics and dimensionality with task acquisition

Merav Stern[1] MS4325@UW.EDU
Shawn Olsen[2] SHAWNO@ALLENINSTITUTE.ORG
Eric Shea-Brown[1] ETSB@UW.EDU
Yulia Oganian[3] YULIA.OGANIAN@UCSF.EDU
Sahar Manavi[2] SAHARM@ALLENINSTITUTE.ORG

[1]University of Washington
[2]Allen Institute for Brain Science
[3]University of California, San Francisco

When we learn a new task, changes in our neural activity take place in order to accumulate and act upon relevant information. These changes can appear with different magnitudes in multiple brain areas. To understand the dynamics and ultimately the mechanisms of these changes, we follow mice as they learn to perform a visual change detection task and use wide-field GCaMP signaling to record their neural activity across the dorsal surface of the cortex. We also study random neural network models with cortical-resembling high-level area structures; by iteratively training these networks to perform the task we assess the similarities and differences in the mouse cortex and artificial recurrent networks. We find that initially, during the naive behavioral stage, the visual cortex alone responds to the changing stimuli. As the learning progresses, frontal areas respond as well, and eventually, at the expert level, the whole mouse cortex responds to task-relevant stimuli. Cortical activity becomes correlated across all areas, and responses in general become more stereotyped with precise temporal dynamics. Moreover, the dimension of this activity decreases as training progresses. Our artificial neural networks show similar learning-related phenomena. Moreover, the dimension of both cortical and artificial neural networks decreases to roughly 4; we note that this, suggestively, is what would be expected from an independent representation of each task-relevant stimulus sequence (AA, AB, BA, BB). All together, we identify three cortex-wide phenomena that emerge during learning of a basic sequential task: task-specific engagement of surprisingly widespread areas across cortex, an increase in the temporal precision and stereotypy of cortical activity, and a reduction of its dimensionality. These phenomena occur in our neural network models as well, suggesting that they may recur across many learning systems and posing intriguing questions for further theoretical work.

III-101. A recurrent neural network model for factoring distributed representations

E. Paxon Frady PAXON.FRADY@GMAIL.COM
Spencer Kent SPENCER.KENT@BERKELEY.EDU
Bruno Olshausen BAOLSHAUSEN@BERKELEY.EDU

University of California, Berkeley

Neural computation in many areas of the brain must have the capability of binding together different concepts. In the case of high-level visual processing, these bindings might capture the classic problem of associating what an object is with where it is in a larger scene, or it might capture the association of an object's shape with properties of its surface such as color or texture. The question of how a high-dimensional neural population in the brain could encode and compute with such compound concepts has been partially addressed by models of Vector Symbolic Algebra. One major shortcoming of these models is that they cannot uncover the components of a compound representation without either comparing against all possible combinations of underlying components or by storing these combinations in an associative memory. In this work we detail an algorithm that makes computation with compound distributed representations significantly more practical. Our algorithm can be cast as a set of simple recurrent neural circuits we call resonators. We take inspiration from Arathorn's Map-seeking Circuit and the insight that it is useful in solving this particular combinatorial optimization problem to maintain a superposition of hypotheses that are iteratively refined over time through competition. Our approach moves beyond this prior work by using high dimensional distributed representations within the larger algebraic system of a vector symbolic algebra. We also use a neural network to infer these representations for simple visual scenes which is an effective method for dealing with correlations and ambiguity in the input space. This work takes an important step towards realizing the power of symbolic computing within the paradigm of neural networks.

III-102. A theory of memory replay and generalization performance in neural networks

Andrew Saxe ASAXE@FAS.HARVARD.EDU
Madhu Advani MADVANI@FAS.HARVARD.EDU

Harvard University

Why might memories be stored in hippocampus before being replayed into neocortex? The Complementary Learning Systems Theory (McClelland, McNaughton, O'Reilly, 1995) holds that this two stage process allows new information to be gradually incorporated without catastrophically interfering with prior knowledge. A variety of experimental evidence supports these distinct roles of the hippocampus and neocortex in memory formation and learning, including observation of replay events from hippocampus to neocortex during quiet rest and sleep (Peyrache et al., 2009). Yet fundamental questions remain: is replay always beneficial? how much replay is optimal? and how much benefit can replay confer? Here we develop a theory of the impact of experience replay on generalization performance based on exact solutions to the average learning dynamics of simple neural networks. We derive exact solutions to the learning dynamics resulting from two learning strategies: online learning, in which each example is used once and discarded; and batch learning, in which all examples are stored (for instance, in hippocampus) and replayed repeatedly (for instance, during sleep). Remarkably, while these two strategies yield similar performance when training experience is abundant, we find that replay can be decisively better when training experience is scarce. Further, there is a potential cost to replay: generalization performance eventually worsens because of overfitting to noise in the specific examples being replayed. There is therefore an optimal amount of replay that depends on the signal-to-noise ratio of the task to be learned. Our theory thus makes predictions about how the amount of replay observed should depend on task parameters if the brain is optimally managing learning. We also explore learning real datasets with non-linear neural networks and verify that our theoretical predictions hold qualitatively. More broadly, our results suggest a normative explanation for a two-stage memory system: replay enables better generalization from limited training experience.

III-103. Understanding camouflage detection

Abhranil Das ABHRANIL@ABHRANIL.NET
Ross Calen Walshe CALEN.WALSHE@GMAIL.COM
Wilson Geisler W.GEISLER@UTEXAS.EDU

The University of Texas at Austin

Occurrences of camouflage in nature evoke fascination and wonder in us. Less appreciated are the forces that shaped their evolution: the visual systems of their predators and prey. Indeed, having been filtered by them, camouflage specimens—no matter how ingenuous—are poised just at the edge of detectability, their inventiveness only testifying to the sophistication of detection machinery that pruned even slightly less crafty variants.

Using theory, computation and experiment, we turn our inquiry to the visual resources and mechanisms that are harnessed for detecting camouflage in nature.

In particular, we consider the scenario where the camouflaging animal has exactly mimicked its background texture. Any visual information usable for detection then lies only at its boundary. We begin therefore by defining the boundary mismatch: a computational measure of visual discontinuity at the boundary that putatively summarizes most of this available information.

We then synthesize artificial stimuli using 1/f noise as the camouflage texture (this shares the same spatial frequency properties as natural images, but lacks further structure), and assess human performance on them with a series of target-detection experiments.

We find regular variation in the detectability of these stimuli as a function of their boundary mismatch, allowing us to measure boundary-mismatch thresholds against variations in task-relevant stimulus dimensions like luminance, contrast and duration. We shall extend this analysis to variations in the size, distance and shape of the target, and with naturalistic texture stimuli (see Portilla and Simoncelli, 2000).

These ideas can also be brought to the question of engineering effective camouflage. The boundary mismatch measure allows us to computationally prescribe the best location on a background to hide against, and compare the effectiveness of different textures towards this goal. These computational results can be connected to actual detectability in such scenarios using the results of our psychophysical experiments.

III-104. Stability of hippocampal spiking sequences during an olfactory working-memory task

Jiannis Taxidis[1] JTAXIDIS@UCLA.EDU
Peyman Golshani[1] PGOLSHANI@MEDNET.UCLA.EDU
Eftychios Pnevmatikakis[2] EPNEVMATIKAKIS@GMAIL.COM
Apoorva Mylavarapu[1] AMYLAVARAPU@UCLA.EDU

[1]University of California, Los Angeles
[2]Simons Foundation

Neuronal spiking sequences are a candidate mechanism for the brain to retain information in memory for short time periods. Studies suggest that the hippocampus is involved in working-memory tasks by encoding time through such temporal sequences. But the circuit mechanisms underlying these population dynamics are not well understood. How do these sequences adjust to increasing memory load? Once formed, are they stable over multiple days? To gain more insight into these questions, we trained Gad2-Cre:Ai9 head-fixed mice to perform an olfactory delayed non-match-to-sample task (DNMS), requiring working-memory activation during the delay period. Through in vivo two-photon calcium imaging, we recorded activity from hundreds of dorsal CA1 neurons over multiple consecutive days, while mice performed the task, using a variety of delay periods. Calcium imaging videos were motion-corrected, segmented and de-convolved by adapting existing analysis software and a new method was developed for the automatic registration of neurons across multiple days. We observed two distinct classes of pyramidal cell activity: (i) cells encoding the presentation of specific odors ('odor-cells'), and (ii) cells encoding time during the delay period after a given stimulus ('delay- cells'). Collectively, these neurons formed spiking sequences through which both time and odor- identity could be decoded. When the delay was extended, these sequences were partly reshaped. Odor cells remained stable whereas delay cells shifted their fields backwards or forwards in time. Similarly, odor cells were more likely to retain their representation for multiple days, whereas delay cells remapped their activity, yielding overall unstable sequences across days. Our work suggests a robust stimulus-representation in CA1, combined with dynamic odor-specific temporal-sequences. Untangling the mechanisms underlying the co-existence of these two neural codes is crucial for understanding temporal population dynamics in the hippocampus.

III-105. Spike inference for genetically encoded calcium indicators with models of multistep binding kinetics

David Greenberg[1] DAVID.GREENBERG@CAESAR.DE
Damian Wallace[1] DAMIAN.WALLACE@CAESAR.DE
Kay-Michael Voit[1] KAY-MICHAEL.VOIT@CAESAR.DE
Silvia Wuertenberger[1] SILVIA.WUERTENBERGER@CAESAR.DE
Uwe Czubayko[1] UWE.CZUBAYKO@CAESAR.DE
Arne Monsees[1] ARNE.MONSEES@CAESAR.DE
Joshua Vogelstein[2] JOVO@JHU.EDU
Reinhard Seifert[1] REINHARD.SEIFERT@CAESAR.DE
Yvonne Groemping[1] YVONNE.GROEMPING@CAESAR.DE
Jason Kerr[1] JASON.KERR@CAESAR.DE

[1]Research Center Caesar
[2]Johns Hopkins University

With genetically encoded calcium indicators (GECIs), fluorescence changes arising from neural activity can be recorded from the same neurons for days to weeks. However, the relationship between action potential (AP) discharge and GECI fluorescence is complex, nonlinear and variable over neurons. To quantitatively characterize this relationship we developed a sequential binding model (SBM) describing the chain of physical effects by which APs cause fluorescence changes. We then used the SBM as a basis for reliable AP detection.

The SBM describes AP-evoked calcium influx, extrusion, endogenous buffering and GECI binding state transitions using classical mass-action kinetics with separate on- and off-rates for each binding reaction. This biophysical framework allows the SBM to capture nonlinearity, describe the rising and falling phases of the AP response and incorporate variability in response amplitude and shape over neurons. In combined optical/electrical recordings of the four-binding-site indicator GCaMP6s in mouse visual cortex, the SBM reliably and quantitatively predicted fluorescence signals from AP sequences. We also tested the SBM with in vitro binding assays, demonstrating that the same parameters can fit in vitro and in vivo data.

We inferred AP times and neuron-specific SBM parameters using sequential Monte Carlo (SMC), an approach based on data-driven simulation. Our method incorporated novel sampling and resampling strategies specifically designed to improve SMC performance on AP inference, and was implemented for efficient GPU computation. The algorithm outperformed previous AP-detection methods on GCaMP6s, with higher detection rates, fewer false positives, increased timing precision and reduction of systematic errors across neurons. Together, these results demonstrate the potential for model-based, biophysically grounded approaches in the analysis of complex biological data.

III-106. Deep neuronal networks with recursive learning

Jonathan Kadmon[1] KADMONJ@STANFORD.EDU
Haim Sompolinsky[2,3] HAIM@FIZ.HUJI.AC.IL

[1]Stanford University
[2]Harvard University
[3]Hebrew University of Jerusalem

Many sensory pathways are organized in an hierarchical fashion (e.g., vision, audition, olfaction and somatosensation), yet the benefits of information processing by multilayer architectures are still largely unknown. State-of-the-art artificial deep networks—inspired by the natural ones—perform very well when trained on real-world data, however they provide little insights towards understanding real neuronal networks. Indeed, in typical deep learning settings, weight matrices are generated through back-propagation of downstream error signals—a procedure with dubious biological plausibility. Furthermore, we do not have theoretical understanding of why and when these

networks perform well. To study the advantage of multilayer processing in biological settings, we introduce a new class of deep networks, inspired by cortical circuits. Here, each neuron is tuned to fire for a unique random subset of stimuli. Synaptic weights are learned locally and independently from the rest of the network, effectively reducing the training problem to multiple perceptrons. This architecture-independent procedure allows rigorous study of the depth-width tradeoff. Importantly, the framework is learning-rule agnostic and allows different local training strategies to be implemented. Using a simple toy problem of classifying clustered signals, we derive a mean-field description for the neuronal activations as the inputs transverse the network. The layer-by-layer dynamics is described by a nonlinear recursive equation of a scalar order parameter. We find that there is a universal power-law scaling relation between the total size of the network and the optimal depth that minimizes the variance within each cluster at the efferent end. Furthermore, the size and optimal depth both scale with the typical size of the input clusters.

III-107. A visual projection neuron class stops forward walking when detecting regressive translational motion

Matthew Isaacson ISAACSONM@JANELIA.HHMI.ORG
Michael Reiser REISERM@JANELIA.HHMI.ORG
Jessica Eliason ELIASONJ@JANELIA.HHMI.ORG
Aljoscha Nern NERNA@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

The perception of visual motion is useful for animal navigation and moving object detection, and flies are a prominent model system for elucidating the neural mechanisms behind this computation. In flies, directionally selective neurons in the optic lobe compute local motion and project retinotopically into motion-direction-specific layers in the Lobula Plate. Several large visual projection neurons have been extensively studied that integrate Lobula Plate local motion information across a large field-of-view of the fly eye and project this information to the central brain. However, the Drosophila visual system also features small-field visual-projection neurons which have received little attention. Here we present the results of studying one such cell type, which we call Lobula Plate Columnar Type 1 (LPC1), in tethered flies positioned in the center of an LED visual arena capable of delivering all types of rotational and translational motion stimuli. We found that regressive translational motion, which specifically caused flies to cease forward walking, no longer had this effect upon silencing LPC1, yet LPC1 silencing did not affect optomotor reactions to rotational visual motion. Optogenetic activation of LPC1 caused the fly to cease forward walking while also not affecting turning behaviors produced by rotational motion stimuli, demonstrating that the control of forward locomotion and turning are largely decoupled in the fly nervous system. We then used 2-Photon imaging to measure visual stimulus evoked calcium responses in the axon terminals of LPC1 cells and found that this cell type responds strongly to regressive motion presented to its corresponding eye, but this response is abolished when the opposite eye simultaneously received progressive motion—a computation that can differentiate between the visual consequence of translational motion and rotational self-motion using motion information from both eyes.

III-108. A spatiotemporally-resolved view of cellular contributions to network chaos

Rainer Engelken[1] RE2365@COLUMBIA.EDU
Fred Wolf[2] FRED@NLD.DS.MPG.DE

[1]Columbia University
[2]Max Planck Institute for Dynamics and Self-organization

Biophysical properties of single cells can have a drastic effect on the dynamic stability of a network. Previous work

the stability in neural circuit dynamics mostly focused on time-averaged quantities, such as Lyapunov exponents. It is, however, crucial to dissect the stability and instability of network activity as it unfolds in time.

To enable this we here present a spatiotemporally-resolved analysis of dynamical instability in neural circuits. Our method enabled us to measure the contribution of each neuron to the network chaos in a time-resolved manner and the number of neurons contributing to unstable behavior. This is achieved by a novel efficient method for obtaining covariant Lyapunov vectors in numerically exact event-based calculations based on analytical expressions for the Jacobians of the flow.

Using the approach, we measure the contribution of each neuron to the network chaos and the number of neurons contributing to it. We uncover a direct link between single neuron dynamics and network chaos: whenever neurons receive a spike while they are in a susceptible regime of their internal dynamics, they contribute strongly to the collective network chaos. As a consequence of such instability events, the covariant Lyapunov vectors localize, meaning that the unstable degrees of freedom are confined to a small but varying group of neurons. At the transition from chaos to stability, the frequency and duration of unstable episodes decrease, but even in the stable regime, we uncover short unstable periods whose effect is on average dominated by stable periods.

In summary, using a novel tool to efficiently analyze the dynamic instability of large cortical circuits in time and space, we find that individual neurons contribute to chaos by transient instability events. This has applications e.g. for understanding and measuring the stabilization of initially unstable trajectories during learning.

III-109. Representation of choice bias in the activity of prearcuate gyrus during perceptual decision making.

Gabriela Mochol[1] GABRIELA.MOCHOL@UPF.EDU
Roozbeh Kiani[2] ROOZBEH@NYU.EDU
Ruben Moreno Bote[1] RUBEN.MORENO@UPF.EDU

[1]Universitat Pompeu Fabra
[2]New York University

Models of learning postulate that perceptual decisions may be made not only based on sensory stimuli but also based on prior history of choices, rewards and stimuli. Such biases may be beneficial when sensory information is weak or ambiguous, especially if subjects are not certain about task structure or when prior history carries relevant information about upcoming stimuli or rewarding actions. Here, we report the existence of history-dependent biases in the behavior of highly-trained monkeys performing a motion direction discrimination task and demonstrate a neuronal representation of the bias in the activity of prearcuate gyrus (PAG) neurons. In our task, stimulus direction and strength varied randomly across trials, making previous history non-informative for future choices. Nonetheless, monkeys showed small but significant biases that fluctuated at two different time scales: slow (tens to hundreds of trials) and fast (previous trial). Fast bias on each trial reflected previous choice and feedback, while slow bias reflected the neighboring choices on trials with erroneous responses or ambiguous stimuli. Knowing these biases significantly improved our ability to predict monkeys' upcoming choice on individual trials (improved accuracy >3 % for more difficult motion coherence with p(correct)$<$0.75, compared to a model that predicted choices based only on the stimulus strength). PAG neural population responses represented both the fast and slow biases prior to the motion onset, indicating a correlate for both types of bias in the prefrontal cortex. Critically, trial-to-trial variability of these neural representations of bias trended toward improving our ability to predict the upcoming choice, suggesting a functional role for these representations. Because the same PAG neurons also represented past choices and feedback, they could offer a compact circuit for computation of prior history signals and leveraging those signals to guide behavior.

III-110. Manifold inference from neural dynamics

Ryan Low RLOW@PRINCETON.EDU
Sam Lewallen LEWALLEN@PRINCETON.EDU
Dmitriy Aronov DA2006@COLUMBIA.EDU
Rhino Nevers RNEVERS@PRINCETON.EDU
David Tank DWTANK@PRINCETON.EDU

Princeton University

Population recordings make it possible to study information processing in neuronal networks, but raise the challenge of uncovering structure and meaning in the nonlinear interactions between neurons over time. We developed a novel manifold learning algorithm incorporating temporal dynamics, in order to characterize population activity as a trajectory on a nonlinear space of possible network states. The manifold's structure captures correlations between neurons and temporal relationships between states—constraints that arise from the network's underlying architecture and inputs. Using measurements of activity over time but no information about animal behavior or other external variables, we obtain low dimensional trajectories that reflect both behavior and structured dynamics that do not arise directly from measured external variables. We validated algorithms using a simulated place cell network, and recovered a low dimensional manifold isomorphic to the physical environment represented by the network. We then applied our algorithms to rat hippocampal recordings. During random foraging, the network's trajectory on the manifold was low dimensional and largely similar to the rat's trajectory in the physical environment. However, the trajectories could diverge during local episodes, possibly reflecting internal processes like prediction, memory, or mental exploration. During a non-spatial task that required adjusting an auditory tone to a goal frequency, manifold coordinates represented tone frequency. In both tasks, many neurons had compact firing fields on the manifold, indicating that they fired preferentially in a local neighborhood of network states. The framework we present here can be used to perform trajectory dependent, nonlinear dimensionality reduction for analyzing and visualizing population activity. It may also be useful as a bridge between data and theoretical models, which could make predictions in terms of manifold properties. Alternatively, phenomenological models could be formulated directly on the manifold, and mapped back to neural activity using our techniques.

III-111. Functional investigation of behavioral circuits using precise photo-stimulation

Joseph Donovan JOE@NEURO.MPG.DE
Marco Dal Maschio DALMASCHIO@NEURO.MPG.DE
Thomas Helmbrecht HELMBRECHT@NEURO.MPG.DE
Herwig Baier HBAIER@NEURO.MPG.DE

Max Planck Institute of Neurobiology

A mechanistic understanding of behavioral circuits requires approaches that can determine the contribution of individual neurons to activity in both local and downstream areas. The combination of optogenetic actuators and precise holographic light shaping enables simultaneous activation of multiple single neurons throughout a 3D volume. Together with multiplane calcium imaging and behavioral tracking, this approach can follow the induced activity across the brain and relate it to the motor output. We applied this approach to investigate a premotor midbrain nucleus in larval zebrafish, which is involved in the control of posture and tail bending. Manipulations at the scale of individual neurons lead to a combinatorial explosion of possible activation patterns and therefore require an efficient strategy to select which neurons to activate. We adopted an iterative approach based on maximizing the behavior induced, which should minimize activity irrelevant to behavior. Starting from larger activations patterns, we iteratively select smaller and smaller subsets, at each step choosing the subset that most effectively drives behavior. This procedure efficiently finds minimal subsets, of only a few neurons, sufficient to drive tail bending. To relate network activity to tail movements, we use regularized regression models, which can predict the motor output for each activity state. To improve and extend this approach, we use simulations to evaluate

the effectiveness different photostimulation selection strategies. In conclusion, we propose a flexible investigation strategy, towards the goal of combining natural behavior with guided photostimulation to gain mechanistic insights of circuit function.

III-112. Systems consolidation without replay? Learning rules and circuit architectures for consolidation in cerebellar learning

Jay Bhasin[1] JBHASIN@STANFORD.EDU
Jennifer Raymond[1] JENR@STANFORD.EDU
Mark Goldman[2] MSGOLDMAN@UCDAVIS.EDU

[1]Stanford University
[2]University of California, Davis

A common feature of learning and memory systems is that expression of long-term memory after consolidation can depend on neural circuits distinct from those involved in acquisition. While this transformation has been shown phenomenologically—for example, from hippocampus-dependent to hippocampus-independent storage of declarative memories—the mechanism by which changes supporting consolidation are orchestrated across synaptic sites is not well understood.

We address this problem by working in a well-characterized system: the gaze-stabilizing cerebellar and brainstem circuitry underlying the vestibulo-ocular reflex (VOR). When an animal rotates its head in space, it reflexively counter-rotates its eyes so that image motion on its retinas is minimized. By placing the animal in a virtual environment where the feedback is manipulated to suggest that head movements are eliciting erroneously large or small eye movements responses, the system adjusts the amplitude of the reflex. Initially, the expression of learning depends on the cerebellar cortex, but over time becomes cerebellum-independent, presumably being transferred downstream.

Motivated by previous work, we constructed a circuit model of VOR learning which included empirically motivated plasticity rules at two synaptic sites: fast, error-driven plasticity in the cerebellum and a slow Hebbian rule with a sliding threshold downstream in the brainstem, the hypothesized target of consolidation. During simulated VOR training with the standard sliding threshold plasticity rule, the model was able to consolidate learned increases in eye movement amplitude, even in the absence of replay. However, in the presence of noise, consolidation became unstable. We show how stability can be restored if the downstream plasticity rule is tuned to the statistics of the natural vestibular input. This work illustrates how simple learning rules can break down in the presence of noise, and suggests how systems consolidation may occur even in the absence of replay events and be tuned to the statistics of natural signals.

III-113. An integrated hierarchical control architecture compositional in dynamics and policies.

Thomas Ringstrom RINGS034@GMAIL.COM
Paul Schrater SCHRATER@GMAIL.COM

University of Minnesota Twin Cities

A consensus exists within computational neuroscience that the structures underpinning motor control in the brain are organized hierarchically. However, the nature of the hierarchy remains under theoretical development, with specialized hierarchical frameworks proposed for low-level motor control and different hierarchical models for high-level planning and decisions. Additionally, there is no consensus on the representation of policies and how they should be organized, dereferenced, and mixed for planning and action. These issues manifest largely because there is not a coherent integrated theory of high-to-low level control that uses a common set of components and operations—we are missing an integrated compositional framework. Compositionality as a principle is defined by the use operators on a set of learned or primitive elements to create new reusable solutions. Inherently, compositional algorithms are generative models which significantly reduce the computational overhead of finding new solutions. In this work, we argue for an architecture based around a hierarchical extension of the Linearly Solvable Markov Decision Process (hLMDP). The hLMDP is useful for breaking down policies into compositional atoms which can be recomposed in situations with different constraints, rewards, and dynamics. We demonstrate that the hLMDP has intriguing qualities which make it well suited to tackling problems in two domains that are typically studied in isolation: controlling across a factored space of dynamics, and controlling efficiently with time-varying reward structures. The hLMDP offers neuroscientists a common theoretical language that connects important computational variables at different time-scales, from planning to execution. Additionally, it has the potential to provide new interpretations of the role of cortical circuit activity during planning and execution. The hLMDP handles real-time model-based control, learning, time-varying goals and constraints, while crucially preserving analytic solvability and optimality, ultimately paving a way to scale up computational theory to handle the systems-level analysis required for complex tasks.

III-114. Nonlinear impact of structural plasticity on cortical information storage capacity

Andreas Knoblauch[1] KNOBLAUCH@HS-ALBSIG.DE
Friedrich Sommer[2] FSOMMER@BERKELEY.EDU

[1]Albstadt-Sigmaringen University
[2]University of California, Berkeley

It is well understood that structural synaptic plasticity increases information storage capacity of sparsely connected neural networks (Chklovskii et al., Nature 431:782, 2004). Theoretical works counting configurations of actual synapses suggest that this increase would be in proportion to the potential connectivity of the network (or the inverse filling fraction). However, a weakness of such approches is that it remains unclear how much of this configurational information could really be read-out by a plausible neural network model of learning and memory. Here we give a detailed account by testing learning and retrieval in different network models for associative memory. Each network has the size and connectivity of a generic cortical macrocolumn (1cmm, n=100000 neurons, P=0.1 connection probability). After learning an increasing number of memory patterns by structural plasticity, we test retrieval of the stored memories. Thereby, we are able to estimate precisely the storage capacity of such networks. Interestingly, we observed a linear increase of storage capacity with potential connectivity (as predicted by the previous theories) only for non-sparse neural activity. By contrast, for sparse activity patterns as observed in experiments (Waydo et al., J.Neurosci 26:10232, 2006), sparsely connected networks can store and retrieve only a few memories over a large range of potential connectivity (e.g., 0.1-0.5), whereas storage capacity increases sharply over several orders of magnitude if potential connectivity exceeds some threshold value (e.g., around 0.5) well consistent with experimental measurements of the filling fraction. This effect is not limited to synaptic pruning in simple Willshaw-type models as reported in previous works, but holds as well in more realistic and/or efficient models of synaptic learning and ongoing structural plasticity. Thus, we argue that previous works have strongly underestimated the impact of structural plasticity on learning and memory. We also outline implications for technical applications and models of cognition.

III-115. A minimal model for coherent chaos in a neural network

Itamar D Landau[1] ITAMAR.LANDAU@MAIL.HUJI.AC.IL
Haim Sompolinsky[2,1] HAIM@FIZ.HUJI.AC.IL

[1]Hebrew University of Jerusalem
[2]Harvard University

Firing-rate fluctuations and irregular spiking are ubiquitous in the neocortex. Theoretical models accounting for these observations have their foundations in the chaotic dynamics of recurrent networks, whether through excitation-inhibition balance in spiking models or the more abstract models of rate chaos. Yet a key emergent feature of these models is the decorrelation of neural activity such that the macroscopic, population activity remains nearly constant in time. Indeed a major challenge to theoreticians has been to produce network models which generate macroscopic, spatially coherent fluctuations which can account for broad correlations observed in cortical activity, at the level of spiking activity and especially at the macroscopic level of LFP or EEG.

Here we present an analytically tractable minimal model for macroscopic chaos generated internally by a recurrent network. In particular we show that a standard chaotic rate network endowed with a rank-one structured connectivity component results in chaotic single neuron activity with spatially coherent correlations at the macroscopic level. We identify two regimes: a moderately coherent regime in which the dynamics are stationary and can be solved analytically via dynamic mean-field theory, and a strongly coherent regime in which stationarity breaks down, globally coherent fluctuations dominate, and various time-scales emerge. With this minimal model as a platform we present a number of models with biological interest, introducing excitatory and inhibitory neurons and hierarchically structured subpopulations. We furthermore reveal an intriguing relationship between statistics of the eigenspectrum of the connectivity matrix and the resulting chaotic dynamics. Our work establishes the theoretical groundwork for recurrent neural networks that internally generate spatially coherent fluctuations.

III-116. Only a subset of asynchronous irregular network states are responsive

Zahara Girones[1] ZAHARAGIRONES@GMAIL.COM
Matteo di Volo[2] M.DIVOLO@GMAIL.COM
Alain Destexhe[2] DESTEXHE@UNIC.CNRS-GIF.FR

[1]University of Oregon
[2]CNRS - UNIC

Desynchronized brain states are known to be associated with arousal and increased awareness, but why such brain activity is so widely seen, particularly in awake and attentive mammals, is not known. Intracellular studies have shown that, during desynchronized states, cortical neurons are depolarized and display a large membrane conductance and intense fluctuations. Balanced networks of excitatory and inhibitory neurons displaying asynchronous irregular (AI) states, are models faithfully reproducing desynchronized states. At the single-cell level, it has previously been shown that neurons subject to balanced and noisy synaptic inputs can display enhanced responsiveness. By scanning a large number of networks in AI states, here we show that such enhanced responsiveness is also present at the network level, but only when the conductance state of single neurons and their membrane potential fluctuations are within the biological range. In such states, the entire population of neurons is globally influenced by the external input. We use a mean-field model to characterized the network responsiveness based on three variables: the mean conductance state, the mean membrane potential, and the amplitude of the membrane potential fluctuations of the network neurons. As seen in the network, optimal responsiveness is obtained only when these three parameters are consistent with experimental measurements. Thus, there exists a continuum of AI states, only some of which are responsive. Given the dependence of the responsiveness on the neuronal conductance state, this work shows that intracellular measurements are needed to distinguish the AI states that are responsive.

III-117. Sensory cortex is optimised for prediction of future input

Yosef Singer YOSSING@GMAIL.COM
Yayoi Teramoto YAYOI.TERAMOTO@BALLIOL.OX.AC.UK
Ben Willmore BENJAMIN.WILLMORE@DPAG.OX.AC.UK
Jan Schnupp WSCHNUPP@CITYU.EDU.HK
Andrew King ANDREW.KING@DPAG.OX.AC.UK
Nicol Harper NICOL.HARPER@DPAG.OX.AC.UK

University of Oxford

Neurons in sensory cortex are tuned to diverse features in natural scenes. But what determines which features neurons become selective to? Here we explore the idea that neuronal selectivity is optimised to represent features in the recent past of sensory input that best predict immediate future inputs. We tested this hypothesis using simple feedforward neural networks, which were trained to predict the next few video or audio frames in clips of natural scenes. The networks developed receptive fields that closely matched those of real cortical neurons, including the oriented spatial tuning of primary visual cortex, the frequency selectivity of primary auditory cortex and, most notably, in their temporal tuning properties. The temporal prediction model provides a principled approach to understanding the temporal aspects of RFs. Previous models, based on sparsity, were successful in accounting for many spatial aspects of V1 RF structure, and had some success in accounting for spectral aspects of A1 RF structure. However, these models do not account well for the temporal structure of V1 or A1 RFs. Notably, for both vision and audition, the envelopes of real neuronal RFs tend to be asymmetric in time, with greater sensitivity to very recent inputs compared to inputs further in the past. Here we show that these shortcomings are largely overcome by the temporal prediction approach. Furthermore, the better a network predicted future inputs the more closely its receptive fields tended to resemble those in the brain. This suggests that sensory processing is optimised to extract those features with the most capacity to predict future input.

III-118. Mixed selectivity and population coding in primary and secondary auditory cortex.

Joshua Downer[1] JOSHUA.DOWNER@UCSF.EDU
Jennifer L. Mohn[2] JLMOHN@UCDAVIS.EDU
Kevin N. O'Connor[2] KNOCONNOR@UCDAVIS.EDU
Mitchell Sutter[2] MLSUTTER@UCDAVIS.EDU

[1]University of California, San Francisco
[2]University of California, Davis

Although sensory cortex is mostly associated with the processing of basic stimulus features, many studies have revealed effects of non-sensory variables. For instance, wakefulness, locomotion, attention and other factors have been shown to modulate sensory representations even at the earliest auditory cortical stage, A1. Generally, these studies have examined the effects of a single non-sensory variable on the responses of neurons to a single sensory variable. However, real-world environments contain sounds that vary along multiple feature dimensions and behavioral demands change constantly. In order to illuminate how auditory cortex supports flexible behavior in rich acoustic environments, we recorded single neurons in both A1 and a secondary auditory cortical field (ML) while rhesus macaques performed a feature-selective attention task. We find that single neurons exhibit significant mixed selectivity for sensory and non-sensory variables, with A1 exhibiting more linear and ML more non-linear interactions between variables. Interestingly, in neither field does attention enhance the selectivity for attended features relative to unattended features, as is commonly found in sensory neurons. In order to obtain a compact description of how the observed mixed selectivity in A1 and ML may contribute to task performance, we performed population state-space analyses using targeted dimensionality reduction techniques. These analyses revealed a dynamic de-mixing of task variables over the course of a trial, leading to enhanced population-level representation of attended features. Although population state-space trajectories qualitatively differ between A1 and ML, both

areas carry equivalent amounts of task-relevant information at the population level across the course of a trial. These results are compellingly similar to findings in pre-frontal cortex, where single neuron mixed selectivity gives rise to flexible population coding of dynamic task variables. Thus, early sensory cortex participates in behavior beyond simply filtering and relaying sensory information.

III-119. Controlling burst activity in cortical microcircuits with a homeostatic inhibitory plasticity rule

Filip Vercruysse[1] FILIPVERCRUYSSE@GMAIL.COM
Richard Naud[2] RNAUD@UOTTAWA.CA
Henning Sprekeler[3,4] H.SPREKELER@TU-BERLIN.DE

[1]Technical University Berlin
[2]University of Ottawa
[3]Berlin Institute of Technology
[4]Bernstein Center for Computational Neuroscience

The existence of specialized mechanisms for burst generation in pyramidal cells (PCs) suggests that bursts are likely to be an important temporal feature of neural signals. Bursts appear to be correlated with sensory processing and perception[1], are involved in reliable transmission of action potentials and long-term potentiation, and have been proposed as a cellular mechanism to combine external and internal information[2]. In layer 5 PCs, bursts occur at a low, but consistent rate, and are thought to arise from active dendritic processes. Recent theoretical work suggests that such a low, but finite average burst rate is required for optimal burst coding[3] and for the biological implementation of error backpropagation in neural networks[4]. Given that burst activity relies on dendritic threshold mechanisms, it appears likely that low burst activity require homeostatic control, but the underlying mechanisms are not resolved. Given that inhibitory Martinotti cells regulate dendritic activity[5], we hypothesized that plasticity of inhibitory connections onto apical dendrites can regulate burst firing. To investigate this hypothesis, we used a two-compartmental model of L5 pyramidal cells that was fitted to in vitro data. Our results show that a simple Hebbian plasticity rule on inhibitory synapses leads to robust and self-organized control of dendritic and burst activity. The dendritic learning rule we propose is inspired by a homeostatic rule we previously proposed to control somatic spiking activity[6] and therefore inherits properties such as a balance of excitation and inhibition. The learning rule also supports realistic firing patterns in a recurrent network with dendritic inhibition, and may hence form an important building block for the self-organized regulation of cortical microcircuits with different inhibitory cell types.

III-120. Understanding functional clusters in the larval zebrafish brain using neural circuit models

Yu Hu[1,2] YU.HU@MAIL.HUJI.AC.IL
Florian Engert[3] FLORIAN@MCB.HARVARD.EDU
Misha Ahrens[4] AHRENSM@JANELIA.HHMI.ORG
Haim Sompolinsky[3,1] HAIM@FIZ.HUJI.AC.IL
Yu Mu[4] MUY@JANELIA.HHMI.ORG
Xiuye Chen[3] XIUYECHEN@GMAIL.COM
Aaron Kuan[3] AARONTKUAN@GMAIL.COM
Maxim Nikitchenko[3] NIKITCHM@GMAIL.COM
Owen Randlett[3] OWEN.RANDLETT@GMAIL.COM
Alexander Schier[3] SCHIER@FAS.HARVARD.EDU

[1]Hebrew University of Jerusalem
[2]ELSC

[3]Harvard University
[4]HHMI Janelia Research Campus

Whole-brain calcium imaging allows for an unprecedented view of the activity of populations of neurons, but brings with it new challenges for extracting biological meaning from such recordings. Unsupervised clustering methods, which attempt to recover groups of neurons with similar activity patterns and thus reduce the complexity of the data, are powerful for analyzing the functional organization of the larval zebrafish brain [Vladimirov et al 2014] at the cellular level [Chen et al 2015]. The clusters show a plethora of activity patterns covering the sensory-tuned and motor-tuned and stages in-between. However, an abstraction of homogeneous clusters will take away the individuality of neurons and may mask important dynamics of small groups of neurons. A better mechanistic understanding of the functional properties of these clusters is critical for understanding the transformation from stimulus to behavior and brain-intrinsic dynamics.

Here we investigate the structure of functional clusters, focusing on two aspects: the relation between anatomical and functional structures of clusters, and the variability of the activity of individual neurons within a cluster. First, we find that a cluster often separates into anatomically peripheral and localized core components, the latter coinciding with high functional coherence. We find that the dynamics of peripheral neurons are consistent with reflecting input from long-range projections from the core. On the other hand, the dynamics within a cluster core is consistent with that arising from locally recurrent circuits. In particular, we find the deviations between individual neuron's activity and the cluster average are low dimensional. We develop a recurrent circuit theory that captures the observed variability structure and provides estimates of the effective strengths of local interactions within a cluster. This provides a new theoretical basis for understanding dynamically generated variability and low dimensional population activity structure.

III-121. Gradient descent for spiking neural networks

Dongsung Huh[1] DONGSUNGHUH@GMAIL.COM
Terrence Sejnowski[2] TERRY@SALK.EDU

[1]University of California, San Diego
[2]Salk Institute for Biological Studies

The brain performs real-time computations using dynamic biophysical components: the neurons are coupled by brief impulses, or spikes, which are delivered through the network by the dynamics of axons, dendrites, and synapses. How the brain coordinates the complex dynamics of spiking neural networks (SNNs) to form the basis for computation is the central problem in neuroscience. Deep learning models simplify the problem by assuming static units that produce analog output, which describes the time-averaged firing-rate response of a neuron. These rate-based artificial neural networks (ANNs) have an advantage that the learning rules for training them are widely available, which SNNs currently lack. The recent success of deep learning demonstrates the computational potential of trainable, hierarchical distributed architectures. I investigate the principle for spike-based computation by merging the top-down approaches of deep learning with the bottom-up biophysical spiking neural network architectures. Towards this goal, I derived a general learning algorithm for SNNs from an optimal control principle, representing the first step in harnessing the computational capacity of SNNs. I will discuss the new findings on how the trained SNNs perform various computational problems.

III-122. Temporally varying neural responses to spatially periodic stimuli

Jason Pina JAY.PINA@PITT.EDU
Bard Ermentrout BARD@PITT.EDU

University of Pittsburgh

Certain images that have spatial components in a narrow band of wave numbers have been shown to induce temporally varying neural responses. In pattern sensitive epilepsy, striped lines can trigger epileptic seizures if they are close to 2–5 cycles deg^{-1} (cpd) (Wilkins et al. 1975). Similarly, images – including abstract artwork – with peaks in power near 3 cpd are known to cause aversion in healthy individuals (Fernandez and Wilkins, 2008). Both of these phenomena have been shown to induce abnormal temporal activity in electroencephalography (EEG) or magnetoencephalograhy (MEG) recordings.

Neural fields have proven useful at modeling the spatiotemporal dynamics of ensembles of neurons and capturing many experimentally observed patterns, such as planar and spiral waves. We are thus motivated to consider a spatially extended neural field model where a static, spatially periodic stimulus is provided as input to the excitatory and inhibitory neural populations. By adjusting system parameters such as the amount of recurrent excitation, we may place the stimulus-free system near a so-called Turing-Hopf bifurcation, where the uniform steady state is spontaneously lost to temporally and spatially periodic patterns with wave number m. Simulations and numerical bifurcation analysis for the 1-D system demonstrate the desired resonance, displaying spatially periodic temporal oscillations with very weak stimuli for some wave numbers while requiring much stronger stimuli for others. We analytically show that a weak stimulus with wave number m destabilizes the steady state, and find the stability boundary as a function of the recurrent excitation and stimulus strength. Finally, we present a more realistic 2-D system simulated on a GPU that exhibits this strong sensitivity to the spatial frequency of the stimulus. These 2-D simulations also allow us to demonstrate resonance to noisy images with dominant wave numbers near m, matching experimental findings in visual discomfort.

Author Index