

10th Annual Meeting

COSYNE

Computational and Systems Neuroscience

2013

www.cosyne.org

MAIN MEETING

Feb 28 - Mar 3, 2013

Salt Lake City, UT



Program Summary

Thursday, 28 February

- 4:00 pm **Registration opens**
- 5:30 pm **Welcome reception**
- 6:20 pm **Opening remarks**
- 6:30 pm **Session 1: Keynote**
Invited speaker: W. Bialek
- 7:30 pm **Poster Session I**

Friday, 1 March

- 7:30 am **Breakfast**
- 8:30 am **Session 2: Circuits**
Invited speaker: E. Marder; 3 accepted talks
- 10:30 am **Session 3: Vision**
6 accepted talks
- 12:00 pm **Lunch break**
- 2:00 pm **Session 4: Audition**
Invited speaker: B. Shinn-Cunningham; 3 accepted talks
- 4:00 pm **Session 5: Sensory Processing**
Invited speaker: Y. Fregnac; 2 accepted talks
- 5:15 pm **Dinner break**
- 7:30 pm **Poster Session II**

Saturday, 2 March

- 7:30 am **Breakfast**
- 8:30 am **Session 6: Coding/Computation**
Invited speaker: I. Fiete; 3 accepted talks
- 10:30 am **Session 7: Coding/Computation**
Invited speaker: T. Sejnowski; 3 accepted talks
- 12:00 pm **Lunch break**
- 2:00 pm **Session 8: Decision making**
Invited speaker: C. Brody; 3 accepted talks
- 4:00 pm **Session 9: Networks**
Invited speaker: K. Boahen; 2 accepted talks
- 5:15 pm **Dinner break**
- 7:30 pm **Poster Session III**

Sunday, 3 March

- 7:30 am **Breakfast**
- 8:30 am **Session 10: Learning/Decision making**
Invited speaker: P. Schrater; 3 accepted talks
- 10:30 am **Session 11: Behavior/Motor**
Invited speaker: D. Gordon; 3 accepted talks
- 12:00 pm **Lunch break**
- 2:00 pm **Session 12: Neurons, Stimuli and Perception**
Invited speaker: A. Movshon; 1 accepted talk

Poster Session Topics

<i>Topic Area</i>	Session I <i>Thursday</i>	Session II <i>Friday</i>	Session III <i>Saturday</i>
Cognition: attention, memory	1–3	1–3	1–2
Bayesian, optimality	4–10	4–8	3–9
Neural correlations, population coding	11–14	9–13	10–13
Cognition: decision making, reward, confidence	15–26	14–24	14–25
Neural encoding, decoding	27–33	25–31	26–32
Circuits	34–45	32–42	33–43
Basal ganglia, bird song, grid cells, hippocampus, navigation	46–50	43–48	44–47
Sensory: hearing	51–53	49–50	48–49
Cognition: objects, categories	54–55	51–53	50–52
Learning / plasticity	56–61	54–59	53–57
Motor systems	62–66	60–65	58–63
Oscillations	67–68	66–68	64–67
Sensory: adaptation, statistics, perception	69–71	69–70	68–70
Sensory: chemical, multisensory, somatosensory	72–76	71–74	71–75
Sensory: vision	77–89	75–88	76–88
Networks theory, modeling, computational	90–100	89–100	89–100

RECORD - PROCESS - STORE THOUSANDS OF CHANNELS WITHOUT COMPROMISE

1000's of channels with:

- BROAD BAND DATA STORAGE
- REAL-TIME PROCESSING
- VERSATILE DATA DISTRIBUTION
- BEHAVIORAL INTEGRATION
- POWERFUL VISUALIZATION
- COMPLEX ELECTRICAL STIMULATION

1000's of channels of

- SINGLE UNITS
- MUA
- EEG/EMG
- ECOG

TDT has augmented the powerful Z-Series RZ/PZ recording systems with new software and hardware to facilitate very high channel count visualization and processing as well as direct to disk data streaming on thousands of channels—of course we can do less too—without compromise.


The logo for Tucker-Davis Technologies (TDT) features the letters 'TDT' in a bold, stylized, italicized font. The 'T' and 'D' are connected, and the 'T' has a unique shape with a curved top.

Tucker-Davis Technologies
www.tdt.com

11930 Research Circle
Atachua, Florida 32615
Tel: 386.462.9622
Fax: 386.462.5365
E-mail info@tdt.com



The MIT Press




The Neural Basis of Free Will

CRITERIAL CAUSATION

Peter Ulric Tse

A neuroscientific perspective on the mind-body problem that focuses on how the brain actually accomplishes mental causation.

384 pp., 28 illus., \$38 cloth




Brain and the Gaze

ON THE ACTIVE BOUNDARIES OF VISION

Jan Lauwereyns

A radically integrative account of visual perception, grounded in neuroscience but drawing on insights from philosophy and psychology.

312 pp., 49 illus., \$40 cloth




Discovering the Human Connectome

Olaf Sporns

A pioneer in the field outlines new empirical and computational approaches to mapping the neural connections of the human brain.

224 pp., 17 color plates, 55 b&w illus., \$35 cloth



Principles of Brain Dynamics

GLOBAL STATE INTERACTIONS

edited by Mikhail I. Rabinovich, Karl J. Friston, and Pablo Varona

Experimental and theoretical approaches to global brain dynamics that draw on the latest research in the field.

Computational Neuroscience series
360 pp., 36 color plates, 94 b&w illus., \$65 cloth

How We Remember

BRAIN MECHANISMS OF EPISODIC MEMORY

Michael E. Hasselmo

A novel perspective on the biological mechanisms of episodic memory, focusing on the encoding and retrieval of spatiotemporal trajectories.

336 pp., 8 color plates, 111 b&w illus., \$40 cloth

Visual Population Codes

TOWARD A COMMON MULTIVARIATE FRAMEWORK FOR CELL RECORDING AND FUNCTIONAL IMAGING

edited by Nikolaus Kriegeskorte and Gabriel Kreiman

How visual content is represented in neuronal population codes and how to analyze such codes with multivariate techniques.

Computational Neuroscience series
632 pp., 14 color plates, 151 b&w illus., \$55 cloth

Biological Learning and Control

HOW THE BRAIN BUILDS REPRESENTATIONS, PREDICTS EVENTS, AND MAKES DECISIONS

Reza Shadmehr and Sandro Mussa-Ivaldi

A novel theoretical framework that describes a possible rationale for the regularity in how we move, how we learn, and how our brain predicts events.

Computational Neuroscience series • 400 pp., 133 illus., \$40 cloth

back in print

Sensory Communication

edited by Walter A. Rosenblith

Available again: a landmark collection on sensation and perception that remains influential.


860 pp., \$50 paper



The MIT Press mitpress.mit.edu



is the new centre for Neuro-Electronics Research Flanders.



NERF is a joint basic research initiative, set up by VIB, imec and KU Leuven to understand the function of brain circuits. NERF labs are located in the heart of Europe in Leuven Belgium on imec's nanotechnology campus. At NERF, world-class neuroscientists perform innovative, collaborative, interdisciplinary research, combining nanoelectronics with neurobiology.

Currently, NERF hosts four research laboratories:

- chemosensory processing in zebrafish (Yaksi Lab)
- visual processing in mice (Bonin Lab)
- memory processing and spatial navigation in rodents (Kloosterman Lab)
- learning and decision-making (Haesler lab)
- neural circuitry of vision (Farrow Lab)

and two visiting groups:


- Optical control of neural activity in behaving animals (Battaglia Group)
- Novel neural probes and large-scale anatomy (McNaughton Group)

NERF is looking for enthusiastic, energetic early career researchers (at the PhD and postdoc level) with strong interests in systems neuroscience to join our laboratories. PhD students will have the opportunity to enroll in KU Leuven international PhD programs. For more information, visit our website (www.nerf.be). For applications, please contact individual PIs directly.

Imec - NERF,
Kapeldreef 75,
3001 Leuven Belgium



Also visit us during
the upcoming NERF 2013
Neurotechnology Symposium
on April 22-23
(www.nerf.be/symposium).



Great companies never stop moving forward

Qualcomm is a world leader in providing mobile communications, computing and network solutions that are transforming the way we live, learn, work, and play.

As a pioneering technology innovator, we are committed to continuous evolution and exploration. Our researchers and computational neuroscientists engage in a wide variety of exciting and technically challenging projects—including the application of biologically inspired learning machines to a new breed of neuromorphic computing devices.

We help you work smarter, move faster and reach further. We see success in your future. In fact, we're passionate about it.

We are futurists.

www.qualcomm.com



QUALCOMM



Subscribe to Active Zone

The Cell Press Neuroscience Newsletter



Featuring:


- Cutting-edge neuroscience from Cell Press and beyond
- Interviews with leading neuroscientists
- Special features: Podcasts, Webinars and Review Issues
- Neural Currents - cultural events, exhibits and new books
- *And much more*

Read now at

bit.ly/activezone



Cell
PRESS



Introducing Cerebus Central Suite
Software v6.04

No limitations – cbSDK Freedom

 **BLACKROCK**[®]
MICROSYSTEMS
sales@blackrockmicro.com



SEEKING COMPUTATIONAL SCIENTIST IN SENSORIMOTOR CONTROL

We are currently hiring to fill multiple full-time positions in engineering, and computational modeling of sensorimotor control at Brain Corporation, San Diego, CA.

We are seeking exceptional candidates with practical experience in sensorimotor processing, reinforcement learning or action selection and who have excellent programming skills in C/C++, Python or MATLAB.

TEN REASONS TO JOIN BRAIN CORPORATION

Work on the most exciting scientific challenge of the century.

Outstanding team of computational scientists, programmers, and engineers.

Industry-level salary, stock option grant, matching 401K, end of year bonus.

Health, dental, and vision insurance.

Free breakfast and lunch at work; lots of healthy choices.

Free iPhone5 for new employees; free gadgets twice per year.

Live 5 minutes from the Pacific Ocean. Surf or dive any day of the year. Bike to work.

Bring your family to sunny San Diego: beaches, zoo, Sea World, Legoland, Wild Animal Park.

Family-friendly events, all-expense-paid holiday retreats, deep-sea fishing.


Attend conferences. We even stay an extra day at the Cliff Lodge just to ski.

Do you have substantial experience in any of the following: constructing internal models and motor primitives; modeling locomotion, reaching and grasping; vision for navigation and control; oculomotor systems, proprioceptive systems, spinal cord, cerebellum, basal ganglia, or motor cortex; or computational or theoretical research in spike coding, spiking network dynamics and plasticity?

Please submit your CV/resume and relevant papers to Dr. Eugene M. Izhikevich, CEO at jobs@braincorporation.com

WE APPLY NEUROSCIENCE TO ROBOTICS. WE ARE INTRIGUED BY THE SECRETS OF THE BRAIN AND WE USE SIMULATIONS TO UNDERSTAND NEURAL COMPUTATIONS. WE ARE PASSIONATE ABOUT CHANGING THE WORLD WITH SMARTER ROBOTS THAT HAVE ARTIFICIAL NERVOUS SYSTEMS.

PLEASE VISIT WWW.BRAINCORPORATION.COM
FOR MORE INFORMATION



Plexon is the pioneer and global leader in custom, high performance and comprehensive data acquisition, behavior and analysis solutions specifically designed for neuroscience research.

Innovation for the next 30 years

Plexon Inc

30th

ANNIVERSARY
1983 to 2013

CinePlex®




Behavioral Research

OmniPlex® D



Data Acquisition

PlexBright™



Optogenetics

SAVE
THE
DATE

4th Annual Plexon Neurophysiology Workshop
April 29 – May 2, 2013

www.plexon.com



grapevine

neural interface system

neural interface system
 neuroscience research
 neuroprosthetic development
 Mac OSX, Windows, Linux
 up to 512 electrodes

connect to



amplifier front ends for electrodes

amplify, filter, and digitize biopotential signals



- surf D** 16-channel differential pairs - *EMG*
- surf S** 32-channel single reference - *EEG, ECog, ENG*
- micro** 32-channel single reference - *spikes and LFP*
- micro +stim** 32-channel single reference, record and stimulate - *spikes and LFP*
- nano** rodent microelectrode applications




scout | low-cost laboratory processor option for up to 128 channels.



Contact us for a demonstration in your lab.

toll free | (866) 324-5197
 local | (801) 413-0139
 fax | (801) 413-2874
 email | sales@rppl.com



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 2013 Leadership

Organizing Committee:

- *General Chairs*: Jonathan Pillow (University of Texas at Austin), Nicole Rust (University of Pennsylvania)
- *Program Chairs*: Marlene Cohen (University of Pittsburgh), Peter Latham (The Gatsby Computational Unit, UCL)
- *Workshop Chairs*: Jessica Cardin (Yale), Tatyana Sharpee (Salk)
- *Communications Chair*: Kanaka Rajan (Princeton)

Executive Committee:

- Anne Churchland (Cold Spring Harbor Laboratory)
- Zachary Mainen (Champalimaud Neuroscience Programme)
- Alexandre Pouget (University of Geneva)
- Anthony Zador (Cold Spring Harbor Laboratory)

Advisory Board:

- Matteo Carandini (University College London)
- Peter Dayan (University College London)
- Steven Lisberger (UC San Francisco and HHMI)
- Bartlett Mel (University of Southern California)
- Maneesh Sahani (University College London)
- Eero Simoncelli (New York University and HHMI)
- Karel Svoboda (HHMI Janelia Farm)

Program Committee:

- Misha Ahrens (Harvard University)
- Bruno Averbeck (National Institute of Health)
- Jeff Beck (University of Rochester)
- Matthias Bethge (Max Planck Institute, Tübingen)
- David Freedman (University of Chicago)
- Surya Ganguli (Stanford University)
- Maria Geffen (University of Pennsylvania)
- Andrea Hasenstaub (Salk Institute)
- Vivek Jayaraman (Janelia Farm)
- Adam Kohn (Albert Einstein College of Medicine)
- Mate Lengyel (University of Cambridge)
- Michael Long (New York University)
- Wei Ji Ma (Baylor College of Medicine)
- Christian Machens (Champalimaud Foundation, Lisbon)
- Liam Paninski (Columbia University)
- Anitha Pasupathy (University of Seattle)
- Nicholas Price (Monash University)
- Emilio Salinas (Wake Forest University)
- Elad Schneidman (Weizmann Institute)
- Garrett Stanley (Georgia Institute of Technology)
- Angela Yu (University of California San Diego)

Reviewers:

Ehsan Arabzadeh, Rava Azeredo da Silveira, Wyeth Bair, Scott Brincat, Ethan Bromberg-Martin, Lars Buesing, Yoram Burak, Timothy Buschman, Daniel Butts, Charles Cadieu, Stijn Cassenaer, Matthew Chafee, Zhe Chen, Paul Cisek, Shaun Cloherty, Ruben Coen-Cagli, Anita Disney, Eizaburo Doi, Shaul Druckmann, Jan Drugowitsch, Yasmine El-Shamayleh, Bernhard Englitz, Ila Fiete, Flavio Frohlich, Makoto Fukushima, Sebastian Gerwinn, Jesse Goldberg, Mark Goldman, Arnulf Graf, Michael Graupner, Sebastian Guderian, Robert Guetig, Ralf Haefner, Markus Hietanen, Mark Histed, Gregory Horwitz, Mark Humphries, Robbie Jacobs, Beata Jarosiewicz, Kresimir Josic, Adam Kampff, Steve Kennerley, Minoru Koyama, Andrei Kozlov, Nikolaus Kriegeskorte, Hakwan Lau, Anthony Leonardo, Alexander Lerchner, Nicholas Lesica, Matthieu Louis, Jorg Lucke, Leo Lui, Jakob Macke, Jason MacLean, Olivier Marre, Jamie Mazer, Ferenc Mechler, Daniel Meliza, Hugo Merchant, Paul Miller, Ruben Moreno Bote, Eran Mukamel, Kathy Nagel, Thomas Naselaris, Kristina Nielsen, Yael Niv, John O'Doherty, Shawn Olsen, Gergo Orban, Srdjan Ostojic, Gonzalo Otazu, Stephanie Palmer, Rony Paz, Bijan Pesaran, Biljana Petreska, Xaq Pitkow, Eftychios Pnevmatikakis, Kanaka Rajan, Alfonso Renart, Constantin Rothkopf, Christopher Rozell, Cristina Savin, Paul Schrater, Odelia Schwartz, John Serences, Peggy Series, Maoz Shamir, Pradeep Shenoy, Fabian Sinz, Matthew Smith, Samuel Sober, Samuel Solomon, Alireza Soltani, Martin Stemmler, Ian Stevenson, Thibaud Taillefumier, Ching-Ling Teng, Taro Toyozumi, Wilson Truccolo, Tomoki Tsuchida, Naotsugu Tsuchiya, Joji Tsunada, Richard Turner, John Tuthill, Tim Vogels, Ed Vul, Ross Williamson, Byron Yu, Kechen Zhang, Shunan Zhang, Davide Zoccolan, Joel Zylberberg

Conference Support

Program and Publicity Support:

- Tomáš Hromádka, Cold Spring Harbor Laboratory

Administrative Support, Registration, Hotels:

- Denise Soudan, Conference and Events Office, University of Rochester

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. Three award granting programs were available in 2013.

The generosity of our sponsors helps make these travel grant programs possible.

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.

This program is supported by the following corporations and foundations:



- The Gatsby Charitable Foundation
- Qualcomm Incorporated
- Brain Corporation
- Cell Press/Neuron
- Evolved Machines
- National Science Foundation (NSF)

The 2013 recipients are:

Gautam Agarwal, Evan Archer, Gautam Awatramani, Arjun Bansal, Yuval Benjamini, Ralph Bourdoukan, Johanni Brea, Charles Cadieu, Simone Carlo Surace, James Cavanagh, Matthew Chalk, Julio Chapeton, Christine Constantinople, Sean Escola, Limor Freifeld, Felipe Gerhard, Arnulf Graf, Matjaz Jogan, Mel Win Khaw, Robert Kerr, Kamesh Krishnamurthy, Guillaume Lajoie, Cheng-Hang Liu, Ashok Litwin-Kumar, Gabriela Mochol, Mayur Mudigonda, Douglas Ollerenshaw, Marius Pachitariu, Marino Pagan, Cengiz Pehlevan, Ari

About Cosyne

Rosenberg, Cristina Savin, Shreejoy Tripathy, Balazs Ujfalussy, KiJung Yoon, Yunguo Yu, Adam Zaidel, Joel Zylberberg.

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.



This program is supported by a grant from the National Science Foundation.

The 2013 recipients are:

Lilach Avitan, Manisha Bhardwaj, Sara Constantino, Maria Dadarlat, Emily Denton, Rong Guo, Ann Hermundstad, Komal Kapoor, Karin Knudson, Radhika Madhavan, Alejandro Ramirez, Friederike Schuur, Neda Shahidi, Adhira Sunkara, Nergis Toemen, Fleur Zeldenrust.

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 (“NSF Cosyne Mentors”).



This program is supported by a grant from the National Science Foundation.

The 2013 NSF Cosyne Mentors are listed below, each followed by their mentee:

John Cunningham and Gamaleldin Elsayed, Andrea Hasenstaub and Elizabeth Phillips, Wei Ji Ma and Shan Shen.

Program

Note: Printed copies of this document do not contain the abstracts; they can be downloaded at:

http://cosyne.org/c/index.php?title=Cosyne2013_Program.

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

Thursday, 28 February

- 4:00 pm **Registration opens**
- 5:30 pm **Welcome reception**
- 6:20 pm **Opening remarks**

Session 1: Keynote

(Chair: Peter Latham)

- 6:30 pm [Are we asking the right questions?](#)
William Bialek, Princeton University (**invited**) 23
- 7:30 pm **Poster Session I**

Friday, 1 March

- 7:30 am **Continental breakfast**

Session 2: Circuits

(Chair: Garrett Stanley)

- 8:30 am [The impact of degeneracy on system robustness](#)
Eve Marder, Brandeis University (**invited**) 25
- 9:15 am [The cortical network can sum inputs linearly to guide behavioral decisions](#)
M. H. Histed, J. Maunsell, Harvard Medical School 37
- 9:30 am [Thalamic drive of deep cortical layers](#)
C. Constantinople, R. Bruno, Columbia University 27
- 9:45 am [Whole-brain neuronal dynamics during virtual navigation and motor learning in zebrafish](#)
M. B. Ahrens, K. Huang, D. Robson, J. Li, M. Orger, A. Schier, R. Portugues, F. Engert,
Janelia Farm Research Campus 44
- 10:00 am **Coffee break**

Program

Session 3: Vision

(Chair: Emilio Salinas)

10:30 am	Lateral interactions tune the early stages of visual processing in <i>Drosophila</i> L. Freifeld, D. Clark, H. Yang, M. Horowitz, T. Clandinin, Stanford University	30
10:45 am	Visual speed information is optimally combined across different spatiotemporal frequency channels M. Jogan, A. Stocker, University of Pennsylvania	35
11:00 am	A neural encoding model of area PL, the earliest face selective region in monkey IT C. Cadieu, E. Issa, J. DiCarlo, Massachusetts Institute of Technology	36
11:15 am	Chromatic detection from cone photoreceptors to individual V1 neurons to behavior in rhesus monkeys C. Hass, J. Angueyra, Z. Lindbloom-Brown, F. Rieke, G. Horwitz, University of Washington	40
11:30 am	The impact on mid-level vision of statistically optimal V1 surround normalization R. Coen-Cagli, O. Schwartz, University of Geneva	39
11:45 am	Learning to infer eye movement plans from populations of intraparietal neurons A. Graf, R. Andersen, California Institute of Technology	33
12:00 pm	Lunch break	

Session 4: Audition

(Chair: Bruno Averbeck)

2:00 pm	Peripheral and central contributions to auditory attention Barbara Shinn-Cunningham, Boston University (invited)	26
2:45 pm	The dynamics of auditory-evoked response variability and co-variability across different brain state G. Mochol, S. Sakata, A. Renart, L. Hollender, K. Harris, J. de la Rocha, IDIBAPS	41
3:00 pm	Stimulus-response associations shape corticostriatal connections in an auditory discrimination task P. Znamenskiy, Q. Xiong, A. Zador, Cold Spring Harbor Laboratory	43
3:15 pm	Self-supervised neuronal processing of continuous sensory streams R. Guetig, Max Planck Institute of Experimental Medicine	41
3:30 pm	Coffee break	

Session 5: Sensory Processing

(Chair: Maria Geffen)

4:00 pm	Hidden complexity of synaptic receptive fields in cat V1 Yves Fregnac, CNRS - UNIC (invited)	27
4:45 pm	A proposed role for non-lemniscal thalamus in cortical beta rhythms: from mechanism to meaning S. Jones, C. Moore, Brown University	38
5:00 pm	Multiple perceptible signals from a single olfactory glomerulus M. Smear, A. Resulaj, J. Zhang, T. Bozza, D. Rinberg, Janelia Farm Research Campus	32
5:15 pm	Dinner break	
7:30 pm	Poster Session II	

Saturday, 2 March

7:30 am **Continental breakfast**

Session 6: Coding/Computation

(Chair: Matthias Bethge)

8:30 am	Evidence of a new (exponentially strong) class of population codes in the brain Ila Fiete, University of Texas, Austin (invited)	24
9:15 am	Feed-forward inhibition in hippocampal microcircuits: adaptation to spike-based computation B. B. Ujfalussy, M. Lengyel, University of Cambridge	30
9:30 am	Structured chaos and spike responses in stimulus-driven networks G. Lajoie, K. K. Lin, E. Shea-Brown, University of Washington	42
9:45 am	Triple-spike-dependent synaptic plasticity in active dendrites implements error-backpropagation M. Schiess, R. Urbanczik, W. Senn, University of Bern	31
10:00 am	Coffee break	

Session 7: Coding/Computation

(Chair: Surya Ganguli)

10:30 am	Suspicious coincidences in the brain: beyond the blue Brain. Terrence Sejnowski, Salk Institute for Biological Studies (invited)	27
11:15 am	Conditioned interval timing in V1 by optogenetically hijacking basal forebrain inputs C. Liu, Johns Hopkins University	31
11:30 am	Rank-penalized nonnegative spatiotemporal deconvolution and demixing of calcium imaging data E. A. Pnevmatikakis, T. Machado, L. Grosenick, B. Poole, J. Vogelstein, L. Paninski, Columbia University	37
11:45 am	Temporal basis for predicting sensory consequences of motor commands in a cerebellum-like structure A. Kennedy, G. Wayne, P. Kaifosh, K. Alvina, L. Abbott, N. B. Sawtell, Columbia University	42
12:00 pm	Lunch break	

Session 8: Decision making

(Chair: Wei Ji Ma)

2:00 pm	Neural substrates of decision-making in the rat Carlos Brody, Princeton University (invited)	24
2:45 pm	Neural variability and normalization drive biphasic context-dependence in decision-making K. Louie, M. W. Khaw, P. Glimcher, New York University	34
3:00 pm	On the role of neural correlations in decision-making tasks N. Parga, F. Carnevale, V. de Lafuente, R. Romo, Universidad Autonoma de Madrid	32
3:15 pm	Saccadic choices without attentional selection during an urgent-decision task E. Salinas, G. Costello, D. Zhu, T. Stanford, Wake Forest University	38
3:30 pm	Coffee break	

Program

Session 9: Networks

(Chair: Jeff Beck)

4:00 pm	Neurogrid: A hybrid analog-digital platform for simulating large-scale neural models Kwabena Boahen, Stanford University (invited)	23
4:45 pm	Spiking networks learning to balance excitation and inhibition develop an optimal representation. R. Bourdoukan, D. G. Barrett, C. Machens, S. Deneve, Ecole Normale Supérieure	33
5:00 pm	Formation and maintenance of multistability in clustered, balanced neuronal networks A. Litwin-Kumar, B. Doiron, Carnegie Mellon University	39
5:15 pm	Dinner break	
7:30 pm	Poster Session III	

Sunday, 3 March

7:30 am **Continental breakfast**

Session 10: Learning/Decision making

(Chair: Mate Lengyel)

8:30 am	Learning worth in an uncertain world: probabilistic models of value Paul Schrater, University of Minnesota (invited)	26
9:15 am	Evidence for decision by sampling in reinforcement learning M. W. Khaw, N. Daw, New York University	28
9:30 am	Inference rather than selection noise explains behavioral variability in perceptual decision-making J. Drugowitsch, V. Wyart, E. Koechlin, INSERM	29
9:45 am	An investigation of how prior beliefs influence decision-making under uncertainty in a 2AFC task D. Acuna, M. Berniker, H. Fernandes, K. Kording, Northwestern University	34
10:00 am	Coffee break	

Session 11: Behavior/Motor

(Chair: Misha Ahrens)

10:30 am	Collective regulation in ant colonies Deborah Gordon, Stanford University (invited)	25
11:15 am	On the duality of motor cortex: movement representation and dynamical machine P. N. Sabes, University of California, San Francisco	45
11:30 am	Evidence for a causal inverse model in an avian song learning circuit R. Hahnloser, N. Giret, J. Kornfeld, S. Ganguli, Institute of Neuroinformatics UZH / ETHZ	28
11:45 am	A network model for learning motor timing in songbirds C. Pehlevan, F. Ali, T. M. Otchy, B. Olveczky, Harvard University	44
12:00 pm	Lunch break	

Session 12: Neurons, Stimuli and Perception

(Chair: Marlene Cohen)

2:00 pm	Neural correlates of visual orientation constancy A. Rosenberg, D. Angelaki, Baylor College of Medicine	36
2:15 pm	Functional models and functional organisms for systems neuroscience Anthony Movshon, New York University (invited)	25
3:00 pm	Closing remarks	

I-1. [Independent pools of visual short-term memory resource for different features](#)
Hongsup Shin, Ronald Van den Berg, Wei Ji Ma, Baylor College of Medicine 45

I-2. [Tactile working memory in rat and human: Prior competes with recent evidence](#)
Athena Akrami, Arash Fassihi, Vahid Esmaeili, Mathew E Diamond, School for Advanced Studies (SISSA) 46

I-3. [Mechanisms and circuitry underlying basal forebrain enhancement of top-down and bottom-up attention](#)
Michael Avery, Nikil Dutt, Jeffrey Krichmar, University of California, Irvine 46

I-4. [Semi-parametric Bayesian entropy estimation for binary spike trains](#)
Evan Archer, Il M Park, Jonathan W Pillow, University of Texas at Austin 47

I-5. [Sparse coding model and population response statistics to natural movies in V1](#)
Mengchen Zhu, Ian Stevenson, Urs Koster, Charles Gray, Bruno Olshausen, Christopher J Rozell, Georgia Institute of Technology 48

I-6. [Efficient hierarchical receptive field estimation in simultaneously-recorded neural populations](#)
Kamiar Rahnama Rad, Carl Smith, Liam Paninski, Columbia University 48

I-7. [Bayesian decision-making using neural attractor networks](#)
Ingmar Kanitscheider, Jeffrey Beck, Alexandre Pouget, University of Geneva 49

I-8. [Robust online estimation of noisy sparse signals by threshold-linear and integrate-and-fire neurons](#)
Tao Hu, Alexander Genkin, Dmitri Chklovskii, Janelia Farm Research Campus 49

I-9. [A generative model of natural images as patchworks of textures](#)
Matthias Bethge, Niklas Luedtke, Debapriya Das, Lucas Theis, Max Planck Institute 50

I-10. [A non-parametric Bayesian prior for causal inference of auditory streaming](#)
Ulrik Beierholm, Tim Yates, University of Birmingham 50

I-11. [Correlations strike back \(again\): the case of associative memory retrieval](#)
Cristina Savin, Peter Dayan, Mate Lengyel, University of Cambridge 51

I-12. [Timescale-dependent shaping of correlations by stimulus features in the auditory system](#)
Ana Calabrese, Sarah M N Woolley, Columbia University 51

I-13. [Compensation for neuronal circuit variability](#)
Tilman Kispersky, Eve Marder, Brandeis University 52

I-14. [Precise characterization of multiple LIP neurons in relation to stimulus and behavior](#)
Jacob Yates, Il M Park, Lawrence Cormack, Jonathan W Pillow, Alexander Huk, University of Texas at Austin 52

I-15. [Common neural correlates of adaptive control in the anterior cingulate cortex of rats and humans](#)
Nandakumar Narayanan, James Cavanagh, Marcelo Caetano, Michael Frank, Mark Laubach, University of Iowa 53

I-16. [Neural correlates of target detection in the human brain](#)
Arjun Bansal, Alexandra Golby, Joseph Madsen, Gabriel Kreiman, Boston Children's Hospital 53

I-17. [A computational theory of action-based decisions](#)
Vassilios Christopoulos, Paul Schrater, California Institute of Technology 54

I-18. [Individual FEF neurons code both task difficulty and an internally generated error signal](#)
Tobias Teichert, Dian Yu, Vincent Ferrera, Columbia University 55

I-19. [Perceptual decision-making in a sampling-based neural representation](#)
Ralf Haefner, Pietro Berkes, Jozsef Fiser, Brandeis University 55

I-20. [Model-free reinforcement learning predicts attentional selection during foraging in the macaque](#)
Matthew Balcarras, Salva Ardid, Daniel Kaping, Stefan Everling, Thilo Womelsdorf, York University 56

I-21. How can single neurons predict behavior? Xaq S Pitkow, Sheng Liu, Yong Gu, Dora Angelaki, Greg DeAngelis, Alexandre Pouget, University of Rochester	56
I-22. Confidence based learning in a perceptual task: how uncertainty and outcome influence choice Eric DeWitt, Andre G Mendonca, Adam Kepecs, Zachary F. Mainen, Champalimaud Neuroscience Programme	57
I-23. A model of interactive effects of learning and choice incentive in the striatal dopamine system Anne Collins, Michael Frank, Brown University	58
I-24. Circuit mechanisms underlying cognitive flexibility in rats Ioana Carcea, Natalya Zaika, Robert C. Froemke, New York University	58
I-25. Activity in mouse pedunculo-pontine tegmental nucleus reflects priors for action value John Thompson, Gidon Felsen, University of Colorado, Denver	59
I-26. Striatal optogenetic stimulation reveals distinct neural computation for value-based decision making A. Moses Lee, Lunghao Tai, Antonello Bonci, Linda Wilbrecht, University of California, San Francisco	59
I-27. Encoding of yaw in the presence of roll or pitch: Studies in a fly motion sensitive neuron Suva Roy, Shiva Sinha, Rob de Ruyter van Steveninck, Indiana University	60
I-28. Beyond Barlow: a Bayesian theory of efficient neural coding Jonathan W Pillow, Il M Park, University of Texas at Austin	60
I-29. Error statistics and error correction: Evidence for multiple coordinate encodings Todd Hudson, Michael Landy, New York University	61
I-30. Input dependence of local field potential spectra: experiment vs theory Francesca Barbieri, Alberto Mazzoni, Nikos K. Logothetis, Stefano Panzeri, Nicolas Brunel, ISI Foundation	61
I-31. Object-based spectro-temporal analysis of auditory signals Mikio Aoi, Yoonseob Lim, Uri Eden, Timothy Gardner, Boston University, Department of Math & Stats	62
I-32. Heterogeneity increases information transmission of neuronal populations Eric Hunsberger, Matthew Scott, Chris Eliasmith, University of Waterloo	63
I-33. Constraining cortical population coding of visual motion in area MT by smooth pursuit behavior Stephanie E Palmer, Leslie Osborne, University of Chicago	63
I-34. The role of Inhibitory STDP in disinaptic feedforward circuits Florence Kleberg, Matthieu Gilson, Tomoki Fukai, RIKEN BSI	64
I-35. Rapid development of feed forward inhibition drives emergence of the alert cortical state Matthew Colonnese, George Washington University	64
I-36. Laminar differences in receptive fields measured by reverse correlation of intracellular recordings. Alejandro Ramirez, Eftychios A. Pnevmatikakis, Josh Merel, Ken Miller, Liam Paninski, Randy Bruno, Columbia University	65
I-37. T-type calcium channels promote predictive homeostasis in thalamocortical neurons of LGN Su Hong, Haram Kim, Sora Yun, Christopher Fiorillo, KAIST	65
I-38. Neural implementations of motion detection in the fly medulla: insights from connectivity and theory Arjun Bharioke, Shin-ya Takemura, Shiv Vitaladevuni, Richard Fetter, Zhiyuan Lu, Stephen Plaza, Louis Scheffer, Ian Meinertzhagen, Dmitri Chklovskii, Janelia Farm Research Campus, HHMI	66
I-39. Evidence for a nonlinear coupling between firing threshold and subthreshold membrane potential Skander Mensi, Christian Pozzorini, Olivier Hagens, Wulfram Gerstner, EPFL	67
I-40. Spatial distribution and efficacy of thalamocortical synapses onto layer 4 excitatory neurons Carl Schoonover, Juan Tapia, Verena Schilling, Verena Wimmer, Richard Blazeski, Carol Mason, Randy Bruno, Columbia University	67
I-41. A plasticity mechanism for the robustness of the oculomotor integrator Pedro Goncalves, University College, London	68

I-42. [Robust formation of continuous attractor networks](#)
Daniel Robles-Llana, Ila Fiete, University of Texas at Austin 68

I-43. [Connectomic constraints on computation in networks of spiking neurons](#)
Venkatakrisnan Ramaswamy, Arunava Banerjee, The Hebrew University of Jerusalem 69

I-44. [Dynamics of random clustered networks](#)
Merav Stern, Haim Sompolinsky, Larry Abbott, Columbia University 69

I-45. [NeuroElectro: a community database on the electrophysiological diversity of mammalian neuron types](#)
Shreejoy J Tripathy, Richard Gerkin, Judy Savitskaya, Nathaniel Urban, Carnegie Mellon University 70

I-46. [Traveling waves of the hippocampal theta rhythm encode rat position](#)
Gautam Agarwal, Ian Stevenson, Kenji Mizuseki, Gyorgy Buzsaki, Friedrich Sommer, Redwood Center for Theoretical Neuroscience 71

I-47. [Are aperiodic 1D grid cell responses consistent with low-dimensional continuous attractor dynamics?](#)
KiJung Yoon, Daniel Robles-Llana, Amina Kinkhabwala, David Tank, Ila Fiete, University of Texas at Austin 71

I-48. [A single computational mechanism for both stability and flexibility in vocal error correction](#)
Samuel Sober, Conor Kelly, Emory University 72

I-49. [Optimal foraging and multiscale representation by hippocampal place cells](#)
Jason Prentice, Vijay Balasubramanian, Princeton University 72

I-50. [Slow time-scales in songbird neural sequence generation.?](#)
Jeffrey Markowitz, Gregory Guitchounts, Timothy Gardner, Boston University 73

I-51. [Criterion dynamics in the ferret optimise decision bias during an auditory detection task](#)
Robert W Mill, Christian Sumner, MRC Institute of Hearing Research 73

I-52. [A multiplicative learning rule for auditory discrimination learning in mice](#)
Brice Bathellier, Sui Poh Tee, Christina Hrovat, Simon Rumpel, Research Institute of Molecular Pathology 74

I-53. [Strategies for encoding competing acoustic events by single neurons in primate auditory cortex](#)
Yi Zhou, Xiaojin Wang, Arizona State University 75

I-54. [View-invariance and mirror-symmetric tuning in a model of the macaque face-processing system](#)
Joel Z Leibo, Fabio Anselmi, Jim Mutch, Akinori Ebihara, Winrich Freiwald, Tomaso Poggio, Massachusetts Institute of Technology 75

I-55. [Ventral stream models that solve hard object recognition tasks naturally exhibit neural consistency.](#)
Daniel Yamins, Ha Hong, Ethan Solomon, James DiCarlo, Massachusetts Institute of Technology 76

I-56. [Similarity between spontaneous and sensory-evoked activity does suggest learning in the cortex](#)
Cristina Savin, Pietro Berkes, Chiayu Chiu, Jozsef Fiser, Mate Lengyel, University of Cambridge 76

I-57. [Uncertainty and the striatum: How tonically active neurons may aid learning in dynamic environments.](#)
Nicholas Franklin, Michael Frank, Brown University 77

I-58. [How neurogenesis and the scale of the dentate gyrus affect the resolution of memories](#)
James Aimone, Craig Vineyard, Sandia National Laboratories 77

I-59. [A not so bad tradeoff: perception and generalization of aversive memories](#)
Rony Paz, Weizmann Institute 78

I-60. [Spatial learning in *C. elegans*](#)
Adam Calhoun, Sreekanth Chalasani, Tatyana O. Sharpee, University of California, San Diego 78

I-61. [The effect of STDP temporal structure on the learning of single excitatory and inhibitory synapses](#)
Maoz Shamir, Yotam Luz, Ben Gurion University of the Negev 79

I-62. [Quantifying representational and dynamical structure in large neural datasets](#)
Jeffrey Seely, Matthew Kaufman, Adam Kohn, Matthew Smith, Anthony Movshon, Nicholas Priebe, Stephen Lisberger, Stephen Ryu, Krishna Shenoy, Larry Abbott, Columbia University 80

I-63. Motor coding in the supplementary motor area of humans and monkeys Gal Chechik, Hadas Taubman, Ariel Tankus, Eilon Vaadia, Itzhak Fried, Rony Paz, Gonda Brain Research Center	80
I-64. Transient collective dynamics in inhibition-stabilized motor circuits Guillaume Hennequin, Tim P Vogels, Wulfram Gerstner, School of Computer and Communication Sciences	81
I-65. Eye movements depend on the intrinsic reward structure in a natural navigation task Constantin A Rothkopf, Dana Ballard, Frankfurt Institute for Advanced Studies	82
I-66. Behavior-modulated correlation of cerebellar Purkinje neuron and network activity Sungho Hong, Mario Negrello, Erik De Schutter, Okinawa Institute of Science and Technology	82
I-67. Neuronal avalanches in the resting meg of the human brain Oren Shriki, Jeffrey Alstott, Frederick Carver, Tom Holroyd, Richard Henson, Marie Smith, Edward Bullmore, Richard Coppola, Dietmar Plenz, NIMH/NIH	83
I-68. Network dynamics amplifies the effects of weak electric fields on gamma and slow-wave activity Lucas C Parra, Davide Reato, The City College of CUNY	83
I-69. Correspondence between perceptual salience of 4th-order visual textures and natural scene statistics John Briguglio, Ann Hermundstad, Mary M. Conte, Jonathan Victor, Gasper Tkacik, Vijay Balasubramanian, University of Pennsylvania	84
I-70. Multiple temporally-tuned mechanisms control visual adaptation to contrast Stephen Engel, Min Bao, University of Minnesota	85
I-71. When is sensory precision variable? Shan Shen, Ronald Van den Berg, Wei Ji Ma, Baylor College of Medicine	85
I-72. Adaptive shaping of feature selectivity in the rodent vibrissa system He Zheng, Douglas Ollerenshaw, Qi Wang, Garrett Stanley, Georgia Institute of Technology	86
I-73. Neural coding of perceived odor intensity Yevgeniy Sirotnin, Roman Shusterman, Dmitry Rinberg, The Rockefeller University	86
I-74. How is duration information from multiple sensory sources combined? Mingbo Cai, David Eagleman, Baylor College of Medicine	87
I-75. Controlling the trade-off between categorization and separation via resonance in the olfactory bulb Maximilian Puelma Touzel, Michael Monteforte, Fred Wolf, Max Planck Institute for Dynamics and Self-Or	87
I-76. Multisensory integration of vision and intracortical microstimulation for sensory feedback. Maria Dadarlat, Joseph O'Doherty, Philip N. Sabes, University of California, San Francisco	88
I-77. Predicting V1 neural responses to natural movies using the shift-invariant bispectrum Mayur Mudigonda, Ian Stevenson, Urs Koster, Christopher Hillar, Charles Gray, Bruno Olshausen, Redwood Center For Theoretical Neuroscience	89
I-78. Spatial structure and organization of nonlinear subunits in primate retina Jeremy Freeman, Greg Field, Peter Li, Martin Greschner, Lauren Jepson, Neil Rabinowitz, Eftychios A. Pnevmatikakis, Deborah Gunning, Keith Mathieson, Alan Litke, New York University	89
I-79. Modulatory signals from V1 extra-classical receptive fields with distinct spatio-temporal dynamics Christopher Henry, Michael Hawken, New York University	90
I-80. Silencing V2/V3 reduces spiking variability in MT: implications for excitatory/inhibitory balance Camille Gomez, Alexandra Smolyanskaya, Gabriel Kreiman, Richard Born, Harvard Medical School	91
I-81. Distinct neural selectivity for 3D directions of visual motion. Thaddeus B Czuba, Lawrence Cormack, Alexander Huk, Adam Kohn, Albert Einstein College of Medicine	91
I-82. Maximum-entropy modeling of population calcium activity in the visual cortex Dimitri Yatsenko, Kresimir Josic, Andreas Tolias, Baylor College of Medicine	92

Posters I

I-83. A spatio-temporal lattice filter model for the visual pathway Hervé Rouault, Karol Gregor, Dmitri Chklovskii, Janelia Farm Research Campus	92
I-84. 3D Random access ultrafast two-photon imaging reveals the structure of network activity. R. James Cotton, Emmanouil Froudarakis, Peter Saggau, Andreas Tolias, Baylor College of Medicine	93
I-85. Energy and Information in insect photoreceptors Nikon Rasumov, Jeremy E. Niven, Simon Laughlin, University of Cambridge	93
I-86. Kernel regression for receptive field estimation Urs Koster, Bruno Olshausen, Christopher Hillar, University of California, Berkeley	94
I-87. Trade-off between curvature tuning and position invariance in visual area V4 Tatyana O. Sharpee, Minjoon Kouh, John Reynolds, Salk Institute for Biological Studies	95
I-88. Multi-stability in motion perception combines multiple underlying neural mechanisms Andrew Meso, James Rankin, Pierre Kornprobst, Guillaume S Masson, CNRS/Aix-Marseille Université	95
I-89. Changes in laminar synchrony in V1 reflect perceptual decisions. Neda Shahidi, Ariana Andrei, Ming Hu, Valentin Dragoi, University of Texas	96
I-90. Lag normalization in an electrically-coupled neural network Gautam Awatramani, David Schwab, Stuart Trenholm, Vijay Balasubramanian, University of Victoria	96
I-91. Generalized reward-modulated STDP: a model of operant conditioning of cortical neurons Robert Kerr, David B. Grayden, Doreen A. Thomas, Matthieu Gilson, Anthony N. Burkitt, University of Melbourne	97
I-92. Network motifs and collective spiking in neuronal networks Yu Hu, James Trousdale, Kresimir Josic, Eric Shea-Brown, University of Washington	98
I-93. Quadratic programming of tuning curves: from spiking dynamics to function David GT Barrett, Sophie Deneve, Christian Machens, École Normale Supérieure	98
I-94. Learning activity patterns in recurrent networks of visible and hidden spiking neurons Johanni Brea, Walter Senn, Jean-Pascal Pfister, University of Bern	99
I-95. Associative memory encoding in bump attractor networks Vladimir Itskov, Carina Curto, University of Nebraska - Lincoln	99
I-96. Recursive conditional means (RCM): A powerful mechanism for combining sensory estimates Wilson Geisler, Jeffrey Perry, University of Texas at Austin	100
I-97. Energy constraints link structure and function in thin axons in the brain's wiring Ali Neishabouri, Aldo Faisal, Imperial College London	100
I-98. Adaptive gaussian poisson process: a model for in vivo neuronal dynamics Simone Carlo Surace, Jean-Pascal Pfister, University of Bern	101
I-99. Rapid update in state estimation accounts for fast feedback control Frederic Crevecoeur, Stephen H Scott, Centre for Neurosciences Studies	101
I-100. Dynamical entropy production in cortical circuits with different network topologies Rainer Engelken, Michael Monteforte, Fred Wolf, MPI for Dynamics and Self-Organisation	102

Poster Session II

7:30 pm Friday 1st March

II-1. Long-term memory . . . now longer than ever Marcus Benna, Stefano Fusi, Columbia University	102
II-2. Here's Waldo! A mechanistic model of visual search predicts human behavior in an object search task Thomas Miconi, Children's Hospital Boston	103
II-3. NMDA-mediated feedback accounts for effects of visual spatial attention in Neurogrid simulations Nicholas Steinmetz, Ben Benjamin, Kwabena Boahen, Stanford University	103
II-4. The laminar origin of sensitivity to high-order image statistics in macaque visual cortex Yunguo Yu, Anita Schmid, Jonathan Victor, Weill Cornell Medical College	104
II-5. A biophysical model of Bayesian inference and MCMC sampling in neural circuits Alessandro Ticchi, Aldo Faisal, Imperial College London	104
II-6. Adaptive estimation of firing rate maps under super-Poisson variability Mijung Park, J. Patrick Weller, Gregory Horwitz, Jonathan W Pillow, University of Texas at Austin	105
II-7. Introducing MEDAL: a Matlab Environment for Deep Architecture Learning Dustin Stansbury, Jack Gallant, University of California, Berkeley	105
II-8. Encoding the direction of ego-motion in the fly brain James Trousdale, Sam Carroll, Kresimir Josic, University of Houston	106
II-9. Representing correlated retinal population activity with Restricted Boltzmann Machines David Schwab, Kristina Simmons, Jason Prentice, Vijay Balasubramanian, Princeton University	106
II-10. Dendritic nonlinearities shape beyond-pairwise correlations and improve (spiking) population coding Joel Zylberberg, N. Alex Cayco-Gajic, Eric Shea-Brown, University of Washington	107
II-11. The perils of inferring information from correlations Jeffrey Beck, Ingmar Kanitscheider, Xaq S Pitkow, Peter Latham, Alexandre Pouget, University of Rochester	108
II-12. Temporal decorrelation by power-law adaptation in pyramidal neurons Christian Pozzorini, Richard Naud, Skander Mensi, Wulfram Gerstner, EPFL, BMI	108
II-13. Theory of higher-order correlations in neural network with columnar structure Yasuhiko Igarashi, Masato Okada, University of Tokyo	109
II-14. Perceptual decisions are limited primarily by variability in early sensory cortex Charles Michelson, Jonathan W Pillow, Eyal Seidemann, University of Texas at Austin	109
II-15. Different response properties of rat parietal and frontal cortices during evidence accumulation Timothy Hanks, Chunyu Ann Duan, Jeffrey Erlich, Bingni Brunton, Carlos Brody, HHMI	110
II-16. Orbitofrontal cortex is sensitive to natural behavioral categories during social exploration Geoffrey Adams, John Pearson, Michael Platt, Duke University	110
II-17. Humans exploit the uncertainty in priors to improve direction perception. Steeve Laquitaine, Justin Gardner, Riken BSI	111
II-18. Hebbian mechanisms underlying the learning of Markovian sequence probabilities Kristofer Bouchard, Surya Ganguli, Michael Brainard, University of California, San Francisco	111
II-19. Controllability and resource-rational planning Falk Lieder, University of Zurich	112
II-20. Neural responses in the rat parietal cortex during decision formation and movement David Raposo, Matthew Kaufman, Anne Churchland, Cold Spring Harbor Laboratory	112
II-21. Do humans account for stimulus correlations in visual perception? Manisha Bhardwaj, Ronald Van den Berg, Wei Ji Ma, Kresimir Josic, University of Houston	113

Posters II

II-22. Neural correlates of arbitration between model-based and model-free reinforcement learning systems	
Sang Wan Lee, Shinsuke Shimojo, John O'Doherty, California Institute of Technology	113
II-23. The neural correlates of counterfactual-Q-learning in a strategic sequential investment task	
Rong Guo, Michael Tobia, Wendelin Böhmer, Tobias Sommer, Christian Büchel, Klaus Obermayer, Technische Universität Berlin	114
II-24. Dynamics of decision and action in rat posterior parietal cortex	
Matthew Kaufman, David Raposo, Anne Churchland, Cold Spring Harbor Laboratory	115
II-25. Normalization predicted by optimal inference in sensory systems	
Matthew Chalk, Paul Masset, Boris Gutkin, Sophie Deneve, École Normale Supérieure	115
II-26. The synaptic sampling hypothesis	
Laurence Aitchison, Peter Latham, University College, London	116
II-27. Information-theoretic limits on encoding over diverse populations	
O. Ozan Koyluoglu, Ila Fiete, University of Texas at Austin	116
II-28. Matching encoding and decoding with spiking neurons	
Fleur Zeldenrust, Sophie Deneve, Boris Gutkin, École Normale Supérieure	117
II-29. Thalamic synchrony drives cortical feature selectivity in standard and novel visual stimuli	
Sean Kelly, Jens Kremkow, Jianzhong Jin, Yushi Wang, Qi Wang, Jose-Manuel Alonso, Garrett Stanley, Georgia Institute of Technology	117
II-30. Encoding and decoding stimuli that generate persistent activity	
Kyriaki Sidiropoulou, Panayiota Poirazi, University of Crete	118
II-31. Evaluation of single unit error contribution to neural state space dynamics in linear BMI decoders	
Islam Badreldin, Karthikeyan Balasubramanian, Mukta Vaidya, Joshua Southerland, Andrew Fagg, Nicholas Hatsopoulos, Karim Oweiss, Michigan State University	118
II-32. The geometry of excitation-inhibition balance in human neocortex	
Nima Dehghani, Adrien Peyrache, Eric Halgren, Sydney Cash, Alain Destexhe, Harvard University	119
II-33. State-dependent impact of distinct interneuron types on visual contrast gain	
Ulf Knoblich, Jessica Cardin, Yale University	119
II-34. Receptive field formation by interacting excitatory and inhibitory plasticity	
Claudia Clopath, Henning Sprekeler, Tim P Vogels, Columbia University	120
II-35. In vivo dissection of layer 1 inputs in the barrel cortex	
Wanying Zhang, Randy Bruno, Columbia University	120
II-36. Out of the zoo? Interneuron subtypes encode specific behavioral variables in the cingulate cortex	
Sachin Ranade, Duda Kvitsiani, Balazs Hangya, Z. Josh Huang, Adam Kepecs, Cold Spring Harbor Laboratory	121
II-37. A cellular mechanism for system memory consolidation	
Urs Bergmann, Michiel Remme, Susanne Schreiber, Henning Sprekeler, Richard Kempster, Humboldt University Berlin	121
II-38. Temporally evolving surround suppression helps decoding in a spiking model of motion processing	
Philip Meier, Micah Richert, Jayram Moorkanikara Nageswaran, Eugene Izhikevich, Brain Corporation	122
II-39. Functional bases for multidimensional neural computations	
Joel Kaardal, Jeff Fitzgerald, Michael J. Berry II, Tatyana O. Sharpee, Salk Institute for Biological Studies	123
II-40. Synaptic plasticity shapes the excitation/inhibition balance during ongoing cortical activity	
Aryeh Taub, Yonatan Katz, Ilan Lampl, Weizmann Institute	123
II-41. Attention improves information processing by tuning cortical networks towards critical states	
Nergis Toemen, Udo Ernst, University of Bremen	124

[II-42. Beyond acute experiments: automated long-term tracking of socially-housed mice](#)
 Shay Ohayon, Ofer Avni, Adam Taylor, Roian Egnor, Pietro Perona, California Institute of Technology . . . 124

[II-43. Modeling adaptive changes in the motor program underlying birdsong](#)
 Baktash Babadi, Bence Olveczky, Harvard University 125

[II-44. Do orientation preference maps arise from hexagonal retinal ganglion cell mosaics?](#)
 Manuel Schottdorf, Wolfgang Keil, Michael Schnabel, David M. Coppola, Siegrid Löwel, Leonard E. White, Matthias Kaschube, Fred Wolf, Max Planck Inst. Dynamics & Self-organization 125

[II-45. Loss of theta modulation of the hippocampal cell firing is accompanied by deterioration of episode f](#)
 Yingxue Wang, Eva Pastalkova, Janelia Farm Research Campus 126

[II-46. The sense of place: grid cells in the brain and the transcendental number e](#)
 Vijay Balasubramanian, Jason Prentice, Xuexin Wei, University of Pennsylvania 127

[II-47. Characteristics prediction of STDP in hippocampal CA1 network by mutual information maximization](#)
 Ryota Miyata, Keisuke Ota, Toru Aonishi, Tokyo Institute of Technology 127

[II-48. Time-scales of neural integration constrain a songbird operant behavior](#)
 Yoonseob Lim, Barbara Shinn-Cunningham, Timothy Gardner, Boston University 128

[II-49. Unsupervised learning of binaural features from naturalistic stereo sounds](#)
 Wiktor Młynarski, Max Planck Institute for Math in the Sciences 128

[II-50. Closing the loop; Inverse-model learning with a nonlinear avian syrinx and sparse auditory coding](#)
 Alexander Hanuschkin, Surya Ganguli, Richard Hahnloser, INI, University of Zurich and ETH Zurich . . . 129

[II-51. Timing of invariant object recognition in humans](#)
 Leyla Isik, Ethan Meyers, Joel Z Leibo, Tomaso Poggio, Massachusetts Institute of Technology 129

[II-52. Object-vision models that better explain IT also categorize better, but all models fail at both](#)
 Seyed-Mahdi Khaligh-Razavi, Nikolaus Kriegeskorte, University of Cambridge 130

[II-53. Size-invariant shape coding in visual area V4](#)
 Yasmine El-Shamayleh, Anitha Pasupathy, University of Washington 130

[II-54. A general theory of learning and memory with complex synapses](#)
 Subhaneil Lahiri, Surya Ganguli, Stanford University 131

[II-55. Multi-step decision tasks for dissociating model-based and model-free learning in rodents.](#)
 Thomas Akam, Peter Dayan, Rui Costa, Champalimaud Neuroscience Program 131

[II-56. Ghrelin modulates phasic dopamine evoked by food-reward via action in the lateral hypothalamus](#)
 Jackson Cone, Mitchell Roitman, University of Illinois at Chicago 132

[II-57. Dopamine modulates functional communication of the human striatum subdivision](#)
 Payam Piray, Marieke van der Schaaf, Ivan Toni, Roshan Cools, Radboud University Nijmegen 133

[II-58. Volitional control by a person with tetraplegia of high gamma local field potentials \(LFPs\) recorded](#)
 Tomislav Milekovic, Daniel Bacher, Blaise Yvert, Leigh Hochberg, Emad Eskandar, Sydney Cash, John Donoghue, Brown University 133

[II-59. Value-updating interaction among contexts in choice behaviors of rats](#)
 Akihiro Funamizu, Makoto Ito, Kenji Doya, Ryohei Kanzaki, Hirokazu Takahashi, JSPS research fellow . . 134

[II-60. Solving the secretary problem](#)
 Vincent Costa, Bruno Averbeck, NIMH/NIH 134

[II-61. Single neuron contributions to motor variability](#)
 Kris Chaisanguanthum, Helen Shen, Philip N. Sabes, University of California, San Francisco 135

[II-62. An origin for coarticulation of speech sequences in human sensory-motor cortex](#)
 Kristofer Bouchard, Keith Johnson, Edward Chang, University of California, San Francisco 135

[II-63. Large-scale optical imaging reveals structured network output in isolated spinal cord](#)
 Timothy Machado, Liam Paninski, Thomas M. Jessell, Columbia University 136

[II-64. Novel motion illusion evidence for perception-action coupling](#)
Dongsung Huh, University College, London 137

[II-65. A brain-machine interface for control of medically-induced coma](#)
Maryam Shanechi, Jessica Chemali, Max Liberman, Ken Solt, Emery Brown, Massachusetts Institute of Technology 137

[II-66. Spike-triggered local field potential as an indicator of the human epileptogenic zone](#)
Beth Lopour, Itzhak Fried, Dario Ringach, University of California, Los Angeles 138

[II-67. Directed communication-through-coherence during synchronous transients](#)
Agostina Palmigiano, Annette Witt, Demian Battaglia, Theo Geisel, MPI for Dynamics and Self-Organization 138

[II-68. State dynamics of the epileptic brain and the influence of seizure focus](#)
Sridevi Sarma, Sam Burns, Sabato Santaniello, William Stan Anderson, Johns Hopkins University 139

[II-69. The role of adaptation in intrinsic dynamics of primary visual cortex](#)
Yashar Ahmadian, Sandro Romani, Amiram Grinvald, Misha Tsodyks, Ken Miller, Columbia University . . 139

[II-70. Strategies for optomotor control in free flying Drosophila](#)
Leonardo Mesquita, Shiva Sinha, Rob de Ruyter van Steveninck, Indiana University 140

[II-71. Supervised cue calibration relies on the multisensory percept](#)
Adam Zaidel, Wei Ji Ma, Dora Angelaki, Baylor College of Medicine 140

[II-72. Population decoding in rat barrel cortex](#)
Ehsan Arabzadeh, Mehdi Adibi, Colin WG Clifford, Australian National University 141

[II-73. Hemodynamic responses in the somatosensory cortex during locomotion](#)
Bingxing Huo, Yurong Gao, Patrick Drew, Pennsylvania State University 142

[II-74. Rapid detection of odors based on fast, sampling-based inference](#)
Agnieszka Grabska-Barwinska, Jeffrey Beck, Alexandre Pouget, Peter Latham, University College, London 142

[II-75. Neuronal nonlinearity explains greater visual spatial resolution for dark than for light stimuli](#)
Jens Kremkow, Jianzhong Jin, Stanley Jose Komban, Yushi Wang, Reza Lashgari, Michael Jansen, Xiaobing Li, Qasim Zaidi, Jose-Manuel Alonso, SUNY-Optometry 143

[II-76. Encoding of natural scene statistics in the primary visual cortex of the mouse](#)
Emmanouil Froudarakis, Philipp Berens, R. James Cotton, Alexander S. Ecker, Peter Saggau, Matthias Bethge, Andreas Tolias, Baylor College of Medicine 143

[II-77. Recording the entire visual representation along the vertical pathway in the mammalian retina](#)
Philipp Berens, Tom Baden, Matthias Bethge, Thomas Euler, BCCN 144

[II-78. Spontaneous emergence of simple and complex receptive fields in a spiking model of V1](#)
Eugene Izhikevich, Filip Piekiewicz, Jayram Moorkanikara Nageswaran, Csaba Petre, Micah Richert, Sach Sokol, Philip Meier, Marius Buiabas, Dmitry Fisher, Botond Szatmari, Brain Corporation 145

[II-79. Eye's imaging process explains ganglion cells anisotropies](#)
Daniela Pamplona, Jochen Triesch, Constantin A Rothkopf, Frankfurt Institute for Advanced Studies . . . 145

[II-80. Circuit mechanisms revealed by spike-timing correlations in macaque area MT](#)
Xin Huang, University of Wisconsin, Madison 146

[II-81. Hierarchical shape processing and position tolerance in rat lateral visual cortex](#)
Ben Vermaercke, Gert Van den Bergh, Florian Gerich, Hans P. Op de Beeck, KU Leuven 146

[II-82. Anesthesia amplifies visual responses by suppressing cortical state dynamics](#)
Flavio Frohlich, Kristin Sellers, Davis Bennett, UNC - Chapel Hill 147

[II-83. Computational models of contour detection: role of lateral connections, inhibition and normalization](#)
David A Mély, Thomas Serre, Brown University 148

[II-84. Is there a critical area size for the transition from interspersed to columnar V1 architecture?](#)
Wolfgang Keil, Fred Wolf, Juan-Daniel Florez-Weidinger, Matthias Kaschube, Michael Schnabel, David M. Coppola, Siegrid Löwel, Leonard E. White, MPI for Dynamics & Self-organization 148

II-85. [Modeling cortical responses to mixture stimuli reveals origins of orientation tuning variation](#)
 Robbe Goris, Eero Simoncelli, Anthony Movshon, New York University 149

II-86. [Psychophysical evidence for a sampling-based representation of uncertainty in low-level vision](#)
 Marjena Popovic, Ralf Haefner, Mate Lengyel, Jozsef Fiser, Brandeis University 149

II-87. [Effects of attention on spatio-temporal correlations across layers of a single column in area V4](#)
 Tatiana Engel, Nicholas Steinmetz, Tirin Moore, Kwabena Boahen, Stanford University 150

II-88. [Population codes for topography in the zebrafish optic tectum](#)
 Lilach Avitan, Zac Pujuc, Ethan Scott, Geoffrey Goodhill, University of Queensland, Australia 151

II-89. [Robust estimation for neural state-space models](#)
 Lars Buesing, Jakob H Macke, Maneesh Sahani, University College, London 151

II-90. [A design procedure for hierarchical neural control.](#)
 Gregory Wayne, Larry Abbott, Columbia University 152

II-91. [Robust learning of low dimensional dynamics from large neural ensembles](#)
 David Pfau, Eftychios A. Pnevmatikakis, Liam Paninski, Columbia University 152

II-92. [Beyond GLMs: a generative mixture modeling approach to neural system identification](#)
 Lucas Theis, Dan Arnstein, André Maia Chagas, Cornelius Schwarz, Matthias Bethge, Centre for Integrative Neuroscience 153

II-93. [Low-rank connectivity induces firing rate fluctuations in a chaotic spiking model](#)
 Brian DePasquale, Mark M Churchland, Larry Abbott, Columbia University 153

II-94. [Got a moment or two? Neural models and linear dimensionality reduction](#)
 Il M Park, Evan Archer, Nicholas Priebe, Jonathan W Pillow, University of Texas at Austin 154

II-95. [Spike train entropy-rate estimation using hierarchical Dirichlet process priors](#)
 Karin Knudson, Jonathan W Pillow, University of Texas at Austin 154

II-96. [The neural ring: an algebraic tool for analyzing the intrinsic structure of neural codes](#)
 Nora Youngs, Alan Veliz-Cuba, Vladimir Itskov, Carina Curto, University of Nebraska - Lincoln 155

II-97. [Flexible probabilistic approach for measuring model learning independent of behavioral improvement.](#)
 Nathaniel Powell, Paul Schrater, University of Minnesota 155

II-98. [Fast missing mass approximation for the partition function of stimulus driven Ising models](#)
 Robert Haslinger, Demba Ba, Ziv Williams, Gordon Pipa, Massachusetts Institute of Technology 156

II-99. [Diversity of timescales in network activity](#)
 Rishidev Chaudhuri, Alberto Bernacchia, Xiao-Jing Wang, Yale University 156

II-100. [Modeling inter-neuron inhibition with determinantal point processes](#)
 Jasper Snoek, Ryan P Adams, Richard Zemel, University of Toronto 157

III-1. Frontal neurons enable retrieval of memories over widely varying temporal scales Ziv Williams, Harvard Medical School	157
III-2. The neuronal input channel switched by attention reflects Routing by Coherence Iris Grothe, David Rotermund, Simon Neitzel, Sunita Mandon, Udo Ernst, Andreas Kreiter, Klaus Richard Pawelzik, Strüngmann Institute	158
III-3. Optimal speed estimation in natural image movies Johannes Burge, Wilson Geisler, University of Texas at Austin	159
III-4. Dynamic calibration of the influence of priors and sensory input for perceptual estimation Kamesh Krishnamurthy, Matthew Nassar, Shilpa Sarode, Joshua Gold, University of Pennsylvania	159
III-5. Constraining a Bayesian model of orientation perception with efficient coding Xuexin Wei, Alan Stocker, University of Pennsylvania	160
III-6. Some work and some play: a normative, microscopic approach to allocating time between work & leisure Ritwik Niyogi, Yannick Breton, Rebecca Solomon, Kent Conover, Peter Shizgal, Peter Dayan, University College, London	160
III-7. Optimally fuzzy memory Karthik Shankar, Center for Memory and Brain	161
III-8. Sparse Bayesian inference and experimental design for synaptic weights and locations Ari Pakman, Carl Smith, Liam Paninski, Columbia University	161
III-9. Inferring functional connectivity with priors on network topology Scott Linderman, Ryan P Adams, Harvard University	162
III-10. Recurrent generalized linear models with correlated Poisson observations Marius Pachitariu, Lars Buesing, Maneesh Sahani, University College, London	162
III-11. Distinct coherent ensembles reflect working memory processes in primate PFC David Markowitz, Bijan Pesaran, New York University	163
III-12. Dimensionality, dynamics, and correlations in the motor cortical substrate for reaching Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, Surya Ganguli, Stanford University	163
III-13. Optimal neural tuning for arbitrary stimulus priors with Gaussian input noise Zhuo Wang, Alan Stocker, Haim Sompolinsky, Daniel Lee, University of Pennsylvania	164
III-14. From metacognition to statistics: relating confidence reports across species Balazs Hangya, Joshua Sanders, Adam Kepecs, Cold Spring Harbor Laboratory	165
III-15. Eyetracking and pupillometry reflect dissociable measures of latent decision processes James Cavanagh, Thomas Wiecki, Angad Kochar, Michael Frank, Brown University	165
III-16. Deviations from the matching law reflect reward integration over multiple timescales Kiyohito Iigaya, Leo P. Sugrue, Greg S. Corrado, Yonatan Loewenstein, William T. Newsome, Stefano Fusi, Columbia University	166
III-17. Changing subjects' fallacies by changing their beliefs about the dynamics of their environment Friederike Schuur, Brian P Tam, Laurence T Maloney, New York University	166
III-18. Past failures bias human decisions Arman Abrahamyan, Justin Gardner, RIKEN Brain Science Institute	167
III-19. The effects of prior, rewards and sensory evidence in perceptual decision making Yanping Huang, Rajesh Rao, University of Washington	167
III-20. The effect of novelty-based exploration on reinforcement learning in humans. Audrey Houillon, Robert Lorenz, Andreas Heinz, Jürgen Gallinat, Klaus Obermayer, BCCN	168

III-21. How does the brain compute value? A rational model for preference formation. Paul Schrater, Nisheeth Srivastava, University of Minnesota	169
III-22. Phasic locus coeruleus activity changes with practice: a pupillometry study Alexander Petrov, Taylor R. Hayes, Ohio State University	169
III-23. Dynamic integration of sensory evidence and diminishing reward in perceptual decision making Alireza Soltani, Ya-Hsuan Liu, Shih-Wei Wu, Stanford University	170
III-24. Neuronal signatures of strategic decisions in posterior cingulate cortex David Barack, Michael Platt, Duke University	170
III-25. Reconciling decisions from description with decisions from experience Komal Kapoor, Nisheeth Srivastava, Paul Schrater, University of Minnesota	171
III-26. Hierarchy of intrinsic timescales across primate cortex John Murray, Alberto Bernacchia, Tatiana Pasternak, Camillo Padoa-Schioppa, Daeyeol Lee, Xiao-Jing Wang, Yale University	171
III-27. Sparse coding and dynamic suppression of variability in balanced cortical networks Farzad Farkhooi, Martin Nawrot, FU Berlin	172
III-28. Combining feed-forward processing and sampling for neurally plausible encoding models Jorg Lucke, Jacquelyn A. Shelton, Jorg Bornschein, Philip Sterne, Pietro Berkes, Abdul-Saboor Sheikh, Goethe-University of Frankfurt	172
III-29. Visual target signals are computed via a dynamic 'and-like' mechanism in IT and perirhinal cortex Marino Pagan, Nicole Rust, University of Pennsylvania	173
III-30. On the information capacity of spike trains and detectability of rate fluctuations Shinsuke Koyama, The Institute of Statistical Mathematics	173
III-31. Response properties of sensory neurons artificially evolved to maximize information Brian Monson, Yaniv Morgenstern, Dale Purves, Duke University	174
III-32. A novel method for fMRI analysis: Inferring neural mechanisms from voxel tuning Rosemary Cowell, David Huber, John Serences, University of California, San Diego	174
III-33. Successful prediction of a physiological circuit with known connectivity from spiking activity Felipe Gerhard, Tilman Kispersky, Gabrielle J Gutierrez, Eve Marder, Mark Kramer, Uri Eden, Brain Mind Institute	175
III-34. Evidence for presynaptic inhibition in shaping retinal processing Yuwei Cui, Yanbin Wang, Jonathan Demb, Daniel A Butts, University of Maryland	176
III-35. Memory maintenance in calcium-based plastic synapses in the presence of background activity David Higgins, Michael Graupner, Nicolas Brunel, École Normale Supérieure	176
III-36. Dendritic subunits: the crucial role of input statistics and a lack of two-layer behavior DJ Strouse, Balazs B Ujfalussy, Mate Lengyel, Princeton University	177
III-37. Physiology and impact of horizontal connections in rat neocortex Philipp Schnepel, Ad Aertsen, Arvind Kumar, Clemens Boucsein, Bernstein Center Freiburg	177
III-38. Gain-control via shunting-inhibition in a spiking-model of leech local-bend. Edward Frady, William Kristan, University of California, San Diego	178
III-39. Detecting and quantifying topographic order in neural maps Stuart Yarrow, Khaleel Razak, Aaron Seitz, Peggy Series, University of Edinburgh	178
III-40. Efficient and optimal Little-Hopfield auto-associative memory storage using minimum probability flow Christopher Hillar, Jascha Sohl-Dickstein, Kilian Koepsell, Redwood Center for Theoretical Neuroscience	179
III-41. Signal processing in neural networks that generate or receive noise Isao Nishikawa, Kazuyuki Aihara, Taro Toyoizumi, University of Tokyo	179

III-42. [A spiking model of superior colliculus for bottom-up saliency](#)
 Jayram Moorkanikara Nageswaran, Micah Richert, Csaba Petre, Filip Piekiewicz, Sach Sokol, Botond Szatmary, Eugene Izhikevich, Brain Corporation 180

III-43. [Using optogenetics to probe neural circuit mechanisms in the alert, behaving non-human primate](#)
 Jonathan Nassi, Ali Cetin, Anna Roe, John Reynolds, Salk Institute for Biological Studies 180

III-44. [Being balanced: the role of the indirect pathway of the basal ganglia in threshold detection](#)
 Wei Wei, Jonathan Rubin, Xiao-Jing Wang, Yale University 181

III-45. [The amygdala is critical for reward encoding in the orbital, but not medial, prefrontal cortex](#)
 Peter Rudebeck, Andrew Mitz, Joshua Ripple, Ravi Chacko, Elisabeth Murray, National Institute of Mental Health 181

III-46. [Spatial representation in the ventral hippocampus](#)
 Alexander Keinath, Joshua Dudman, Isabel Muzzio, University of Pennsylvania 182

III-47. [Data assimilation of individual HVC neurons using regularized variational optimization.](#)
 Mark Kostuk, Daniel Meliza, Hao Huang, Alian Nogaret, Daniel Margoliash, Henry Abarbanel, University of California, San Diego 183

III-48. [Decoding sound source location using neural population activity patterns](#)
 Mitchell Day, Bertrand Delgutte, Harvard Medical School 183

III-49. [Feedforward inhibition controls tuning in auditory cortex by restricting initial spike counts](#)
 Ashlan Reid, Tomas Hromadka, Anthony Zador, Cold Spring Harbor Laboratory 184

III-50. [Toward a mechanistic description of shape-processing in area V4](#)
 Anirvan Nandy, Tatyana O. Sharpee, John Reynolds, Jude Mitchell, Salk Institute for Biological Studies 184

III-51. [Object selectivity and tolerance to variation in object appearance trade off across rat visual corti](#)
 Sina Tafazoli, Houman Safaai, Matilde Fiorini, Gioia De Franceschi, Davide Zoccolan, School for Advanced Studies (SISSA) 185

III-52. [Invariant population representations of objects are enhanced in IT during target search](#)
 Noam Roth, Margot Wohl, Nicole Rust, University of Pennsylvania 186

III-53. [Distinct neuronal responses in the human substantia nigra during reinforcement learning](#)
 Ashwin Ramayya, Kareem Zaghloul, Bradley Lega, Christoph T Weidemann, Michael Kahana, Gordon Baltuch, University of Pennsylvania 186

III-54. [Single neurons vs. population dynamics: What is changed through learning and extinction?](#)
 Anan Moran, Donald B Katz, Brandeis University 187

III-55. [Associative learning as an emergent property of spatially extended spiking neural circuits with STDP](#)
 Pulin Gong, John Palmer, University of Sydney 187

III-56. [Self-tuning spike-timing dependent plasticity curves to simplify models and improve learning](#)
 Micah Richert, Botond Szatmary, Eugene Izhikevich, Brain Corporation 188

III-57. [Formation and regulation of dynamic patterns in two-dimensional spiking neural circuits with STDP](#)
 John Palmer, Pulin Gong, University of Sydney 188

III-58. [Spinal presynaptic inhibition promotes smooth limb trajectories during reaching](#)
 Andrew Fink, Eiman Azim, Katherine Croce, Z. Josh Huang, Larry Abbott, Thomas M. Jessell, Columbia University 189

III-59. [Integrative properties of motor cortex pyramidal cells during quiet wakefulness and movement](#)
 Paolo Puggioni, Miha Pelko, Mark van Rossum, Ian Duguid, University of Edinburgh 190

III-60. [Neural dynamics following optogenetic disruption of motor preparation](#)
 Daniel O'Shea, Werapong Goo, Paul Kalanithi, Ilka Diester, Charu Ramakrishnan, Karl Deisseroth, Krishna Shenoy, Stanford University 190

III-61. [Characterization of dynamical activity in motor cortex](#)
 Gamaleldin Elsayed, Matthew Kaufman, Stephen Ryu, Krishna Shenoy, Mark M Churchland, John Cunningham, Washington University in Saint Louis 191

III-62. [Decoding arm movements from hybrid spike-field potentials in human motor cortex](#)
 Janos Perge, Brown University 191

III-63. [A model of hierarchical motor control and learning: from action sequence to muscular control](#)
 Ekaterina Abramova, Luke Dickens, Daniel Kuhn, Aldo Faisal, Imperial College London 192

III-64. [Gamma band activity in the human parahippocampal gyrus predicts performance in a sequence task](#)
 Radhika Madhavan, Hanlin Tang, Daniel Millman, Nathan Crone, Joseph Madsen, William Stan Anderson, Gabriel Kreiman, Harvard Medical School 193

III-65. [Temporal windowing of stimulus processing in V1 by saccade-driven alpha oscillations](#)
 James McFarland, Adrian Bondy, Bruce Cumming, Daniel A Butts, University of Maryland 193

III-66. [Theta and gamma activity during human episodic memory retrieval.](#)
 John Burke, Michael Kahana, University of Pennsylvania 194

III-67. [A recurrent neural network that produces EMG from rhythmic dynamics](#)
 David Sussillo, Mark M Churchland, Matthew Kaufman, Krishna Shenoy, Stanford University 194

III-68. [Natural firing patterns reduce sensitivity of synaptic plasticity to spike-timing](#)
 Michael Graupner, Srdjan Ostojic, New York University 195

III-69. [Local edge statistics provides significant information regarding occlusions in natural scenes](#)
 Kedarnath Vilankar, James Golden, Damon Chandler, David Field, Cornell University 195

III-70. [Singular dimensions in spike triggered ensembles of correlated stimuli](#)
 Johnatan Aljadeff, Ronen Segev, Michael J. Berry II, Tatyana O. Sharpee, University of California, San Diego 196

III-71. [The adaptive trade-off between discriminability and detectability in the vibrissa system](#)
 Douglas Ollerenshaw, He Zheng, Qi Wang, Garrett Stanley, Georgia Institute of Technology 197

III-72. [How humans and rats accumulate information in a tactile task, and a putative neuronal code](#)
 Arash Fassihi, Athena Akrami, Vahid Esmaeili, Mathew E Diamond, School for Advanced Studies (SISSA) 197

III-73. [Interplay of confidence and value in the transformation of olfaction-to-action](#)
 Gil Costa, Zachary F. Mainen, Champalimaud Neuroscience Programme 198

III-74. [Transformation of a temporal sequence to a spatial pattern of activity in piriform cortex.](#)
 Honi Sanders, Brian Kolterman, Dmitry Rinberg, Alexei Koulakov, John Lisman, Brandeis University . . . 198

III-75. [Seeing turbulence - real-time visualization of turbulent odor plumes](#)
 Venkatesh Gopal, Robert Morton, Alexander Grabenhofer, Elmhurst College 199

III-76. [Representation of natural images in V4 using feature invariances](#)
 Yuval Benjamini, Julien Mairal, Ben Willmore, Michael Oliver, Jack Gallant, Bin Yu, University of California, Berkeley 199

III-77. [Tuning to low-level visual features is conserved during mental imagery](#)
 Thomas Naselaris, Cheryl Olman, Dustin Stansbury, Jack Gallant, Kamil Ugurbil, Medical University of South Carolina 200

III-78. [Optimal de-noising and predictive coding account for spatiotemporal receptive field of LGN neurons](#)
 Ziqiang Wei, Tao Hu, Dmitri Chklovskii, Janelia Farm Research Campus, HHMI 200

III-79. [Compensation of heading tuning for rotation in area VIP: Retinal and extra-retinal contributions](#)
 Adhira Sunkara, Greg DeAngelis, Dora Angelaki, Washington University in St. Louis 201

III-80. [Emergence of bottom-up saliency in a spiking model of V1](#)
 Botond Szatmary, Micah Richert, Jayram Moorkanikara Nageswaran, Csaba Petre, Filip Piekniewski, Sach Sokol, Eugene Izhikevich, Brain Corporation 202

III-81. Gating of retinal information transmission by saccadic eye movements Pablo Jadzinsky, Stephen Baccus, Stanford University	202
III-82. The role of the pulvinar in visual processing and attention Ethan Meyers, Robert Schafer, Ying Zhang, Tomaso Poggio, Robert Desimone, Massachusetts Institute of Technology	203
III-83. Motion estimation involves high-order correlations differentially tuned for light and dark edges Justin Ales, Damon Clark, James Fitzgerald, Daryl Gohl, Marion Silies, Anthony Norcia, Thomas Clan- dinin, Stanford University	203
III-84. S cone isolating stimuli evoke express saccades in a chromatic contrast dependent manner Nathan Hall, Carol Colby, University of Pittsburgh	204
III-85. Motion-induced gain modulation in V1 improves contour detection Torsten Lüdge, Robert Urbanczik, Walter Senn, Bern	204
III-86. Stimulus-dependence of membrane potential and spike count variability in V1 of behaving mice Gergo Orban, Pierre-Olivier Polack, Peyman Golshani, Mate Lengyel, University of Cambridge	205
III-87. Fly photoreceptors are tuned to detect and enhance higher-order phase correlations Daniel Coca, Uwe Friederich, Roger Hardie, Stephen Billings, Mikko Juusola, University of Sheffield	206
III-88. A computational model of the role of eye-movements in object disambiguation Lena Sherbakov, Gennady Livitz, Aisha Sohail, Anatoli Gorchetchnikov, Ennio Mingolla, Massimiliano Versace, Boston University	206
III-89. How biased are maximum entropy models of neural population activity? Jakob H Macke, Iain Murray, Peter Latham, Max Planck Institute Tübingen	207
III-90. Unsupervised learning of latent spiking representations Ryan P Adams, Geoffrey Hinton, Richard Zemel, Harvard University	207
III-91. Internally generated transitions in multi-state recurrent neural networks Sean Escola, Larry Abbott, Columbia University	208
III-92. Scalability properties of multimodular networks with dynamic gating Daniel Marti, Omri Barak, Mattia Rigotti, Stefano Fusi, Columbia University	208
III-93. A mathematical theory of semantic development Andrew Saxe, James L. McClelland, Surya Ganguli, Stanford University	209
III-94. Disentangling serial adaptation using a biophysically constrained model Bongsoo Suh, Stephen Baccus, Stanford University	209
III-95. Efficient associative memory storage in cortical circuits of inhibitory and excitatory neurons Julio Chapeton, Armen Stepanyants, Northeastern University	210
III-96. Probabilistic inference reveals synapse-specific short-term plasticity in neocortical microcircuits Rui P. Costa, P. Jesper Sjöström, Mark van Rossum, University of Edinburgh	210
III-97. A spiking input-output relation for general biophysical neuron models explains observed 1/f re- sponse Daniel Soudry, Ron Meir, Electrical Engineering, Technion	211
III-98. Storing structured sparse memories in a large-scale multi-modular cortical network model Alexis Dubreuil, Nicolas Brunel, UMR 8118, CNRS & Universite Paris Descartes	211
III-99. Conditional random fields for spiking populations of neurons Emily Denton, Richard Zemel, University of Toronto	212
III-100. The influence of network structure on neuronal dynamics Patrick Campbell, Michael Buice, Duane Nykamp, University of Minnesota	212

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

This PDF can be downloaded at: http://cosyne.org/c/index.php?title=Cosyne2013_Program

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. Are we asking the right questions?

William Bialek

Princeton University

What features of our recent sensory experience trigger the responses of neurons? In one form or another, this question has dominated our exploration of sensory information processing in the brain, from the first description of center-surround organization and lateral inhibition in the retina to face- and object-selective cells in IT. As it becomes possible to record from more neurons simultaneously, we ask analogs of this question in larger populations. Much has been learned along this path, but ... :

Although we (as external observers) can correlate neural responses with sensory inputs, the brain cannot do this. Although neurons respond to the past (causality!), the only information of value to the organism is information that helps make predictions about the future.

Although we talk about successive layers of sensory processing as if they were operating on the inputs (extracting edges, connected contours, etc.), once we leave the primary sensory neurons, all neural computation operates on patterns of spikes.

Although we can generalize our questions from single neurons to populations, many models predict that sufficiently large groups of neurons should exhibit collective behaviors that are not simply extrapolated from what we see in single cells or even in small groups.

Although many of us are especially interested in the function of large brains (like our own, and those of our primate cousins), in the interests of quantitative analysis we often force behavior into small boxes, and we analyze experiments with a very limited mathematical apparatus.

In this talk I will try to turn these somewhat vague concerns into more precise questions, and suggest some paths to interesting answers. I should warn you that I am more sure about the questions than about the answers.

T-2. Neurogrid: A hybrid analog-digital platform for simulating large-scale neural models

Kwabena Boahen

Stanford University

Large-scale neural models link high-level cognitive phenomena to low-level neuronal and synaptic mechanisms, enabling neuroscientists to hypothesize how cognition emerges from the brain's biophysics. These models have been simulated using a digital approach ever since Hodgkin and Huxley pioneered ion-channel modeling in the 1940s. Computer performance has increased over a billionfold since then, enabling a 60-processor Beowulf

cluster to simulate a model with a million neurons connected by half a billion synapses. This cluster takes one minute to simulate one second of the model's behavior - sixty times slower than real-time - and consumes nine kilowatts. Moreover, current projections are that simulating human-neocortex-scale models (10^{10} neurons and 10^{14} synapses) in real-time will require an exascale computer (10^{18} flops) that consumes close to a gigawatt. In this talk, I will describe a hybrid analog-digital approach that makes it possible to now simulate a million neurons and billions of synaptic connections in real-time while consuming a few watts - a five order-of-magnitude improvement in energy-efficiency (power x time). I will demonstrate the success of this method, realized in Neurogrid, by presenting simulations of various cortical cell types, active dendritic behaviors, and spatially selective top-down attentional modulation of one cortical area by another. By providing an energy-efficient method to simulate large-scale neural models, this work takes a significant step toward making the computational power required to link cognition to biophysics widely available to neuroscientists.

T-3. Neural substrates of decision-making in the rat

Carlos Brody
Princeton University

Gradual accumulation of evidence is thought to be a fundamental component of decision-making. Over the last 16 years, research in non-human primates has revealed neural correlates of evidence accumulation in parietal and frontal cortices, and other brain areas. However, the mechanisms underlying these neural correlates remain unknown. Reasoning that a rodent model of evidence accumulation would allow a greater number of experimental subjects, and therefore experiments, as well as facilitate the use of molecular tools, we developed a rat accumulation of evidence task, the "Poisson Clicks" task. In this task, sensory evidence is delivered in pulses whose precisely-controlled timing varies widely within and across trials. The resulting data are analyzed with models of evidence accumulation that use the richly detailed information of each trial's pulse timing to distinguish between different decision mechanisms. The method provides great statistical power, allowing us to: (1) provide compelling evidence that rats are indeed capable of gradually accumulating evidence for decision-making; (2) accurately estimate multiple parameters of the decision-making process from behavioral data; and (3) measure, for the first time, the diffusion constant of the evidence accumulator, which we show to be optimal (i.e., equal to zero). In addition, the method provides a trial-by-trial, moment-by-moment estimate of the value of the accumulator, which can then be compared in awake behaving electrophysiology experiments to trial-by-trial, moment-by-moment neural firing rate measures. Based on such a comparison, we describe data and a novel analysis approach that reveals differences between parietal and frontal cortices in the neural encoding of accumulating evidence. Finally, using semi-automated training methods to produce tens of rats trained in the Poisson Clicks accumulation of evidence task, we have also used pharmacological inactivation to ask, for the first time, whether parietal and frontal cortices are required for accumulation of evidence, and we are using optogenetic methods to rapidly and transiently inactivate brain regions so as to establish precisely when, during each decision-making trial, it is that each brain region's activity is necessary for performance of the task.

T-4. Evidence of a new (exponentially strong) class of population codes in the brain

Ila Fiete
University of Texas at Austin

Neural representation is inherently noisy. Representational accuracy is increased by encoding variables in the activities of a large population of neurons. Most known population codes for continuous variables can at best reduce squared error by a factor of N , where N is the number of neurons involved in the representation. I'll argue that we should consider the existence in the brain of a qualitatively different class of population codes,

whose accuracy increases exponentially, rather than polynomially, with N . To make my argument, I'll present data and a model of persistence and capacity in human short-term memory, showing that human performance is consistent with the possibility that the brain exploits exponentially strong population codes (EPC's). We will also look under the hood at a specific cortical code, the grid cell representation of location, and see that this code enables exponential gains in accuracy with N . These examples suggest that we should begin to search for, and expect to find, codes in the brain that have strong error-correcting capabilities, in the sense of Shannon.

T-5. The impact of degeneracy on system robustness

Eve Marder

Brandeis University

Biological neurons have many voltage and time-dependent currents, many of which may have overlapping functions. Neuronal circuits often have parallel pathways that connect a pair of neurons via several routes. I will describe recent computational and experimental data that argue that degenerate cellular and circuit mechanisms contribute to robust performance despite the considerable variability seen in biological systems.

T-6. Functional models and functional organisms for systems neuroscience

Anthony Movshon

New York University

Systems neuroscience seeks to explain human perception, cognition and action. Animal models are an essential component of this explanation, because it is difficult to explore the underlying biology in humans. The effective use of animal models hinges on two tactical choices: what level of explanation to seek, and what model system to use. I will argue that the appropriate level of explanation for many problems is the functional or algorithmic level, because it is not yet possible or necessary to seek more detail. I will also argue that the appropriate animal model for this kind of functional approach is the nonhuman primate, because of its behavioral sophistication and its closeness to man in brain structure and function. While lower organisms now offer technical advantages, more finely resolved molecular and optical approaches to studying the brain will become commonplace in nonhuman primates the future. Systems neuroscientists will then be much less pressed to choose animals for study on the basis of the available toolkit for each species, and will instead be able to use models that better approximate human brain function and behavior.

T-7. Collective regulation in ant colonies

Deborah Gordon

Stanford University

Ant colonies operate without central control, using networks of simple interactions to regulate foraging activity and adjust to current ecological conditions. In harvester ants, the probability that an outgoing forager leaves the nest depends on its rate of interaction, in the form of brief antennal contact, with returning successful foragers. This links foraging behavior to food availability because the rate of forager return, which can be modelled as a Poisson process, depends on search time; the more food is available, the shorter the search time and the higher the rate of forager return. The accumulation of olfactory signals made by outgoing foragers deciding whether to leave the nest to forage suggests a stochastic accumulator model. Pools of available foragers are refilled from

reserves deeper in the nest, in a manner analogous to synaptic vesicle trafficking. The 11,000 species of ants show staggering ecological diversity. My recent work on two other ant species suggests that decision-making algorithms may be equally diverse, leading to interesting insights from the similarities and differences in neural and ant colony behavior.

T-8. Learning worth in an uncertain world: Probabilistic models of value

Paul Schrater

University of Minnesota

While it is fair to say we choose what we value, the relative ease with which we make choices and actions masks deep uncertainties and paradoxes in our representation of value. For example, ambiguous and uncertain options are typically devalued when pitted against sure things - however, curiosity makes uncertainty valuable. In general, ecological decisions can involve goal uncertainty, uncertainty about the value of goals, and time/state-dependent values. When a soccer player moves the ball down the field, looking for an open teammate or a chance to score a goal, the value of action plans like passing, continuing or shooting depends on conditions like teammate quality, remaining metabolic energy, defender status and proximity to goal all of which need to be integrated in real time. In this talk, we explicate three challenging aspects of human valuation using hierarchical probabilistic value representations. Hierarchical probabilistic value representations provide a principled framework for complex, contextual value learning and for the conversion of different kinds of value by representing more abstract goals across a hierarchy. Curiosity can result from value learning in a hierarchical Bayesian Reinforcement learning, controlled by high level uncertainty about the number and location of rewarded states. Preference reversals are generated from rational value learning with hierarchical context, including anchoring and similarity effects. Finally we show how probabilistic representations of value can solve the problem of converting and integrating heterogeneous values, like metabolic costs vs. scoring a soccer goal. By modeling values in terms of probabilities of achieving better outcomes, we decompose complex problems like the soccer player into weighted mixture of control policies, each of which produces a sequence of actions associated with more specific goal. Critically, the weights are inferences that integrate all the time-varying probabilistic information about the relative quality of each policy. We use the approach to give a rational account for a set of reaching and oculomotor experiments with multiple goals.

T-9. Peripheral and central contributions to auditory attention

Barbara Shinn-Cunningham

Boston University

Imagine yourself at a rollicking party. One of the wonderful feats of auditory perception is that if you find yourself in dull conversation about, say, motor homes, you can satisfy social expectations while really listening to the racy gossip about your colleague coming from your left. This real-life example illustrates the importance of auditory attention in allowing us to communicate in noisy settings. Early psychoacoustic work focused on understanding how you process a message when there is only one source demanding your attention, such as in a phone conversation or when listening to speech in a background of steady-state noise. However, that is not how you operate when you navigate a busy street or argue some controversial issue at a faculty meeting. The perceptual ability to focus auditory attention, which varies greatly even amongst 'normal-hearing' listeners, depends not only on cortical circuitry that directs volitional attention, but also on the fidelity of the subcortical sensory coding of sound. Results from behavioral tests illustrate how different spectro-temporal sound features influence the ability to segregate and select a target sound from an acoustic mixture, while neuroimaging results reveal the power of top-down attention in modulating the neural representation of the acoustic scene.

T-10. Suspicious coincidences in the brain: beyond the blue brain.

Terrence Sejnowski

Salk Institute for Biological Studies

Brains need to make quick sense of massive amounts of ambiguous information with minimal energy costs and have evolved an intriguing mixture of analog and digital mechanisms to allow this efficiency. Spike coincidences occur when neurons fire together at nearly the same time. In the visual system, rare spike coincidences can be used efficiently to represent important visual events in the early stages of visual processing. This can be implemented with analog VLSI technology, creating a new class of cameras.

T-11. Hidden complexity of synaptic receptive fields in cat V1

Yves Fregnac

CNRS-UNIC

Two types of functional receptive fields (RF), in the mammalian visual cortex (V1), are classically distinguished: (i) Simple cells whose spiking response is selective to the position and contrast polarity of oriented stimuli and (ii) Complex cells which respond selectively to specific oriented contours but regardless of their polarity and precise location (Hubel and Wiesel, 1962). In terms of spike-based computation and RF model architecture, Simple RFs are described as a single linear spatiotemporal filter followed by a half-rectifying nonlinearity ('LN model' in Mancini et al, 1990; Heeger, 1992). In contrast, Complex RFs are modeled as two parallel LN branches, whose linear spatiotemporal filters are at 90° spatial phase one of another, and whose output nonlinearities are fully rectifying ('energy model' in Adelson and Bergen, 1985). This classical dichotomy thereby postulates that Simple RFs come down to their linear RF component and that higher-order nonlinearities, if any, are primarily determined by the linear RF selectivity. Complex RFs are fully described by their second-order nonlinearities, i.e. they respond specifically to pair-wise combinations of stimulus positions. Recent research efforts in my lab performed in collaboration with Julien Fournier (primary contributor), Cyril Monier, Manuel Levy, Olivier Marre and Marc Pananceau have addressed complementary aspects of the same question: Does the functional separation between Simple and Complex receptive field types, expressed at the spiking output level, correspond to structural differences, or, on the contrary, to context-dependent weighting of the same substratum of synaptic sources? A related problem is to characterize synaptically the relative impact of intracortical recurrent and lateral processing vs. feedforward drive in visual cortical dynamics. For this purpose, we developed a series of experimental and computational studies based on recordings of intracellular membrane potential dynamics and Volterra decomposition of synaptic responses. We focused on the three following issues: (1) Does the functional expression of the Simple or Complex nature of V1 receptive fields depend, in the same cell, on the spatiotemporal statistics of the visual stimulus? (2) What is the topology of the presynaptic subunits giving rise to the linear and nonlinear components of the sub-threshold (synaptic) receptive field? (3) To what extent the functional diversity expressed in the synaptic inputs of a single cell can be used as a substrate for stimulus-dependent adaptation and associative plasticity?

T-12. Thalamic drive of deep cortical layers

Christine Constantinople

Randy Bruno

Columbia University

CMC2229@COLUMBIA.EDU

RANDYBRUNO@COLUMBIA.EDU

The thalamocortical projection to layer 4 (L4) of primary sensory cortex is thought to be the main route by which information from sensory organs reaches the neocortex. Sensory information is believed to then propagate through the cortical column along the L4; L2/3; L5/6 pathway. However, the same TC axons that arborize so extensively in

L4 also form smaller secondary arbors at the L5-L6 border. Therefore, deep layer neurons are poised to integrate sensory information from at least two classes of inputs: the direct thalamocortical pathway and the indirect L4; L2/3; L5/6 pathway. We sought to determine the relative functional contributions of these two pathways to the sensory responses of neurons in L5/6 in the rat barrel cortex. A substantial proportion of L5/6 neurons exhibit sensory-evoked post-synaptic potentials and spikes with the same latencies as L4. Paired in vivo recordings from L5/6 neurons and thalamic neurons revealed significant convergence of direct thalamocortical synapses onto diverse types of infragranular neurons. Pharmacological inactivation of L4 had no effect on sensory-evoked synaptic input to L5/6 neurons, and responsive L5/6 neurons continued to discharge spikes. In contrast, inactivation of thalamus suppressed sensory-evoked responses. We conclude that L4 is not a distribution hub for cortical activity, contrary to longstanding belief, and that thalamus activates two separate, independent 'strata' of cortex in parallel. Our results show that the effectiveness of a projection in activating a target region cannot be inferred from the strengths or relative numbers of individual synapses. Finally, our data suggest that the upper and lower layers have different functional roles, and are not separate steps in a single serial processing chain (L4, then L2/3, then L5/6).

T-13. Evidence for decision by sampling in reinforcement learning

Mel Win Khaw
Nathaniel Daw
New York University

MWK247@NYU.EDU
NDD204@NYU.EDU

Predominant accounts of learning in trial-and-error decision tasks, such as prediction error theories of the dopamine system, envision that the brain maintains a net value for each option and compares them at choice time. However, there has recently been considerable attention to mechanisms that construct decision variables anew at choice time, drawing e.g. on episodic memories of previous experiences with an option. Theories that construct decision variables by sampling outcomes sparsely from previous episodes can account for aggregate characteristics of preferences, such as loss aversion and risk sensitivity (Stewart, Chater & Brown, 2005; Erev, Ert & Yechiam, 2008), but have not been quantitatively compared to choice adjustments trial-by-trial, an area where prediction error models have enjoyed success. We analyzed choice timeseries from humans (n=20) performing a four-armed bandit decision task, comparing the fit of a traditional prediction error learning model to a sampling model. The models were matched in all respects save the construction of the decision variables: whereas error-driven learning averages previous outcomes with weights declining exponentially in their lags, our model samples previous outcomes with a recency bias of the same exponential form. Thus (given the same learning rate parameter controlling the sharpness of the exponential decay) the models predict the same decision variables on average, but distinct patterns of variability around them. Fitting free parameters by maximum likelihood, marginalizing the samples drawn, and using BIC to estimate model evidence, sampling was favored for 17/20 participants, with an average log Bayes factor of 4.71 (112 times more likely, given the data, than the prediction error model). The best-fitting number of samples per option was, for most subjects, just one. These results challenge deeply held neurocomputational accounts of learning in this class of tasks, and suggest the involvement of additional neural systems, notably those associated with episodic memory.

T-14. Evidence for a causal inverse model in an avian song learning circuit

Richard Hahnloser¹
Nicolas Giret²
Jörgen Kornfeld²
Surya Ganguli³

RICH@INI.PHYS.ETHZ.CH
NIGIRET@INI.PHYS.ETHZ.CH
JOERGENK@INI.PHYS.ETHZ.CH
SGANGULI@STANFORD.EDU

¹Institute of Neuroinformatics UZH / ETHZ

²University of Zurich / ETH Zurich

³Stanford University

Imitation learning, as in speech, requires that sensory targets must be able to instruct the brain's motor codes. The computational principles underlying this instruction remain largely unknown. In a vocal learner we explore the closed sensory-motor loop or inverse-model hypothesis according to which auditory signals feed into vocal motor areas by inverting the causal mappings from motor commands to sounds to auditory feedback. Causal inverse models are appealing because they constitute the simplest known neural mechanism capable of explaining motor feats such as single-trial imitation. Causal inverse models predict that sensory inputs to motor areas lag motor responses with a temporal offset given by the total loop delay, i.e., the sum of auditory and motor response latencies. We test for existence of such models in adult male zebra finches by chronically recording from the cortical output area of a basal-ganglia pathway. At many single and multi-unit sites, sensory responses tend to mirror motor-related activity with a temporal offset of about 40 ms, in accordance with minimal loop delays estimated using electrical and auditory stimulation. We show that vocal-auditory mirroring arises from a simple eligibility-weighted Hebbian learning rule that constitutes a generative mechanism for inverse models and that can explain several puzzling aspects of auditory sensitivity in motor areas, including selectivity for the bird's own song, lack of sensitivity to distortions of auditory feedback, and dependence of mirroring offsets on firing variability. Namely, variable motor sequences as in the cortical area we study (the lateral magnocellular nucleus of the anterior nidopallium, LMAN) give rise to large mirroring offsets and to causal inverse models (which map sensation to the same action), whereas stereotyped motor sequences as found elsewhere (HVC) give rise to zero mirroring offsets and to less powerful predictive inverse models (that map sensation to future action).

T-15. Inference rather than selection noise explains behavioral variability in perceptual decision-making

Jan Drugowitsch^{1,2}
Valentin Wyart³
Etienne Koechlin⁴

JDRUGO@GMAIL.COM
VALENTIN.WYART@ENS.FR
ETIENNE.KOECHLIN@UMPC.FR

¹INSERM

²École Normale Supérieure

³INSERM, ENS

⁴INSERM, ENS, UPMC Paris

During perceptual decision-making, human behavior often varies beyond what can be explained by variability in the underlying sensory evidence. It has been hypothesized that this additional variability stems from noise in the inference process or from noise in the selection process, but no clear evidence has been presented in favor of either alternative. In particular, it has recently been proposed that humans decide by drawing samples from their posterior belief, which is akin to introducing selection noise at the last stage of the decision process. In order to pinpoint the origin of behavioral variability, we designed a multi-sample categorization task in which human subjects had to decide which of two alternatives was the generative orientation of a sequence of 2 to 16 high-contrast Gabor patterns. Critically, varying the length of the sequence between trials allowed us to distinguish between a single source of noise at the selection stage and one that affects the incremental inference of the posterior belief based on the pattern sequence. We designed an ideal observer model of the task that allowed us to characterize in detail how these different sources of noise were expected to affect behavior. Fitting human behavior to the model revealed both qualitatively and quantitatively that neither noise in the selection process nor in the prior expectation of either alternative being correct could explain the data. Instead, behavioral variability was best explained by noise in the inference process, followed by choosing the alternative perceived as being most likely correct. This pattern of results was confirmed in two additional experimental conditions: one in which human subjects had three alternatives to choose from instead of two, and one in which the prior probability of either alternative being correct changed gradually from trial to trial.

T-16. Lateral interactions tune the early stages of visual processing in drosophila

Limor Freifeld¹
 Damon Clark²
 Helen Yang¹
 Mark Horowitz¹
 Thomas Clandinin¹

LIMORB@STANFORD.EDU
 DAMON.CLARK@YALE.EDU
 HELEN.H.YANG@STANFORD.EDU
 HOROWITZ@STANFORD.EDU
 TRC@STANFORD.EDU

¹Stanford University
²Yale University

Early stages of visual processing must transform complex and dynamic inputs to guide behavior. While peripheral neurons often exploit natural scene statistics to minimize information loss via efficient encoding, downstream neurons extract specific features at the expense of other information. Where does such specialization first arise? How similar are these strategies across taxa? Using two-photon Ca² imaging in *Drosophila*, combined with high-throughput methods for spatiotemporal receptive field (RF) characterization, we investigate a first order interneuron, L2, that provides input to a pathway specialized for detecting moving dark edges. We find that L2 cells have an antagonistic center-surround RF and differentially respond to large and small dark objects. These responses are successfully described via a center-surround model as a sum of two inputs associated with different time-constants. This spatiotemporal coupling suggests an efficient strategy for encoding motion cues associated with dark objects. Furthermore, this coupling may speed tune L2, thus affecting the response properties of downstream motion detectors and their sensitivities to different types of moving objects. We conclude that the functional properties of L2 represent an early step in the specialization of downstream visual circuitry. Next, using genetic and pharmacological manipulations, we identify neural mechanisms that shape L2 responses. We find that GABAergic lateral interactions, mediated at least in part pre-synaptically via receptors on photoreceptors, provide surround antagonism. GABAergic circuits also affect response kinetics and are required for L2 cells to strongly respond to decrements. We also find that cholinergic interactions establish the extent of the spatial RF, including both its center and surround components. Remarkably, this detailed characterization reveals striking similarities between the functional properties of L2 and first order interneurons in the vertebrate retina, demonstrating that early visual processing circuits across taxa employ a similar set of solutions to transform complex visual information

T-17. Feed-forward inhibition in hippocampal microcircuits: adaptation to spike-based computation

Balazs B Ujfalussy
 Mate Lengyel
 University of Cambridge

BBU20@CAM.AC.UK
 M.LENGYEL@ENG.CAM.AC.UK

Inhibitory micro-circuits are essential ingredients of neuronal networks but their computational roles are largely unknown. One experimentally particularly well-explored, but computationally little understood system is the hippocampal CA1 region. Specifically, the modulatory effects of perforant path (PP) inputs on the responses to Schaffer collateral (SC) inputs show complex changes, from dominant excitation to overall suppression mediated by feed-forward inhibition, depending on their relative timing. Here we show that these paradoxical effects can be understood as a result of the adaptation of the hippocampal circuit to computing analogue quantities using spike-based communication. We develop a theory in which postsynaptic CA1 pyramidal cells perform computations based on the graded activity of their presynaptic counterparts in the face of uncertainty because the analogue activities of presynaptic cells are only reflected in their discrete and stochastic firing. Since each spike alone conveys only limited information about the underlying activity of the cell that fired it, the postsynaptic neuron has to combine information from many presynaptic neurons. Importantly, the optimal way to combine information from many presynaptic sources depends on the joint statistics of those presynaptic neurons. We apply this theory to the case when the joint statistics of the presynaptic activities show patterns characteristic of the PP and SC inputs

of CA1 pyramidal cells in vivo. Namely, both PP and SC inputs are assumed to be strongly modulated by the theta oscillation, with a 110 degree phase delay between the two areas. We then predict the experimentally observed timing-dependent modulation of SC inputs by PP stimulation to be optimal for computing for just such presynaptic statistics. Thus, our theory allows us to understand a puzzling aspect of hippocampal circuit dynamics as an adaptation to spike-based computations under in vivo-like conditions.

T-18. Conditioned interval timing in V1 by optogenetically hijacking basal forebrain inputs

Cheng-Hang Liu

CHENGLIU08@GMAIL.COM

Johns Hopkins University

Visually-evoked neural activity in rodent primary visual cortex (V1) following behavioral conditioning comes to report the time of expected reward after visual cue presentation; so called reward-timing. Three general forms of reward-timing activity have been characterized in V1: neurons with 1) sustained increased, 2) sustained decreased activity until the time of expected reward, or, 3) neurons with peak (phasic) activity at the expected reward time. We hypothesized that a neuromodulatory system conveys the receipt of reward, providing a reinforcement signal necessary for V1 cortical plasticity to establish reward-timing activity. As neurons in basal forebrain (BF) nuclei are activated by the acquisition of reward and provide prominent neuromodulatory inputs to V1, we explored their role in imparting the putative reinforcement signal. We found that spatiotemporally restricted optogenetic activation of BF projections within V1 at fixed delays following visual stimulation was sufficient to recapitulate the effects of behavioral conditioning. Evidenced by single-unit recordings, all three general forms of 'reward-timing' activity were observed to accord with the expected timing of BF activation. Accurate report at the population level of the conditioned intervals (1 or 2 second delay) is observed in respective experimental groups. Subsequently changing the conditioning interval demonstrates that the 'reward-timing' activity in V1 is bidirectionally modifiable, yielding an appropriate update of the neural report of expected delay. Our data also suggest that the precision of optogenetically entrained time can be fine-tuned by prior experience, and its operation conforms to the widely observed scalar timing property. These findings demonstrate that optogenetically hijacking BF innervation of V1 is sufficient to engender conditioned interval timing.

T-19. Triple-spike-dependent synaptic plasticity in active dendrites implements error-backpropagation

Mathieu Schiess

Robert Urbanczik

Walter Senn

University of Bern

SCHIESS@PYL.UNIBE.CH

URBANCZIK@PYL.UNIBE.CH

SENN@PYL.UNIBE.CH

Neurons with active dendrites have been suggested to functionally represent a 2-layer network of point neurons. We point out that this interpretation offers a biological implementation of error-backpropagation. This learning rule for multilayer neural networks can be implemented in a single neuron based on dendritic and somatic spikes which propagate to the synaptic location. We particularly show that learning in neurons involving dendritic structures fails if synaptic plasticity would not take account of dendritic spikes. Only if the presynaptic, dendritic and somatic spikes jointly modulate plasticity can the representational power of dendritic processing be exploited. We propose a learning rule which is suitable to reinforcement as well as supervised learning scenarios. The learning rule shows how the timing among these spikes and the postsynaptic voltage optimally determines synaptic plasticity. We show that the rule maximizes the expected reward in the context reinforcement learning and it optimizes a lower-bound of the log-likelihood in the context of supervised learning. The theory offers a unifying framework to encompass spike-timing-dependent plasticity and dendritic processes under the perspective of learning.

T-20. Multiple perceptible signals from a single olfactory glomerulus

Matthew Smear^{1,2}
Admir Resulaj³
Jingji Zhang⁴
Thomas Bozza⁴
Dmitry Rinberg⁵

SMEARM@JANELIA.HHMI.ORG
ADMIR.RESULAJ@NYUMC.ORG
JINGJI-ZHANG@NORTHWESTERN.EDU
BOZZA@NORTHWESTERN.EDU
DMITRY.RINBERG@NYUMC.ORG

¹Janelia Farm Research Campus
²Howard Hughes Medical Institute
³New York University
⁴Northwestern University
⁵NYU Neuroscience Institute

Glomeruli are the input channels of the olfactory system, where olfactory sensory neurons (OSNs) connect to the brain. The mouse olfactory bulb contains roughly 2,000 glomeruli, each receiving input from OSNs that express a specific odorant receptor gene. However, odors typically activate many glomeruli in complex combinatorial patterns. This complicates efforts to define the contribution of individual glomeruli to olfactory function. To study, for the first time, the signaling capacity of a single glomerulus, we used gene-targeted mice that express ChR2 from a defined odorant receptor gene, M72. We find that mice detect photostimulation of one glomerulus with near-perfect performance. Furthermore, activation of a single glomerulus can also be detected on an intense odor background. When the odor is a known M72 ligand, the odor masks light detection. In contrast, when the odor is not an M72 ligand, mice easily discriminate odor from light paired with odor. Mice can thus detect the smallest possible change in glomerular input patterns, indicating that much of the combinatorial capacity of the glomerular array is available to perception. In addition, mice can discriminate different intensities of light, and the timing of light input through one glomerulus. This demonstrates that identical patterns of glomerular input can be discriminated on the basis of non-spatial signaling parameters. The operation of identity, intensity, and temporal coding within single olfactory input channels may enable mice to efficiently smell natural olfactory stimuli, such as faint scents in the presence of intense odor backgrounds.

T-21. On the role of neural correlations in decision-making tasks

Nestor Parga¹
Federico Carnevale¹
Victor de Lafuente²
Ranulfo Romo²

NESTOR.PARGA@UAM.ES
FEDERICO.CARNEVALE@UAM.ES
LAFUENTE@UNAM.MX
RROMO@IFC.UNAM.MX

¹Universidad Autonoma de Madrid
²Universidad Nacional Autonoma de Mexico

Simultaneous recordings of pairs of cortical neurons have shown that spike-count correlation coefficients (CCs) cover a wide range of values. According to recent experimental evidence (Renart et al., 2010; Ecker et al., 2010) cortical networks are able to decorrelate neural activity producing very low CCs. Theoretical work shows that this is possible even if synaptic efficacies are strong and neurons are densely connected. This is because correlations between the external, excitatory and inhibitory inputs cancel (Renart et al., 2010). However little is known about the origin of correlations and analysis based on recordings of cortical activity of animals performing non-trivial tasks are scarce. Here we describe the role of spike-count correlations in monkeys performing a perceptual decision-making task consisting in the detection of somatosensory stimuli (de Lafuente and Romo, 2005). The main results are: 1)The temporal profile of the spike-count CCs is modulated during the task stages in a condition-dependent way. At the end of the task, before the subject reports its decision, CCs are rather small. In agreement with theoretical predictions, this is true even for large firing rates. 2)An important source of the temporal modulations of CCs is a random internally-generated signal. Computational modeling shows that this signal is responsible of errors in the subject's decisions (Carnevale et al., 2012). 3)Choice probability (CP) can

be precisely computed from spike-count correlations. They are obtained from the difference between correlations evaluated using all type of trials and correlations computed segregating trials according to the subject's choice. This implies that even very small correlations in trials with a fixed choice are compatible with single neurons having significant CP. 4) We developed a neural population analysis to decode the subject's choice. This allowed us to find combinations of firing rates of frontal lobe neurons fully correlated with the decision.

T-22. Spiking networks learning to balance excitation and inhibition develop an optimal representation.

Ralph Bourdoukan¹
David GT Barrett^{1,2}
Christian Machens³
Sophie Deneve

RALPH.BOURDOUKAN@GMAIL.COM
DAVID.BARRETT@ENS.FR
CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
SOPHIE.DENEVE@ENS.FR

¹École Normale Supérieure

²Champalimaud Centre for the Unknown, Lisbon

³Champalimaud Foundation

Cortical activity is typically irregular, asynchronous, and Poisson-like (Shadlen, M.N. and Newsome, W.T., 1998). This variability seems to be predominantly caused by a balance of excitatory and inhibitory neural input (Haider, B. et al., 2006). However, the learning mechanisms that develop this balance and the functional purpose of this balance are poorly understood. Here we show that a Hebbian plasticity rule drives a network of integrate-and-fire neurons into the balanced regime while simultaneously developing an optimal spike-based code. The remarkable coincidence of balance and optimality in our model occurs when synaptic plasticity is proportional to the product of the postsynaptic membrane potential and presynaptic firing rate. This plasticity rule acts to minimise the magnitude of neural membrane potential fluctuations. Balance develops because, without it, membrane potential fluctuations are too large. Meanwhile, an optimal representation develops because membrane potentials correspond to representation errors for a signal encoded by the network (Boerlin, M. and Denève, S., 2011). This signal may be a sensory signal or the result of some network computation. It can be extracted from the network spike trains with a fixed linear decoder (a summation of postsynaptic potentials), with a precision on the order of $1/N$, where N is the number of neurons. This is much more precise than a typical rate model. Our work suggests that several of the features measured in cortical networks, such as the high trial-to-trial variability, the balance between excitation and inhibition, and spike-time-dependent plasticity are all signatures of an optimal spike-based representation.

T-23. Learning to infer eye movement plans from populations of intraparietal neurons

Arnulf Graf
Richard Andersen
California Institute of Technology

GRAF@VIS.CALTECH.EDU
ANDERSEN@VIS.CALTECH.EDU

Learning can induce plasticity by shaping the response properties of individual neurons throughout cortex. It is still unclear how learning-related plasticity occurs at the level of neuronal ensembles. We recorded populations of lateral intraparietal (LIP) neurons while the animal made instructed delayed memory saccades. We predicted eye movements from the responses of neuronal ensembles using Bayesian inference. To examine whether neuronal populations can undergo learning-related changes, these instructed trials were daily followed by brain-control trials. In these trials, the animal maintained fixation for the entire trial and was rewarded when the planned eye movement to the remembered target location was accurately predicted in real-time based solely on the neuronal population response and without an actual eye movement. The animal had to learn to adjust his neuronal re-

sponses to maximize the prediction accuracy without performing any overt behavior. We first examined whether learning to control the brain-machine interface resulted in changes in the neuronal activity. We found that the responses of neuronal ensembles in the instructed and brain-control trials could be accurately discriminated, thus showing a change in the population activity. Second, we investigated whether these changes were reflected in the accuracy of eye movement predictions. We showed that during the brain-control trials the prediction accuracy of eye movement planning activity increased, thus providing evidence for learning-related plasticity. This result is, to our knowledge, the first evidence for plasticity using brain-control experiments in an eye movement area. We also found that this learning effect was strongest for eye movement plans that were difficult to predict, showing that learning effects can be strongest for conditions 'where there is something to learn'. In summary, we report for the first time that populations of neurons in LIP can learn to rapidly shift their response characteristics to increase the prediction accuracy of oculomotor plans.

T-24. Neural variability and normalization drive biphasic context-dependence in decision-making

Kenway Louie
Mel Win Khaw
Paul Glimcher

New York University

KLOUIE@CNS.NYU.EDU
MWK247@NYU.EDU
GLIMCHER@CNS.NYU.EDU

Understanding the neural code is critical to linking brain and behavior. In sensory systems, the algorithm known as divisive normalization appears to be a canonical neural computation, observed in areas ranging from retina to cortex and mediating processes including contrast adaptation, surround suppression, visual attention, and multisensory integration. Recent electrophysiological studies have extended these insights beyond the sensory domain, demonstrating an analogous algorithm for the value signals that guide decision-making, but the effects of normalization on choice behavior are unknown. Here, we show that simple spike count models of decision-making incorporating normalization and stochastic variability in value coding generate significant - and classically irrational - inefficiencies in choice behavior. Notably, these models predict a novel biphasic form of value-dependent contextual modulation: the relative choice between two given options varies as a non-monotonic function of the value of other alternatives. Exploration of parameter space shows that these results depend critically on both normalization and stochastic variability. To test these predictions, we examined the behavior of 40 human subjects in an incentivized value-guided trinary choice task. We found that relative choice behavior between any two options depends on the value of a third option; consistent with model predictions, this influence is negative at low values and positive at high values. These findings suggest that the specific form of neural value representation critically influences stochastic choice behavior and that normalized value coding can provide a generalizable quantitative framework for examining context effects in decision-making.

T-25. An investigation of how prior beliefs influence decision-making under uncertainty in a 2AFC task

Daniel Acuna^{1,2}
Max Berniker¹
Hugo Fernandes¹
Konrad Kording¹

¹Northwestern University

²Rehabilitation Institute of Chicago

DANIEL.ACUNA@NORTHWESTERN.EDU
MBERNIKE@NORTHWESTERN.EDU
HUGOGUH@GMAIL.COM
KK@NORTHWESTERN.EDU

The two-alternative-forced-choice (2AFC) paradigm is the workhorse of psychophysics, used widely in fields ranging from neuroscience to economics. The data obtained from the 2AFC task is used to compute the just-

noticeable-difference (JND), which is generally assumed to quantify sensory precision, independent of a subject's beliefs (their prior) and how they represent uncertainty. However, this interpretation is only true for some specific theories of how humans make decisions under uncertainty. There are a host of alternative decision-making theories that make different predictions for how people behave in the 2AFC task and how to interpret the resulting data. Here we mathematically examine some prominent theories of how the brain represents uncertainty, as well as common experimental protocols, to determine what the JND measures. We show that if, after combining sensory and prior information, subjects choose the option that is most likely to be correct (i.e. the maximum of the posterior estimate, MAP) and the distributions are unimodal, then the JND correctly measures a subject's sensory precision. However, if, as sampling and matching theories suggest, subjects choose relative to the proportion of their posterior estimates, then the JND measures something altogether different that depends on their prior. We designed an experiment to test these varying assumptions using interleaved estimation and 2AFC tasks. We found that changes in the subjects' prior uncertainty had no influence on the measured JNDs. This finding is predicted by MAP decision-making, and argues against a range of prominent models for the representation of uncertainty. We show that, in general, the 2AFC task is not a straightforward tool for measuring subject's sensory precision. Instead, the 2AFC task can be used to falsify theories of the neural representation of uncertainty.

T-26. Visual speed information is optimally combined across different spatiotemporal frequency channels

Matjaz Jogan
Alan Stocker

University of Pennsylvania

MJOGAN@SAS.UPENN.EDU
ASTOCKER@SAS.UPENN.EDU

Humans have the ability to optimally combine sensory cues across different perceptual modalities (Ernst & Banks, 2002). Here, we tested whether optimal cue-combination also occurs within a single perceptual modality such as visual motion. Specifically, we studied how the human visual system computes the perceived speed of a translating intensity pattern that contains motion energy at multiple spatiotemporal frequency bands. We assume that this stimulus is encoded in a set of spatiotemporal frequency channels where the response of each channel represents an individual cue. We formulate a Bayesian observer model that optimally combines the likelihood functions computed for individual channel responses together with a prior for slow speeds. In order to validate this model, we performed a visual speed discrimination experiment. Stimuli were either drifting sinewave gratings with a single frequency at various contrasts, or pairwise linear combinations of those gratings in two different phase configurations that resulted in different overall pattern contrasts. The measured perceptual speed biases and discrimination thresholds show the expected Bayesian behavior where stimuli with larger thresholds were perceived to move slower (Stocker & Simoncelli, 2006). For the combined stimuli, discrimination thresholds were typically smaller compared to those measured for the independent components alone, which is a key feature of optimal cue combination. Finally, the two phase configurations of the combined stimuli did not lead to any significant difference in terms of both bias and threshold, supporting the notion of independent channels. Our observer model provided a good account of the data when jointly fit to all conditions. However, the fits were significantly better when we assumed that individual channel responses were normalized by the overall channel activity. Our findings suggest that perceived speed of more complex stimuli can be considered a result of optimal signal combination across individual spatiotemporal frequency channels.

T-27. A neural encoding model of area PL, the earliest face selective region in monkey IT

Charles Cadieu
Elias Issa
James DiCarlo

Massachusetts Institute of Technology

C.CADIEU@GMAIL.COM
ISSA@MIT.EDU
DICARLO@MIT.EDU

Progress has been made in understanding early sensory areas (e.g. olfactory bulb, retinae, V1) by constructing encoding models at the neural level. However, our understanding of human face processing has been hindered by a lack of such models. Building models of face selective neurons would lead to a mechanistic explanation of cognitive phenomena observed in the psychophysics and fMRI communities. We provide a first step towards this goal by modeling the earliest face selective region in monkey IT, area PL, which has been proposed as a gateway to face processing and may serve as a face detection module. We tested a wide array of image-based encoding models and found that hierarchical models that pool over local features captured PL responses across 3043 images at 87% cross-validated explained variance (65% mean explained variance for sites, $n=150$). Models with the highest explanatory power incorporated localized sub-features, feature rectification, and a tolerance operation over space and scale. Those models also demonstrated similar properties to PL according to a phenomenological 'scorecard' (e.g. similar rankings across face parts and non-face images). We compared these models with the 'word model' in the field – that 'face neurons' signal face presence – by measuring human judgements of 'faceness' ($n=210$ subjects). These judgements did not match the phenomenological scorecard and correlated poorly with PL responses (22% explained variance). Together these results provide new perspective on the early stages of face processing in IT: PL is better viewed as a non-linear image operation than as a cognitive indicator of face presence. In summary, this work is the first to create image-computable encoding models of face selective neurons. These models may bridge the gap between the cognitive understanding of face processing and a mechanistic understanding of the neural basis of face processing.

T-28. Neural correlates of visual orientation constancy

Ari Rosenberg
Dora Angelaki

Baylor College of Medicine

ROSENBERG@CNS.BCM.EDU
ANGELAKI@CABERNET.CNS.BCM.EDU

As we examine the world visually, movements of the eyes, head, and body change how the scene projects onto the retina. Despite this changing retinal image, perception of the environment remains highly stable — oriented along the gravitational vector termed earth vertical. Consider an upright observer fixating a vertical bar. The bar is correctly perceived to be oriented vertically in space, and its image runs along the retina's vertical meridian. If the observer's head then rolls to one side while maintaining fixation, perception of the bar remains near vertical even though its retinal image is now oblique. This visual orientation constancy reveals the influence of extra-retinal sources such as vestibular signals on visual perception, without which, the head-rolled observer could misinterpret the vertical bar as oblique. Where and how this is achieved in the brain remains unknown. Electrophysiological studies conducted in primary visual cortex have yielded conflicting results, and human studies suggest parietal cortex may be involved. Here we examine this possibility by recording extracellularly from 3D surface orientation selective neurons in the caudal intraparietal area (CIP) of macaque monkeys. Tilt tuning curves, describing how responses depend on the direction in which a plane leans towards the observer, were recorded for each cell with the animal upright and rolled ear down. Relative to the upright tilt tuning curve, about 40% of the head/body rolled tuning curves shifted significantly in the direction preserving the preferred tilt relative to earth vertical. This shift was generally larger than the ocular counter-roll but smaller than the roll amplitude (partially compensating, as often observed in multisensory integration) and sometimes accompanied by a gain change. Our findings demonstrate that the responses of CIP neurons correlate with visual orientation constancy, providing a novel look into how multisensory integration unifies and stabilizes perception of the environment.

T-29. The cortical network can sum inputs linearly to guide behavioral decisions

Mark H Histed
John Maunsell
Harvard Medical School

MARK_HISTED@HMS.HARVARD.EDU
JOHN_MAUNSELL@HMS.HARVARD.EDU

Individual neurons are sensitive to the strength and synchrony of inputs. For example, spike threshold and active dendritic channels are nonlinearities that make neurons more likely to spike when inputs arrive synchronously. These properties have led to many proposals for how temporal patterns of neural activity can encode information. On the other hand, theoretical models and studies of behaving animals show that information can be represented by the total number of spikes, with a smaller role for millisecond-scale temporal synchrony. To study this apparent contradiction, we trained mice to behaviorally report changes in the spiking of a population of excitatory neurons induced with optogenetics. We show that animals' behavior does not depend on the temporal pattern of inputs within a 100 ms time window, as if the neural population is a near-perfect linear integrator of added input. In particular, no deviation from linearity is seen when we synchronize spiking across frequencies from theta to gamma, or use pulsed inputs as short as 1 ms. To explore this insensitivity to input synchrony, we recorded the activity of neurons in response to optical stimulation. We find that near behavioral threshold, even the most responsive cells fire only few extra spikes – approximately one per trial on average. Under these conditions, where added inputs are small relative to the other ongoing inputs cortical neurons receive, these extra inputs gain little advantage from synchrony. Our results show that the cortical network can operate in a linear, behaviorally relevant regime where synchrony between inputs is not of primary importance, because each input has only a weak effect on single cells in the population. The common circuit structure of the cortex — with balanced inhibition and strong local recurrent connectivity — may be constructed to allow linear population coding.

T-30. Rank-penalized nonnegative spatiotemporal deconvolution and demixing of calcium imaging data

Eftychios A. Pnevmatikakis¹
Timothy Machado^{1,2}
Logan Groseknick³
Ben Poole³
Joshua Vogelstein⁴
Liam Paninski¹

EFTYCHIOS@STAT.COLUMBIA.EDU
TAM2138@COLUMBIA.EDU
LOGANG@STANFORD.EDU
BENMPOOLE@GMAIL.COM
JOVO@STAT.DUKE.EDU
LIAM@STAT.COLUMBIA.EDU

¹Columbia University

²Howard Hughes Medical Institute

³Stanford University

⁴Duke University

Calcium imaging is an increasingly powerful and popular technique for studying large neuronal ensembles. However, data interpretation remains challenging; the fast spiking of neurons is indirectly observed through a noisy slower calcium signal, obtained at a low imaging rate. FOOPSI and "peeling" are two algorithms for extracting spikes from imaging data using nonnegative sparse deconvolution. They both use a simple linear model in each pixel: upon each spike, the calcium signal increases by a fast stereotypical transient and then it decays slowly towards a baseline concentration. Although effective, these methods are typically applied on a pixel-by-pixel basis (or summed across the full ROI) and do not combine information optimally across pixels. Here we extend FOOPSI to derive an efficient spatiotemporal deconvolution and demixing algorithm. Our key insight is that under this linear model, the spatiotemporal calcium evolution matrix has rank equal to the (unknown) number of underlying neurons. Our problem can be cast as a rank-penalized estimation of a structured matrix and solved in a relaxed form using convex optimization. Our algorithm can be parallelized by considering nonoverlapping ROIs and scales linearly with time and quadratically with the number of pixels in each ROI. Moreover, we develop

a highly optimized GPU implementation. Our algorithm leads to dramatic denoising compared to non-spatial approaches. We can further apply a nonnegative structured matrix factorization to simultaneously deconvolve and demix the spike trains, even in the presence of spatially overlapping neurons. We introduce a method-of-moments approach to fitting the model parameters that is quicker and more robust than the previous approximate expectation-maximization methods. We also derive and compare several model selection strategies (e.g., BIC, AIC, Cp). We apply our methods to simulated and in-vitro spinal cord data, for which ground truth is available via antidromic stimulation, with promising results.

T-31. Saccadic choices without attentional selection during an urgent-decision task

Emilio Salinas
Gabriela Costello
Dantong Zhu
Terrence Stanford University
Wake Forest University

ESALINAS@WAKEHEALTH.EDU
GCOSTELL@WAKEHEALTH.EDU
DZHU@WAKEHEALTH.EDU
STANFORD@WAKEHEALTH.EDU

This work leverages a recently-developed task, heuristic model, and analytical framework to investigate how the frontal eye field (FEF) links the deployment of spatial attention to the generation of oculomotor choices. Neuron types in FEF range between visual neurons (V), which respond to stimulus presentation, to movement neurons (M), which are activated before saccade onset. Both fire differently depending on whether a target or a distracter is in the response field, but such neural discrimination is thought to represent separate functions: V neurons select a visual goal on the basis of perceptual information whereas M neurons plan specific saccades. Converging lines of evidence indicate that the V differentiation corresponds to the deployment of spatial attention to a target item; for instance, it occurs even when task rules require a spatially incompatible saccadic report or no saccade at all. We propose that, by their very design, the vast majority of prior studies have led to a generally accepted, but we think incorrect, conclusion: that the deployment of spatial attention through V-neuron activation is a prerequisite to the saccadic programming carried out by M cells. We investigate this in monkeys performing a rapid-choice task in which, crucially, motor planning always starts ahead of perceptual analysis, placing a strong temporal constraint on attentional shifts, and alignment is not imposed or encouraged between the locus of attention and the eventual saccadic goal. We find that the choice is instantiated in FEF as a competition between oculomotor plans (i.e., M activity), in agreement with model predictions. Notably, perception strongly influences this ongoing motor activity but has no measurable impact on the V cells, suggesting that rapid saccadic choices may occur without prior attentional selection of the target location. Therefore, the linkage between spatial attention and saccade planning is considerably more flexible than currently thought.

T-32. A proposed role for non-lemniscal thalamus in cortical beta rhythms: from mechanism to meaning

Stephanie Jones
Christopher Moore
Brown University

STEPHANIE_JONES@BROWN.EDU
CHRISTOPHER_MOORE@BROWN.EDU

Cortical Beta oscillations (15-29Hz) in humans are correlated with perception and attention and often altered in diseases, including Parkinson's Disease (PD). Crucial to understanding the computational importance of Beta in health and disease is to discover how it emerges. We have recently combined human brain imaging, computational neural modeling, and electrophysiological recordings in rodents to explore the functional relevance and mechanistic underpinnings cortical Beta rhythms, which arise as part of a complex of Alpha and Beta components in primary somatosensory cortex in humans. In this talk, we will review our quantification of salient characteristics

of this rhythm and its importance in information processing, including its impact on tactile detection, changes with healthy aging and practice, and modulations with attention. Constrained by these data, we developed a novel neural mechanism for the emergence of this rhythm based on a biophysically principled computational model of SI circuitry with layer specific exogenous synaptic drive. The theory predicts that the Alpha/Beta complex emerges from the combination of two stochastic ~10Hz excitatory drives to the granular/infragranular and supragranular layers. Beta requires sufficiently strong supragranular drive and a near simultaneous delay, whereas Alpha can emerge from dominance of lemniscal drive, as commonly suggested, or from activity in the distal dendrites, a more novel prediction. This model accurately reproduces numerous key features of the human MEG-measured rhythm, including its impact on sensory responses. While the model does not presume the specific regions that provide the exogenous drive, candidate sources are lemniscal and non-lemniscal thalamic nuclei, which fit the model requirements particularly well. Preliminary correlative and causal data from mice support the model predictions. Given the supragranular profile of pallidal thalamic projections to the neocortex, this theory also provides a direct and precise prediction as to why Beta is altered in PD.

T-33. The impact on mid-level vision of statistically optimal V1 surround normalization

Ruben Coen-Cagli^{1,2}
Odelia Schwartz²

RUBEN.COENCAGLI@UNIGE.CH
ODELIA.SCHWARTZ@EINSTEIN.YU.EDU

¹University of Geneva

²Albert Einstein College of Medicine

A major feature of the visual system is its hierarchical organization: cortical areas higher in the hierarchy typically have larger neuronal receptive fields (RF), and increasingly complex selectivity. The increase in RF size implies that the nonlinearities that characterize processing of stimuli beyond the RF in one area, influence the inputs received by areas downstream, and therefore their RF selectivity. This key issue has not been studied systematically. We address it here from an image statistics perspective: using a two-stage architecture, we compared how nonlinear models of area V1 differently optimized for natural images affect processing at the next stage. For the V1 nonlinearity, we considered complex cell models with (and without) different forms of surround divisive normalization derived from image statistics, including canonical normalization and a statistically optimal extension that accounted for image nonhomogeneities. Surround normalization is known to reduce dependencies between spatially-separated V1 RFs. However, V1 model complex cells exhibited residual correlations whose precise form depended on the nonlinearity. We assumed that to achieve redundancy reduction, the objective for the second stage, namely the linear V2 RFs, was to learn without supervision a representation that removed such correlations. This approach revealed V2-like feature selectivity (e.g., corners, 3-junctions, and texture boundaries) when we used the optimal normalization and, to a lesser extent, the canonical one, but not in the absence of both. We then considered the implications for perception; while both types of V1 normalization largely improved object recognition accuracy, only the statistically optimal normalization provided significant advantages in a task more closely matched to mid-level vision, namely figure/ground judgment. These results suggest that experiments probing mid-level areas might benefit from using stimuli designed to engage the computations that characterize V1 optimality. Supported by the NIH (CRCNS-EY021371) and Army Research Office (58760LS).

T-34. Formation and maintenance of multistability in clustered, balanced neuronal networks

Ashok Litwin-Kumar¹
Brent Doiron²

ALK@CMU.EDU
BDOIRON@PITT.EDU

¹Carnegie Mellon University

²University of Pittsburgh

Neuronal networks in cortex exhibit complex and variable patterns of activity even in the absence of a stimulus. However, it is unclear what network mechanisms give rise to these fluctuations. One possibility is so-called attractor dynamics, which endows a neural system with a multitude of stable states that are sampled stochastically during spontaneous activity. Recently, we proposed that clustered synaptic connections among members of neuronal assemblies give rise to attractor dynamics in recurrent excitatory/inhibitory networks. Clustered, balanced networks exhibit variable firing rates in spontaneous conditions and a reduction of variability with stimulus application, consistent with recent experimental results. How do cortical networks ensure that a rich repertoire of attractor states, rather than only a few, are sampled during spontaneous conditions? Without mechanisms to ensure this, large neuronal assemblies with many recurrent excitatory synapses may suppress other assemblies through inhibition, preventing transitions between attractor states and leading to uniform winner-take-all behavior. We investigate the influence of a recently proposed inhibitory plasticity rules that regulate inhibitory synaptic strengths, finding that robust multistability can be maintained even with heterogeneous neuronal assemblies. Specifically, we study clustered networks with many assemblies of highly heterogeneous sizes. Without fine tuning of synaptic strengths, very large assemblies exhibit high firing rates, while smaller assemblies remain suppressed. The spontaneous dynamics of these networks lack the rich, long timescale activity of clustered networks with symmetric assembly sizes. We next implement a recently characterized inhibitory rule that has been shown to regulate excitatory firing rates homeostatically, finding that it increases the number of attractor states sampled in spontaneous dynamics. We propose that cortical networks actively regulate their spontaneous dynamics through plasticity, leading to rich, variable patterns of activity.

T-35. Chromatic detection from cone photoreceptors to individual V1 neurons to behavior in rhesus monkeys

Charles Hass
Juan Angueyra
Zachary Lindbloom-Brown
Fred Rieke
Gregory Horwitz
University of Washington

CAHASS@UW.EDU
ANGUEYRA@UW.EDU
ZACKLB@UW.EDU
RIEKE@UW.EDU
GHORWITZ@U.WASHINGTON.EDU

Chromatic sensitivity cannot exceed limits set by noise in the cone photoreceptors. To determine whether cone noise is a bottleneck for cortical and psychophysical sensitivity to chromatic patterns, we developed a computational model of stimulus encoding in the cone outer segments and compared the performance of the model to the psychophysical performance of monkeys and to the sensitivities of individual neurons in the primary visual cortex (V1). The model simulated responses of a realistic mosaic of cones using a temporal impulse response function and a noise power spectral density that were derived from *in vitro* recordings of macaque cones. Behavioral data and neurophysiological recordings from V1 were obtained from three monkeys performing a chromatic detection task. We probed four isoluminant color directions and found that the sensitivity of the simulated cone mosaic, V1 neurons, and the monkeys were tightly yoked. This result suggests that the fidelity of signal transmission from the retina, through V1, to behavior is equivalent across these four color directions. Nevertheless, the absolute sensitivity of the cone mosaic was higher than that of individual V1 neurons, which in turn were slightly less sensitive than the monkey. Additional comparisons of model and behavioral thresholds revealed that the eccentricity dependence and high temporal frequency falloff of luminance flicker detection are well-predicted by the cone model, but that the behavioral insensitivity to low-frequency, achromatic modulation is not. Usefully, our model provides a baseline against which detection efficiency of such varied patterns can be compared fairly. An important direction for future research is to probe the behavioral relevance of signals that are filtered and preserved from the cones to V1 and from V1 to behavior.

T-36. The dynamics of auditory-evoked response variability and co-variability across different brain state

Gabriela Mochol^{1,2}
 Shuzo Sakata³
 Alfonso Renart⁴
 Liad Hollender⁵
 Kenneth Harris⁶
 Jaime de la Rocha¹

GMOCHOL@CLINIC.UB.ES
 SHUZO.SAKATA@STRATH.AC.UK
 ARENART73@GMAIL.COM
 LIAD_HOLLENDER@HOTMAIL.COM
 KENNETH.HARRIS@IMPERIAL.AC.UK
 JROCHAV@CLINIC.UB.ES

¹IDIBAPS

²Nencki Institute of Experimental Biology

³University of Strathclyde

⁴Champalimad Foundation

⁵Rutgers University

⁶Imperial College London

Cortical activity is ubiquitously variable and yet we lack a basic understanding of the mechanisms which cause this variability or of the computational implications it may have. Single neuron variability and pair-wise correlations decrease upon stimulus presentation. Recent theoretical studies have proposed different mechanisms which can account for the drop in variability but the dynamics of co-variability and their dependence on different brain states have not been explored. To gain an understanding of the relations between brain state, variability and co-variability we combined multi-array cortical recordings and analysis of the model network dynamics. We recorded population evoked responses in auditory cortex of urethane-anesthetized rats during periods of Inactivated (UP-DOWN transitions) and Activated (sustained firing) brain state. We presented tones and short clicks and computed time-resolved population averaged firing rates, spike count Fano factors (FF) and pair-wise correlation coefficients (r ; 20 ms window). As expected, baseline FF and r values were larger the higher the degree of Inactivation. Both statistics showed however a decrease after stimulus onset to values which were state independent, making the magnitude of the drop large for Inactivated and very small for Activated periods. We built a rate model with adaptation exhibiting bi-stability. Depending on the strength of the external input and the adaptation magnitude, the system exhibited either fluctuation-driven UP/DOWN transitions or sustained activity. The stimulus presentation quenched the variability of the population rate by transiently setting the system into the small fluctuations mono-stable regime. Conditionally-independent inhomogeneous Poisson spike trains generated from the model population fluctuating rates (double stochasticity) reproduced the experimentally observed drop in FF and r and its state dependence. Our findings show that the dynamics of evoked response variability can be related to the population rate dynamics which are strongly influenced by brain state. Funding: Polish MNISW 'Mobility Plus' Program 641/MOB/2011/0 (G.M)

T-37. Self-supervised neuronal processing of continuous sensory streams

Robert Guetig

GUETIG@EM.MPG.DE

Max Planck Institute of Experimental Medicine

During behavior, a continuous stream of sensory information reaches the central nervous system in the form of a high-dimensional spatio-temporal pattern of action potentials. When processing such activity, many sensory neurons respond with exquisite tuning and high specificity to temporally local stimulus features, such as sounds within a communication call or shapes within a movie. Often the temporal extent of such embedded features is orders of magnitude shorter than the duration of the encompassing, behaviorally meaningful sensory episode. It is commonly hypothesized that the emergence of neuronal feature detectors requires temporal segmentation of the training data, to allow neurons to adjust the strength of their responses to isolated target features. It is unclear, how such temporal supervision is implemented in neuronal circuits, in particular before sensory representations have formed. Here we show that only the number of feature occurrences without any temporal information is

sufficient to train biologically plausible model neurons to detect spatio-temporal patterns of spikes that arrive embedded in long streams of background activity. Intriguingly, neurons can even learn complex continuous tuning functions from aggregate training labels, i.e.~the sum of the desired response strengths over multiple features occurring within a given trial. Neither individual counts nor the times of features are needed. We discovered that the simplicity of such supervisory signaling allows neuronal networks to self-supervise: A single supervisor neuron that feeds back the mean population activity as training label enables neuronal populations to reliably detect reoccurring spike patterns in their input activity without any external supervision. When coupling several of such self-supervised neuronal populations through lateral inhibition we observed the formation of continuous neuronal feature maps. By successfully applying these findings to bird song and human speech recognition tasks we establish a novel network mechanism for self-supervised learning in populations of spiking sensory neurons.

T-38. Temporal basis for predicting sensory consequences of motor commands in a cerebellum-like structure

Ann Kennedy
Gregory Wayne
Patrick Kaifosh
Karina Alvina
Larry Abbott
Nathaniel B. Sawtell
Columbia University

ANNKENNEDY05@GMAIL.COM
GDW2104@COLUMBIA.EDU
PKAIFOSH@GMAIL.COM
KA2332@COLUMBIA.EDU
LFA2103@COLUMBIA.EDU
NS2635@COLUMBIA.EDU

Animals use corollary discharge (CD) to predict the sensory consequences of motor commands. In the passive electrosensory system of mormyrid fish, responses to self-generated electric organ discharges (EODs) are suppressed to improve detection of external electric fields. In the cerebellum-like electrosensory lobe, Purkinje-like medium ganglion cells (MGs) receive both electrosensory input and granule cell (GC) input containing EOD-related signals. GC inputs to MG cells generate a temporally-specific 'negative image' of the electrosensory inputs to the MG cell following an EOD. The negative image cancels unwanted electrosensory inputs time-locked to the EOD [1]. Because these sensory effects far outlast the EOD command, CD signals must be expanded in time to provide an adequate basis for this cancellation. Previous modeling studies showed that experimentally observed anti-Hebbian spike timing-dependent plasticity (aSTDP) at GC-to-MG synapses [2] can produce negative images, but this work relied on the assumption that GC inputs form a delay-line-like temporal basis [3]. Here we examine the actual temporal basis for negative images by recording GC responses, mossy fibers that provide input to GCs, and two classes of interneurons—inhibitory Golgi cells and excitatory unipolar brush cells (UBCs). We find that temporal response patterns in GCs can be explained by excitatory inputs from a few mossy fibers, with little apparent role for Golgi cell inhibition. Though activity in most GCs ends shortly after the EOD, a subset show delayed and temporally diverse responses likely due to input from UBCs. We generated a large population of model GCs by using random combinations of mossy fibers recordings as inputs to integrate-and-fire neurons. Although not in the form of a delay line, the temporal structure of the resulting GC basis is sufficient to allow an aSTDP rule to account for experimentally observed negative images.

T-39. Structured chaos and spike responses in stimulus-driven networks

Guillaume Lajoie¹
Kevin K. Lin²
Eric Shea-Brown¹

GLAJOIE@AMATH.WASHINGTON.EDU
KLIN@MATH.ARIZONA.EDU
ETSB@WASHINGTON.EDU

¹University of Washington

²University of Arizona

Large, randomly coupled networks of excitatory and inhibitory neurons are ubiquitous in neuroscience, and are known to autonomously produce chaotic dynamics. In general, chaos represents a threat to the reliability of network responses: if the same signal is presented many times with different initial conditions, there is no guarantee that the system will entrain to this signal in a repeatable way. As a consequence, it is possible that computations carried out by chaotic networks cannot rely on precise spike timing to carry information and must therefore depend on coarser statistical quantities such as firing rates. This motivates twin questions addressed in the present theoretical work. First, what is the impact of temporal inputs on the presence of chaotic dynamics in balanced networks of spiking, excitable neurons? Second, when networks remain chaotic in the presence of stimuli, how does this impact the structure and repeatability of their spiking output? We find intriguing answers on both counts. First, aperiodic temporal inputs can strongly enhance or diminish the presence of chaos, as assessed by the number and strength of positive Lyapunov exponents compared with those in spontaneously firing balanced networks. Second, and surprisingly, repeatable and precise patterning of spike responses occurs even in the presence of many positive Lyapunov exponents. Importantly, these structured responses arise when synaptic interactions and the stimulus are of similar order, so they are not simply a consequence of overwhelming external drive. Rather, they result from the fact that network interactions force dynamics to evolve along low-dimensional attracting sets. This appears to be a general property of spiking in sparse, balanced networks, and could widen the possibilities by which stimulus information is encoded by such networks.

T-40. Stimulus-response associations shape corticostriatal connections in an auditory discrimination task

Petr Znamenskiy
Qiaojie Xiong
Anthony Zador

ZNAMENSK@CSHL.EDU
XIONG@CSHL.EDU
ZADOR@CSHL.EDU

Cold Spring Harbor Laboratory

Plasticity of corticostriatal connections is thought to underlie reinforcement learning. However, how the changes in corticostriatal transmission establish associations between sensory stimuli and motor responses is not known. To address this question, we investigated the changes in corticostriatal connectivity of the auditory cortex during acquisition of an auditory discrimination task in rats. We first developed a method for interrogation of the strength of corticostriatal connections *in vivo*. We virally expressed Channelrhodopsin-2 in the cortex and implanted optical fibers coupled to tetrodes in the striatum, targeting the axons of corticostriatal neurons. Light stimulation drove neurotransmitter release and generated excitatory currents in striatal neurons, which could be detected in extracellular recordings. Consistent with their synaptic origin, responses peaked ~3 ms after light onset and adapted to high frequency stimulation. We used this light-evoked field potential response as a proxy for the strength of corticostriatal connections at the population level. We trained rats in a two-alternative choice frequency discrimination task and measured the strength of striatal outputs of auditory cortical neurons tuned to different sound frequencies, exploiting the fact that cortical inputs with different frequency tuning are spatially segregated in the striatum. Acquisition of the task resulted in rapid and persistent changes in corticostriatal functional connectivity, selectively potentiating cortical inputs tuned to frequencies associated with contralateral choices. In subjects trained to associate high frequencies with contralateral choices, the magnitude of light-evoked responses was positively correlated with the preferred frequency of the striatal site. The trend was reversed in subjects trained to associate low frequencies with contralateral choices. Therefore, the association between sound frequency and appropriate behavioral response is reflected in the weights of corticostriatal connections. These results suggest a straightforward mechanism, through which sensory stimuli drive action selection in the basal ganglia, consistent with existing models of corticostriatal plasticity in reinforcement-based decision-making tasks.

T-41. Whole-brain neuronal dynamics during virtual navigation and motor learning in zebrafish

Misha B Ahrens^{1,2}

Kuo-Hua Huang

Drew Robson³

Jennifer Li³

Michael Orger⁴

Alexander Schier³

Ruben Portugues³

Florian Engert³

AHRENSM@JANELIA.HHMI.ORG

KHHUANG@MCB.HARVARD.EDU

DREW.ROBSON@GMAIL.COM

JENMBLI@GMAIL.COM

MICHAEL.ORGGER@NEURO.FCHAMPALIMAUD.ORG

SCHIER@FAS.HARVARD.EDU

RPORTUGUES@MCB.HARVARD.EDU

FLORIAN@MCB.HARVARD.EDU

¹Janelia Farm Research Campus

²HHMI

³Harvard University

⁴Champalimaud Centre for the Unknown

We present a paradigm for recording anywhere in the brain, from several thousands of individual neurons at a time, in larval zebrafish behaving in virtual reality environments, during navigation and motor learning. The animals are paralyzed, and bilateral recordings from motor neuron axons in the tail provide sufficient information for decoding intended forward swims and turns. These intended actions are converted into motion in a virtual environment, rendering realistic visual feedback in response to fictive locomotion. Simultaneously, a two-photon microscope scans over the brain of transgenic zebrafish expressing a genetically encoded calcium indicator in all neurons. In this way, activity in large populations of neurons, that may cover the entire brain, can be monitored during diverse behaviors. Whole-brain activity is monitored as fish exhibit three behaviors analogous to the freely swimming counterparts: First, the two dimensional optomotor response; second, darkness avoidance; and third, motor learning. During the 2D optomotor response, whole-hindbrain recordings reveal functional networks involved in forward swimming, left- and right turns. During lateralized swimming, activity is lateralized, but this organization is reversed in part of the cerebellar cortex. During darkness avoidance, neurons in the habenula and the pretectum respond to luminosity in distinct spatial receptive fields. During motor learning, many brain areas are active during different phases of the behavior — the learning and the maintenance periods — with strong neuronal activation in the cerebellum and the inferior olive, brain structures that are involved in motor learning in mammals. Lesioning the latter structure leads to a loss of the behavior. Statistical methods, including dimensionality reduction, reveal multiple temporal profiles of neuronal activity, localizing to distinct brain areas, suggestive of a functional network architecture. Such whole-brain recordings during behavior, in combination with computational techniques for the analysis of these high dimensional data, will generate new insights into circuit function underlying behavior.

T-42. A network model for learning motor timing in songbirds

Cengiz Pehlevan

Farhan Ali

Timothy M. Otchy

Bence Olveczky

Harvard University

CENGIZ.PEHLEVAN@GMAIL.COM

FARNALI@YAHOO.COM.SG

TIM.OTCHY@GMAIL.COM

OLVECZKY@FAS.HARVARD.EDU

Timing precision in motor output is fundamental to the mastery of a variety of motor skills, yet how the nervous system learns and adaptively modifies temporal structure of performance is not well understood. Zebra finches, with their precise learned vocalizations, provide a unique opportunity for addressing this question. To explore how temporal changes in birdsong are implemented in underlying circuitry, we developed an experimental paradigm to induce changes in the duration of targeted song elements. Chronic recordings from premotor nucleus HVC in the context of this paradigm showed neural dynamics in this putative ‘time-keeper’ circuit to co-vary with changes in song timing. Here, we present a biologically plausible computational model of the HVC network that can account

for our observations and, more generally, provide a circuit level explanation for learning in the temporal domain. Our model treats HVC as a synfire chain, and adds different sources of variability to its nodes to account for temporal variability in song, and to allow for ‘exploration’ in the time domain. Synapses between nodes in the chain are updated based on a reward function that is the sum of an externally delivered reinforcement and an internally generated ‘template’ component that serves to preserve the original temporal structure. Competition between these factors explains key observations about dynamics of learning. Finally, we discuss the advantages of coding and learning motor timing in a synfire-chain network as opposed to a fully recurrent network. Our model represents a simple and potentially general network solution to the problem of learning timing in motor output.

T-43. On the duality of motor cortex: movement representation and dynamical machine

Philip N. Sabes

SABES@PHY.UCSF.EDU

University of California, San Francisco

Although primary motor cortex (M1) has been the focus of decades of intense study, there remains little consensus on how the activity in M1 gives rise to coordinated movement. The traditional approach to studying M1 focuses on determining the movement parameters that are encoded by the area, and most researchers in the field design experiments within this representational framework. Recently, a challenge of this framework has come from researchers who view motor cortex as a dynamical machine that generates patterns of activity that form temporal basis functions for downstream control. This dynamical systems framework has received support from the observation that motor cortical activity resembles a linear system with rotational dynamics, suggesting that M1 may form a Fourier basis. Here I show that these views can be reconciled. First, I show that the dynamics of M1 can be largely reproduced by simple, feedforward linear networks. Second, I show that this network creates temporal basis functions that can control a model two-joint arm. This result is consistent with the interpretation of M1 as a non-representational dynamical machine: network initial conditions and network dynamics are arbitrary, within limits, and motor learning only requires finding the right projection from M1 to muscle activation. Importantly, however, this network cannot generalize unless the initial conditions vary smoothly with movement parameters such as target location, i.e. the network must ‘represent’ target in its initial activity. Cosine tuning of target location (a classic representational model) forms a good set of initial conditions for movement generalization. Third, I show that the emphasis on basis functions is not required. Control of arm movements is equally good when network dynamics and downstream projections are fixed, and instead the cosine tuning parameters are learned. This simulation evokes the traditional representational view that pre-movement activity in motor cortex represents movement targets.

I-1. Independent pools of visual short-term memory resource for different features

Hongsup Shin

HSHIN@CNS.BCM.EDU

Ronald Van den Berg

RB2@BCM.EDU

Wei Ji Ma

WJMA@BCM.EDU

Baylor College of Medicine

There is an ongoing debate on whether the basic units of visual short-term memory (VSTM) are objects or features. When an object has two features, for example orientation and color, is each feature remembered equally well as when the object has only one relevant feature, or is memory resource divided over features? This question has not been examined in the context of the currently best available model of VSTM, the variable-precision model (1). In this model, memory resource is a continuous quantity that affects mnemonic precision, and is variable from trial to trial and from item to item. Here, we use a change localization paradigm and the variable-precision model

to examine whether or not mnemonic precision is divided over features. Stimuli were colored, oriented ellipses. Subjects briefly viewed four ellipses with both feature values drawn independently from uniform distributions. After a delay, a second display containing four ellipses appeared, three of which were the same as in the first display and one of which had changed. In the one-feature conditions, the change occurred always in the same feature — either orientation or color, depending on the experimental session. In the two-feature condition, on every trial, the change occurred randomly in either feature. Observers reported the location of the change. We tested two optimal-observer models, which differed in their resource allocation. In the independent-resource model, mean precision for a given feature is identical between the one-feature and two-feature conditions. In the shared-resource model, the mean precision for a given feature is a proportion of the mean precision in the corresponding one-feature condition. We found that the independent-resource model better explains subject behavior. This suggests that the mnemonic precision of a feature is not be affected by the number of features, supporting the idea of objects being the basic units of VSTM.

I-2. Tactile working memory in rat and human: Prior competes with recent evidence

Athena Akrami^{1,2}

Arash Fassih¹

Vahid Esmaeili¹

Mathew E Diamond¹

ATHENA.AKRAMI@GMAIL.COM

FASSIHI@SISSA.IT

VESMAEIL@SISSA.IT

DIAMOND@SISSA.IT

¹School for Advanced Studies (SISSA)

²Princeton

We speculate that during a delayed comparison task, the memory trace of a stochastic stimulus becomes noisy and unreliable and shifts progressively toward the center of a prior distribution built from the past experience. In this scenario, the decision of the subject would be based on a comparison between the second stimulus and the shifting-toward-prior trace of the first stimulus (Fig. 1). We test our hypothesis in a tactile working memory paradigm, adapted to human and rats, by studying the effect of (1) increasing the delay duration, and hence the purported amount of shift toward the prior, and (2) changing the prior distribution. Two sequential vibratory stimuli were characterized by velocity standard deviation, δ_1 and δ_2 ; subjects had to judge whether as $\delta_2 < \delta_1$ or $\delta_2 > \delta_1$. Rats and humans perform equally well for short delays. Longer delays instead result in augmented Contraction Bias (Ashourian and Loewenstein 2011): for small δ values, subjects tend to more frequently judge as $\delta_2 < \delta_1$ even when the actual value of δ_2 was larger than δ_1 . This fits a model where the memory of δ_1 drifts towards the prior expected value, and at long delays rises above the value of δ_2 . These results demonstrate the similarities of rats and human's cognitive system when incorporating expectations and probabilistic inference in perception.

I-3. Mechanisms and circuitry underlying basal forebrain enhancement of top-down and bottom-up attention

Michael Avery

Nikil Dutt

Jeffrey Krichmar

University of California, Irvine

AVERYM@UCI.EDU

DUTT@ICS.UCI.EDU

JKRICHMA@UCI.EDU

Both attentional signals from the frontal cortex and neuromodulatory signals from the basal forebrain (BF) have been shown to have a strong influence on information processing in the primary visual cortex (V1). These two systems are highly interactive and exert complementary effects on their targets, including increasing firing rates and decreasing interneuronal correlations. One interesting dichotomy that has arisen from experimental research, however, is that the cholinergic system is important for increasing V1's sensitivity to both sensory and attentional

information. In order to see how the basal forebrain and top-down attention act together to modulate sensory input, we developed a spiking neural network model of V1 and thalamus that incorporates cholinergic neuromodulation and top-down attention (Figure and model details below). Our model was able to match experimental data showing that neuromodulatory projections from the basal forebrain and top-down attentional signals enhance cortical coding by decreasing interneuronal correlations and increasing between-trial reliability of neurons in the visual cortex. In accordance with recent experiments and models, we further showed that interneuronal decorrelation is primarily mediated by inhibitory neurons. In addition to this, we suggest 'global' and 'local' modes of action by which the basal forebrain may be enhancing bottom-up sensory input and top-down attention, respectively. That is, activation of the basal forebrain itself decreases the efficacy of top-down projections and increases the reliability of bottom-up sensory input by blocking top-down attentional inputs in the thalamus. In contrast, local release of acetylcholine in the visual cortex, which is triggered by top-down glutamatergic projections, can enhance top-down attention with high spatial specificity. These findings will lead to a better understanding of the basal forebrain and its interactions with attentional signals, and provide mechanisms for how the basal forebrain can enhance both top-down attentional signals and bottom-up sensory input.

I-4. Semi-parametric Bayesian entropy estimation for binary spike trains

Evan Archer
Il M Park
Jonathan W Pillow
University of Texas at Austin

EVANARCHER@GMAIL.COM
MEMMING@AUSTIN.UTEXAS.EDU
PILLOW@MAIL.UTEXAS.EDU

The set of possible neural response patterns is frequently so large that its distribution cannot be reliably estimated from limited data. Consequently, information and entropy estimation for neural data presents a significant challenge which has been met by a diverse literature. Most entropy estimators in this literature, however, are general purpose in that they are designed to work on as broad a class of data-generating distributions as possible. For neural applications all such general-purpose estimators have a critical weakness: they ignore by design much of our strong prior knowledge about the structures of spike trains. Neural response patterns, however, are not arbitrary: we can apply our prior knowledge about the basic statistical structure of spike trains to entropy estimation. Here, we augment the nonparametric Bayesian entropy estimation method (Archer et al) by incorporating a simple, parametric model of spike trains. Intuitively, we wish to incorporate our prior knowledge that spikes are rare events, and we assign lower prior probability to words with more spikes. Mathematically, we model a spike word as a vector of independent Bernoulli random variables, each with a probability p of firing. Under this model, for typical values of p , very sparse vectors are much more likely than those with many spikes. Alone, this simple model does not provide a good method for entropy estimation, as it cannot flexibly account for data drawn outside the model class. However, by "centering" a Dirichlet process on this parametric model, we obtain a semi-parametric model that can model arbitrary discrete distributions. We derive a computationally efficient entropy estimator under the model, and for real data, we show that this model outperforms conventional estimators.

I-5. Sparse coding model and population response statistics to natural movies in V1

Mengchen Zhu^{1,2}

Ian Stevenson³

Urs Koster³

Charles Gray⁴

Bruno Olshausen³

Christopher J Rozell¹

¹Georgia Institute of Technology

²Emory University

³University of California, Berkeley

⁴Montana State University

MCZHU@GATECH.EDU

I-STEVENSON@BERKELEY.EDU

URS@BERKELEY.EDU

CMGRAY@CNS.MONTANA.EDU

BAOLSHAUSEN@BERKELEY.EDU

CROZELL@GATECH.EDU

Local populations of sensory cortical cells exhibit a diverse range of activity patterns. However, classical approaches have neither fully accounted for nor characterized this heterogeneity, especially in response to natural stimuli. First, classical single cell recordings suffered from sampling bias and favored highly responsive cells (Olshausen and Field, 2005). Second, common approaches considered mostly the average activity over different cell classes, without a full description of the statistical distribution over the entire population (Wohrer, Humphries, and Machens, 2012). Recent studies started to address these issues (Yen, Baker, and Gray, 2007; Herikstad et al., 2011). In this study, we make further inroads by recording simultaneous single unit activities across cortical layers in cat V1 in response to natural movies using a silicon polytrode, and comparing the population statistics to the predictions from a dynamical system implementation of the sparse coding model (Olshausen and Field, 1996; Rozell et al., 2008). We show that: (1) The population firing rate distribution is close to exponential in both the recorded data and the sparse coding model in response to natural movies; (2) The response correlation between unit activities is small regardless of the size of receptive field overlap, when using a binning window synced to the movie frame. A similar relationship between the response correlation and receptive field overlap is observed in the sparse coding model; (3) A linear-nonlinear model could predict the exponential firing rate distribution, but not the correlation structure.

I-6. Efficient hierarchical receptive field estimation in simultaneously-recorded neural populations

Kamiar Rahnama Rad

Carl Smith

Liam Paninski

Columbia University

KAMIAR@STAT.COLUMBIA.EDU

CAS2207@COLUMBIA.EDU

LIAM@STAT.COLUMBIA.EDU

A major trend in systems neuroscience is to record simultaneously from large neuronal populations. A key objective in statistical neuroscience is to develop scalable and efficient methods for extracting as much information as possible from these recordings. One important direction involves hierarchical statistical modeling: estimating receptive fields (RFs) (or motor preferences) one neuron at a time is highly suboptimal, and in many cases we can do much better by sharing statistical information across neurons. In particular, we can exploit the fact that nearby neurons often have similar receptive fields. Here “nearby” might be defined topographically (e.g., in the case of cat primary visual cortex, where nearby neurons typically have similar orientation preferences) or more abstractly, in terms of, e.g., shared genetic markers. We discuss two approaches for exploiting neighborhood information. The first method maximizes an appropriately penalized likelihood: we penalize deviations between neighboring RFs and compute the corresponding maximum a posteriori RF map. We use a smooth convex penalizer that allows for large occasional breaks or outliers in the inferred RF map. Posterior confidence intervals can be obtained here via “MAP-perturb” trick (Papandreou and Yuille, 2011). The second method is based on direct Gibbs sampling from the posterior, where the prior is of “low-rank” form, which enables fast direct sampling (Smith et al., 2012).

Both approaches are computationally tractable, scalable to very large populations, and avoid imposing any overly restrictive constraints on the inferred RF map that would lead to oversmoothing. The first method is computationally cheaper, but the second method is able to model RFs in non-vector spaces (e.g., orientation). Both methods are equally applicable to multineuronal spike train or imaging data, and can dramatically reduce the experimental time required to characterize RF maps to the desired precision.

I-7. Bayesian decision-making using neural attractor networks

Ingmar Kanitscheider¹
 Jeffrey Beck²
 Alexandre Pouget¹

INGMAR.KANITSCHIEDER@UNIGE.CH
 JBECK@BCS.ROCHESTER.EDU
 ALEX@CVS.ROCHESTER.EDU

¹University of Geneva

²University of Rochester

There is substantial evidence that human and animal decision-making is compatible with Bayesian inference, which predicts that sensory information is combined with prior information according to its reliability (Knill & Richards, 1996). How such a Bayesian decoder is implemented on a neural level at the decision stage still very much remains a matter of debate. Several groups have proposed to implement the prior implicitly by sampling the preferred stimuli of the neurons from the prior distribution (Shi & Griffiths 2009; Fischer & Pena 2011, Ganguli & Simoncelli COSYNE 2012). The drawback of those implementations is that they require tuning curves proportional to the likelihood function, which is in conflict with the experimental finding of contrast-invariance. We investigate here how to implement a prior when the likelihood function is represented by a probabilistic population codes (PPC), a type of code compatible with contrast invariance. In the case of a flat prior, Deneve et al 1999 showed that a line attractor neural network can implement a robust and biologically plausible maximum likelihood decoder, as long as the likelihood function is encoded with a PPC using the exponential family with linear sufficient statistics. We generalized this result to arbitrary prior distribution which we represent in the initial pattern of activity. The problem with this solution is that the network initial activity is far from the stable attractor, while near optimal performance requires that the network starts close to the attractor. This issue can be alleviated by including an additional linear prefiltering step. We show that the resulting network has close-to-optimal performance and that our results are robust to the presence of a baseline and to nuisance parameters like contrast changing from trial to trial. Similar results can be obtained by implementing the prior distribution in the synaptic weights.

I-8. Robust online estimation of noisy sparse signals by threshold-linear and integrate-and-fire neurons

Tao Hu^{1,2}
 Alexander Genkin³
 Dmitri Chklovskii^{1,2}

HUT@JANELIA.HHMI.ORG
 ALEXGENKIN@INAME.COM
 MITYA@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus

²HHMI

³AVG Consulting

Threshold-linear response functions are ubiquitous in the brain, from synaptic transmission as a function of membrane voltage in graded potential neurons (Field & Rieke, 2002, Asari & Meister, 2012) to firing rate as a function of somatic current in spiking neurons (Dayan & Abbott, 2001). What computational role could such non-linearity play? Here, we explore the hypothesis that threshold linear response function arises from robust online estimation of sparse (or Laplacian) signals contaminated by Gaussian noise. Inspired by online convex optimization (OCO) (Zinkevich, 2003, Dar & Feder, 2011) we derive a robust estimation algorithm by balancing two terms: empirical loss (mean square error for Gaussian noise) and l1-norm regularization (for sparse and/or Laplacian signals).

The resulting algorithm is an integrator followed by threshold-linear function. Using the OCO framework allows us to prove performance guarantees for such an algorithm, not just on average for a given signal ensemble but for the worst case scenario. Thus, threshold-linear estimator has exceptional robustness, operating asymptotically no worse than offline estimator not just on a stationary signal but also on non-stationary, chaotic and even adversarial signals, which may explain its ubiquity in the brain. We confirm the superiority of threshold-linear neurons over linear for sparse input signals numerically. Finally, when the estimation signal must be encoded by a binary sequence, the estimator reduces to an integrate-and-fire neuron.

I-9. A generative model of natural images as patchworks of textures

Matthias Bethge¹
Niklas Luedtke^{2,3}
Debapriya Das³
Lucas Theis⁴

MBETHGE@TUEBINGEN.MPG.DE
NIKLU@AOL.COM
D.DEBAPRIYA@GMAIL.COM
LUCAS@BETHGELAB.ORG

¹Max Planck Institute

²BCCN

³CIN

⁴Centre for Integrative Neuroscience

Natural images can be viewed as patchworks of different textures, where the local image statistics is roughly stationary within a small neighborhood but otherwise varies from region to region. In order to model this variability, we first applied the parametric texture algorithm of Portilla and Simoncelli to image patches of 64x64 pixels in a large database of natural images such that each image patch is then described by 655 texture parameters which specify certain statistics, such as variances and covariances of wavelet coefficients or coefficient magnitudes within that patch. To model the statistics of these texture parameters, we then developed suitable nonlinear transformations of the parameters that allowed us to fit their joint statistics with a multivariate Gaussian distribution. We find that the first 200 principal components contain more than 99% of the variance and are sufficient to generate textures that are perceptually extremely close to those generated with all 655 components. We demonstrate the usefulness of the model in several ways: (1) We sample ensembles of texture patches that can be directly compared to samples of patches from the natural image database and can to a high degree reproduce their perceptual appearance. (2) We further developed an image compression algorithm which generates surprisingly accurate images at bit rates as low as 0.14 bits/pixel. Finally, (3) We demonstrate how our approach can be used for an efficient and objective evaluation of samples generated with probabilistic models of natural images.

I-10. A non-parametric Bayesian prior for causal inference of auditory streaming

Ulrik Beierholm
Tim Yates
University of Birmingham

BEIERH@GMAIL.COM
T.A.YATES@BHAM.AC.UK

Segregation and grouping of auditory stimuli is a necessity for effective parsing of auditory information for e.g. recognition or understanding of speech. One example in the psychophysics literature is based on the segregation of a sequence of tones into either one or two streams as a function of the relationship of the tones (Bregman & Campbell 1971). Such perceptual grouping of sequential auditory stimuli has traditionally been modeled using a mechanistic approach (e.g. McCabe & Denham 1997). The problem however is essentially one of source inference; inferring which tones belong within which stream. This is a problem that in the visual and multi-sensory domains have recently been modeled using Bayesian statistical inference (Koerding et al. 2007). These models have so far been based on parametric statistics and thus restricted to performing inference over just one or two

possible sources, however human perceptual systems have to deal with much more complex scenarios. We have developed a Bayesian model that allows an unlimited number of signal sources to be considered: it is general enough to allow any discrete sequential cues, from any modality. The model uses a non-parametric prior, the Chinese Restaurant Process, so that increased complexity of the data does not necessitate more parameters (see e.g. Orbanz & The 2010 for a review). The model not only determines the most likely number of sources, but also specifies the source that each signal is associated with. The model gives an excellent fit to data from an auditory stream segregation experiment in which participants estimated the number of sources generating a sequence of pure tones that varied in pitch and presentation rate. We propose that this mechanism is a general feature of perceptual organizing of stimuli.

I-11. Correlations strike back (again): the case of associative memory retrieval

Cristina Savin¹
Peter Dayan²
Mate Lengyel¹

CS664@CAM.AC.UK
DAYAN@GATSBY.UCL.AC.UK
M.LENGYEL@ENG.CAM.AC.UK

¹University of Cambridge

²University College, London

It has long been recognised that statistical dependencies in neuronal activity need to be taken into account when decoding stimuli encoded in a neural population. It is far less well appreciated that the same decoding challenges arise in the context of autoassociative memory, when retrieving information stored in correlated synapses. Such correlations have been well documented experimentally (Song et al, 2005); here we show how they can arise between synapses that share pre- or post-synaptic partners when any of several well-known additive (Hopfield, 1982) or metaplastic (Fusi et al, 2005) learning rules is applied. To assess the importance of these dependencies for recall, we adopt the strategy of comparing the performance of decoders which either do, or do not, take them into account, but are otherwise optimal, showing that ignoring synaptic correlations has catastrophic consequences for retrieval. We therefore study how recurrent circuit dynamics can implement decoding that is sensitive to correlations. Optimal retrieval dynamics in the face of correlations require substantial circuit complexities. By contrast, we show that it is possible to construct approximately optimal retrieval dynamics that are biologically plausible. The difference between these dynamics and those that ignore correlations is a set of non-linear circuit motifs that have been suggested on experimental grounds, including forms of feedback inhibition and experimentally observed dendritic nonlinearities (Branco et al, 2011). We therefore show how assuaging an old enemy leads to a novel functional account of key biophysical features of the neural substrate.

I-12. Timescale-dependent shaping of correlations by stimulus features in the auditory system

Ana Calabrese
Sarah M N Woolley
Columbia University

AMC2257@COLUMBIA.EDU
SW2277@COLUMBIA.EDU

A current challenge in systems neuroscience is to understand how networks of neurons process and encode sensory information. Whether or not neurons in a network interact to optimally represent an ensemble of stimuli is a matter of intense debate. We investigated this issue in the songbird auditory forebrain by characterizing changes in the structure of spike train correlations in response to different complex sound stimuli. We used multielectrode arrays to record the spiking activity of populations of neurons of anesthetized male zebra finches (ZF) during silence (i.e. spontaneous activity), and during the processing of two stimuli with different spectrotemporal correlations: ZF songs and modulation limited noise (ml-noise), a synthetic sound for which the maximum spectral

and temporal modulations match those of song. Spike-train correlations between simultaneously recorded neurons were then measured. We found distinct populations of putative excitatory (RS) and putative inhibitory (FS) neurons. Populations differed in action potential width (RS>FS), firing rates (FS>RS), latencies (RS>FS), and response variability across trials (RS>FS). Differences in single-cell properties were accompanied by significant differences in the cross-correlograms (CCGs) of cell pairs; CCG peaks were highest and narrowest for FS-FS pairs and smallest and widest for RS-RS pairs. Spatial and spectral profiles of correlations also differed among pair types. Differences across pair types were maintained under all stimulus conditions. The effects of stimulus correlations on response correlations were timescale and pair type specific. At short timescales (1-5 ms), response correlations between FS-FS pairs were higher during song processing than during ml-noise processing. At longer timescales (on the order of 100 ms), correlations were lower during song processing. We apply a simple network model to study the effect of timescale-dependent changes of correlations on encoding and discrimination of songs.

I-13. Compensation for neuronal circuit variability

Tilman Kispersky
Eve Marder
Brandeis University

TILMAN@BRANDEIS.EDU
MARDER@BRANDEIS.EDU

Neuronal variability between individuals of a population is a core feature of nervous systems. In a given circuit, neuronal variability can be seen in the number neurons present, their channel densities, or their synaptic connectivity pattern. Variability is thought to provide robustness to neuronal circuits because it allows complicated circuits multiple ways in which they can implement the same behaviorally relevant outputs. In a functioning circuit, variability is compensated such that overall functional output is maintained. In this work we study whether variance in neuronal circuits is compensated for within neurons themselves, or on the network level via modulation of synapses, or to which degree both of these mechanisms are used.

I-14. Precise characterization of multiple LIP neurons in relation to stimulus and behavior

Jacob Yates
Il M Park
Lawrence Cormack
Jonathan W Pillow
Alexander Huk
University of Texas at Austin

JLYATES@UTEXAS.EDU
MEMMING@AUSTIN.UTEXAS.EDU
CORMACK@MAIL.UTEXAS.EDU
PILLOW@MAIL.UTEXAS.EDU
HUK@UTEXAS.EDU

Over 3 decades of research have elucidated how the primate brain encodes visual motion and have suggested how later stages may accumulate evidence for making perceptual decisions about motion. This work has suggested a circuit model relating motion stimuli of varying direction and coherence, sensory responses in visual area MT, ramping responses in posterior parietal area LIP, and perceptual decisions. We seek to precisely test this model by characterizing all of these facets simultaneously, with the temporal and spatial precision afforded by reverse-correlation methods. We trained a monkey to perform a novel motion discrimination task with statistically independent motion pulses (144 ms) of variable coherence and direction, while we recorded from LIP neurons (often 2-4 simultaneously). This task allows us to measure how temporal variations in motion pulse strength correlate with both the monkeys choice and how it affects the spiking of LIP neurons. Using Bayesian logistic regression to classify the monkeys choices and a generalized linear model (based on Poisson regression) to fit single trial spike trains, we were able to recover and compare the temporal weighting function of both monkey and neuron(s) within individual experimental sessions. Over several experiments (n= 10), on average the monkey

weighted the beginning of the trial more than the end. But there was variation from day to day, such that in some sessions motion pulses at the end of the trial correlated most strongly with the monkeys choice. The corresponding relations between motion pulses and LIP responses were considerably more complex, implying an indirect relation between LIP and decisions. Furthermore, we often observed fine timescale cross-correlations between neurons with overlapping RFs that were time- and decision-dependent, suggesting shared input. These results re-veal the feasibility of characterizations of sensorimotor transformations during motion decisions with greatly increased precision.

I-15. Common neural correlates of adaptive control in the anterior cingulate cortex of rats and humans

Nandakumar Narayanan¹

James Cavanagh²

Marcelo Caetano³

Michael Frank²

Mark Laubach³

NANDAKUMAR.NARAYANAN@GMAIL.COM

JIM.F.CAV@GMAIL.COM

MARCELOSCAETANO@GMAIL.COM

MICHAEL_FRANK@BROWN.EDU

MARK.LAUBACH@GMAIL.COM

¹University of Iowa

²Brown University

³Yale University

Adaptive control enables adjustments in performance based on prior behavioral outcomes. EEG, fMRI and lesion studies in humans have established that the anterior cingulate cortex (ACC) is crucial for adaptive control. Progress on understanding the neurobiological basis of adaptive control, e.g. using multi-electrode single-unit recordings and optogenetic methods, will require the use of common behavioral tasks to assess adaptive control in animals and humans. By using such tasks, it should be possible to determine if there are common neural correlates of adaptive control across species. Here, we show for the first time that rats and people present the same signatures of adaptive control when performing an equivalent behavioral task. Participants performed a timing task in which they pressed on a lever (rats) or button (humans) over a fixed temporal interval and responded promptly at the end of the interval. Errors occurred if responses were made too soon or late. Recordings of spike activity and field potentials were made in the ACC of 8 younger (4-6 mo.) and 4 older (24 mo.) rats. EEG recordings were made in 12 young adult humans. Four main findings were (1) low-frequency oscillations (below 10 Hz) in the human and rat ACC were enhanced after errors and correlated with subsequent improvements in performance, (2) spike activity in the rat ACC encoded prior behavioral outcomes and was coherent with low-frequency field potential oscillations, (3) inactivating the rat ACC diminished adaptive control and eliminated coherence between spikes and fields in the motor cortex, and (4) the aged ACC failed to show dynamic low-frequency oscillations due to prior behavioral outcomes. These findings implicate a novel mechanism for implementing adaptive control via low-frequency oscillations and provide a mechanistic description of how medial frontal networks synchronize neuronal activity in other brain regions to guide performance.

I-16. Neural correlates of target detection in the human brain

Arjun Bansal^{1,2}

Alexandra Golby³

Joseph Madsen⁴

Gabriel Kreiman²

MINDSCIENCE@GMAIL.COM

AGOLBY@PARTNERS.ORG

JOSEPH.MADSEN@CHILDRENS.HARVARD.EDU

GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU

¹Boston Children's Hospital

²Harvard Medical School

³Brigham and Women's Hospital

⁴Childrens Hospital Boston

Target detection constitutes an important step towards visual recognition in a variety of tasks such as visual search and goal-directed recognition. We sought to examine where, when and how target detection modulates the responses during visual shape recognition in the human brain. We recorded intracranial field potential (IFP) activity from 794 electrodes implanted in 10 subjects for epilepsy localization during a target detection task. Subjects were presented with brief flashes (100ms) of images containing either one or two objects belonging to 5 possible categories. In each block (50 trials), one of those 5 categories was the target and subjects reported whether the image contained an object from the target category or not (two alternative forced choice). We observed robust differences in the IFP responses when the target was present compared to when it was absent. In other words, the IFP responses to the same objects were significantly modulated based on whether they were assigned as the target or non-target for each individual trial. Target-dependent responses were predominantly localized in inferior temporal gyrus (19.1%), fusiform gyrus (12.4%), and middle temporal gyrus (12.4%). The average latency of target-dependent modulation was $361.58pm118.14ms$, which is longer than the latency of the visual selective responses (100-150ms). The latency to target-modulation was shorter in inferior temporal gyrus ($315.84pm130.83ms$, $n=17$) and fusiform gyrus ($329.24pm85.77ms$, $n=11$), compared to the supramarginal inferior parietal gyrus ($396.09pm67.91ms$, $n=5$) and middle frontal gyrus ($394.01pm109.54ms$, $n=3$). Target-dependent modulation was observed in the total power as well as in power in different frequency bands. These results demonstrate strong task-dependent signals in cortex and provide spatial and dynamical constraints for the development of computational models describing the role of top-down signals in feature-selective attention and vision.

I-17. A computational theory of action-based decisions

Vassilios Christopoulos¹
Paul Schrater²

VCHRISTO@CALTECH.EDU
SCHRATER@GMAIL.COM

¹California Institute of Technology

²University of Minnesota

Imagine that you are facing with the challenge of deciding whether to go for dinner tonight. One of the options is to drive to downtown to one of your favorite restaurants, and the other one is to walk to a nearby fast food place. How do you decide between the alternative options? Recent experimental findings suggest that decisions are made through a biased competition between alternative actions. This is in line with the 'action-based' decision theory, which suggests that the brain generates several concurrent policies associated with alternative goals that compete for action selection and uses perceptual information to bias this competition. Despite the strong experimental evidence, little is known about the computational mechanisms of this competition. In the current study, we propose a computational framework that provides the first detailed instantiation of the action-based decision theory. We show how competition emerges from an optimal control framework to model behavior in decision tasks with multiple goals. The complex problem of action selection is decomposed into a weighted mixture of individual control policies, each of which produces a sequence of actions associated with a particular goal. The framework integrates online information related to decision variables into an evolving assessment of the desirability of each goal. The desirability reflects 'how' desirable is a goal at a given state and acts as weighted factor of the individual policies. We evaluated the performance of the framework in reaching and saccade tasks. We show that the averaging behavior observed in rapid reaching tasks with competing goals can be qualitatively predicted by our framework. The model provides also insights on how decision variables affect the behavior and how the action competition can lead to errors in behavior. Our findings suggest that this framework can be successfully used to understand the computational mechanisms of action-based decisions.

I-18. Individual FEF neurons code both task difficulty and an internally generated error signal

Tobias Teichert
Dian Yu
Vincent Ferrera
Columbia University

TT2288@COLUMBIA.EDU
DY2140@COLUMBIA.EDU
VPF3@COLUMBIA.EDU

It has been shown that subjects monitor their choices and preferentially revise initially wrong responses. These corrections occur even in the absence of external feedback and are thought to arise from continued post-decision stimulus processing. Though extracranial scalp recordings have revealed internally generated error signals (error related negativity, ERN) that may form the neural substrate of online error correction, the underlying neural circuits are still poorly understood. It has been suggested that error signals arise within circuits that primarily code response conflict (conflict monitoring theory). EEG and fMRI studies cannot test this assumption; while error and conflict signals have been suggested to arise from the same brain regions, the poor spatial resolution does not allow us to distinguish whether the same neurons code both signals. So far, single-cell studies have not confirmed the existence of neurons that code both response conflict (or difficulty) and an internally generated error signal. The current study addresses this issue and reports internally generated error signals in single neurons of the frontal eye fields (FEF) of macaque monkeys performing a reward-biased speed-categorization task. These cells preferentially increased their firing rate after a wrong response. Interestingly, many of these cells also responded more strongly to difficult (correct) trials, and these signals preceded the error signals on average by ~73 ms. These findings provide a deeper understanding of the neural circuits that generate error signals. The joint coding of error and difficulty in the same neurons, as well as the presence of error signals in dorso-lateral prefrontal rather than medial frontal cortex, provide interesting challenges to current decision-making theories.

I-19. Perceptual decision-making in a sampling-based neural representation

Ralf Haefner¹
Pietro Berkes¹
Jozsef Fiser²

RALF@BRANDEIS.EDU
PIETRO.BERKES@GOOGLEMAIL.COM
FISER@BRANDEIS.EDU

¹Brandeis University

²Central European University

Most computational models of the responses of sensory neurons are based on the information in external stimuli and their feed-forward processing. Extrasensory information and top-down connections are usually incorporated on a post-hoc basis only, e.g. by postulating attentional modulations to account for features of the data that feed-forward models cannot explain. To provide a more parsimonious account of perceptual decision-making, we combine the proposal that bottom-up and top-down connections subservise Bayesian inference as the central task of the visual system (Lee & Mumford 2003) with the recent hypothesis that the brain solves this inference problem by implementing a sampling-based representation and computation (Fiser et al 2010). Since the sampling hypothesis interprets variable neuronal responses as stochastic samples from the probability distribution that the neurons represent, it leads to the strong prediction that dependencies in the internal probabilistic model that the brain has learnt will translate into observable correlated neuronal variability. We have tested this prediction by implementing a sampling-based model of a 2AFC perceptual decision-making task and directly comparing the correlation structure among its units to two sets of recently published data. In agreement with the neurophysiological data, we found that: a) noise correlations between sensory neurons dependent on the task in a specific way (Cohen & Newsome 2008); and b) that choice probabilities in sensory neurons are sustained over time, even as the psychophysical kernel decreases (Nienborg & Cumming 2009). Since our model is normative, its predictions depend primarily on the task structure, not on assumptions about the brain or any additional postulated processes. Hence we could derive additional experimentally testable predictions for neuronal correlations, variability and performance as the task changes (e. g. to fine discrimination or dynamic task switching) or due to

perceptual learning during decision-making.

I-20. Model-free reinforcement learning predicts attentional selection during foraging in the macaque

Matthew Balcarras¹
Salva Ardid¹
Daniel Kaping¹
Stefan Everling²
Thilo Womelsdorf¹

MBALCARR@YORKU.CA
SARDID@YORKU.CA
DKAPING@YORKU.CA
SEVERLIN@UWO.CA
THIWOM@YORKU.CA

¹York University

²Western University

The deployment of top-down attention is frequently studied in cued tasks where attentional shifts are triggered by external cues indicating the relevant stimulus feature [1]. In contrast, foraging tasks lack external cues, so attentional shifts are internally triggered. These attentional shifts could rely on internal representations that track the value of stimulus features[2,3]. Here, we identified which feature values are internally encoded and how they are dynamically updated. We propose that attentional selection in each trial is explained by stochastic competition between the up-to-date values for the features appearing in the trial. Two macaques performed a foraging task (>80% rewarded trials) composed of two colored, drifting gratings that transiently rotated during each trial. To receive reward the monkey needed to attentionally select one stimulus to discriminate the stimulus rotation. Reward was associated with color, but the color-reward associations changed across blocks. Location, rotation direction, and the relative time of rotation were three variables with random reward associations, varying independently from the rewarded color. We analyzed monkeys' behavior by applying model-free and model-based reinforcement learning (RL). They differed respectively in whether all features in a trial compete to trigger stimulus selection, or whether only colors compete. Although both RL versions accounted for ~84% of the behavior variability in >1000 blocks, only model-free RL predicted the pattern of unrewarded trials across trials. This sub-optimal behavior naturally emerged from local correlations of non-color features (e.g. particular location and/or rotation) with reward across nearby trials that negatively biased attentional selection against the highly valued color for subsequent (~4) trials. The model-free RL suggests that attentional control emerges from a strong interaction between feature values operating as inputs and a stochastic covert choice operating as output, thereby constraining neural circuit models of attention and choice [2,4].

I-21. How can single neurons predict behavior?

Xaq S Pitkow¹
Sheng Liu²
Yong Gu³
Dora Angelaki²
Greg DeAngelis¹
Alexandre Pouget⁴

XPITKOW@BCS.ROCHESTER.EDU
SLIU@CNS.BCM.EDU
GUYONG@ION.AC.CN
ANGELAKI@CABERNET.CNS.BCM.EDU
GDEANGELIS@CVS.ROCHESTER.EDU
ALEX@CVS.ROCHESTER.EDU

¹University of Rochester

²Baylor College of Medicine

³Institute of Neuroscience

⁴University of Geneva

Responses of individual neurons in the brain are often significantly predictive of behavior in discrimination tasks. This is surprising because the task-relevant sensory information is distributed across many neurons, so one expects any single neuron to make a tiny contribution to the behavioral choice. Past resolutions of this paradox

recognized that correlations strongly restrict the information content of neural populations, so each neuron is predictive of behavior while multiple neurons all predict the same behavior. However, recent theoretical analyses show that, if tuning curves are heterogeneous, as observed in the brain, then broad noise correlations can be readily disentangled from the sensory signal, so a large population can convey a great deal of information. An efficient readout therefore extracts unique information from each neuron, and any single neuron again has an immeasurably small contribution to behavior. Thus the paradox remains unresolved: How is it possible that single neurons substantially predict behavior? We propose that the amount of sensory information in the behavioral readout is severely restricted, and we evaluate two possible causes: optimal readout limited by 'bad' correlations that have been neglected in previous models, or a highly suboptimal readout. These causes have different, observable consequences for how a neuron's response properties are related to choice probability (CP, a common measure of how well a neuron predicts behavior). The former cause yields a simple inverse relationship between CP and a neuron's stimulus discrimination threshold; the latter is dependent instead on the neural tuning. By examining the relationship between CP, neural threshold, and tuning curves, we can theoretically identify how optimal the neural decoding is. We analyze neural responses recorded in the vestibular and cerebellar nuclei (VN/CN), as well as cortical area MSTd, and we find that responses are consistent with a near-optimal readout of the available information.

I-22. Confidence based learning in a perceptual task: how uncertainty and outcome influence choice

Eric DeWitt¹

Andre G Mendonca¹

Adam Kepecs²

Zachary F. Mainen¹

EDEWITT@NEURO.FCHAMPALIMAUD.ORG

ANDRE.MENDONCA@NEURO.FCHAMPALIMAUD.ORG

KEPECS@CSHL.EDU

ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

¹Champalimaud Neuroscience Programme

²Cold Spring Harbor Laboratory

In decision-making tasks subjects learn to use sensory information to choose actions in order to maximize reward. Usually this is considered two separate problems. Decision theory considers the problem of how to infer the state of the world given a sensory signal. Reinforcement learning considers how to estimate future rewards from prior rewards following an action in a given state of the world. Thus, it is typically assumed in perceptual decision-making studies that state-action-reward contingency uncertainty can be ignored while it is typically assumed in reinforcement learning that sensory uncertainty can be ignored. Here, we challenge this separability assumption by showing that in a simple perceptual decision-making task, even after extensive training, a process resembling reinforcement learning contributes to task uncertainty. We studied a task in which binary odor mixtures were associated with different responses according to a categorical boundary and difficulty (uncertainty) was varied by adjusting the distance of the stimuli from that category boundary. Rats were trained to asymptotic performance, around 25,000 trials, to eliminate the effects of task learning. We then fit a trial-by-trial logistic regression model to estimate the influence of prior successes and errors on the current choice. Prior outcomes biased current choices according to the previously chosen side and outcome, with a magnitude that increased with uncertainty of the previous stimulus and the current stimulus. We next fit a delta rule-like model in which reward expectation was modulated by stimulus uncertainty. This model well captured the data when stimulus uncertainty was estimated using a same-trial decision confidence measure. These results suggest that reinforcement learning mechanisms that enable the learning of arbitrary stimulus-outcome associations also contribute to decision uncertainty once the task is well-learned. This suggests that we should integrate three classically separate approaches: statistical decision theory, statistical learning theory and reinforcement learning.

I-23. A model of interactive effects of learning and choice incentive in the striatal dopamine system

Anne Collins
Michael Frank
Brown University

ANNE_COLLINS@BROWN.EDU
MICHAEL_FRANK@BROWN.EDU

The striatum and dopaminergic systems have been strongly implicated in reward-based behavior, with debates focusing on the relative roles of this system in reinforcement learning, motor performance, and incentive motivation. Neural network models implicate the corticostriatal dopamine system at the intersection of all of these functions – not independently, but interactively. Dopaminergic modulations directly influence action selection and choice incentive: the degree to which decisions are based primarily on learned prospective gains vs losses, encoded in D1 and D2 expressing medium spiny neuron populations. Reciprocally, phasic dopamine signals involved in learning progressively modulate synaptic weights and hence activity levels, which in turn influence not only action selection, but also the eligibility for further learning. Although this model captures a variety of findings across species, a formal analysis is lacking, due to its complexity. Here we present a novel algorithmic description dual process model with actor-critic qualities to capture these fundamental interactive properties of the neural implementation, incorporating both incentive and learning effects into a single theoretical framework suitable for formal analysis and quantitative fits. The actor is divided into QG and QN values representing distinct striatal populations which, by virtue of an interactive effect of learning and activity, come to differentially specialize in discriminating values with positive and negative outcomes. The choice rule is a function of the weighted difference between these QG and QN values for each action. Dopamine effects on choice incentive or motivation are captured by modulating the relative gains on the expression of QG vs QN values, while the effects on learning are captured by asymmetry in their learning rates. In contrast to existing algorithmic models, simulations simultaneously capture documented effects of dopamine on both learning and choice incentive across a variety of studies, as well as their interactive effects on motor skill learning.

I-24. Circuit mechanisms underlying cognitive flexibility in rats

Ioana Carcea
Natalya Zaika
Robert C. Froemke
New York University

IOANA.CARCEA@MED.NYU.EDU
NATALYA.ZAIKA@NYUMC.ORG
ROBERT.FROEMKE@MED.NYU.EDU

Learning to adapt to changing environments requires cognitive flexibility. Impairments of this executive function have been documented in a number of psychiatric disorders (e.g. OCD, schizophrenia and autism). Previous studies in primates and rodents indicate that neuronal circuits in the prefrontal cortex (PFC), via their interaction with striatal and neuromodulatory neuronal networks, contribute to cognitive flexibility by selecting and updating behavioral strategies based on available sensory input and past experience. It remains unclear what cellular and circuit mechanisms sustain a flexible representation of meaningful sensory information in the PFC. To address this question, we started by investigating how PFC gains access to sensory information and how it estimates the behavioral significance of various sensory inputs. Our electrophysiological recordings indicate that auditory stimuli evoke responses in a small population of PFC neurons, even in untrained, anesthetized adult animals. Prefrontal acoustic representations are poorly tuned to stimulus features in naïve animals but become sharper after learning to associate auditory stimuli with behavioral meaning. We find that plasticity in the prefrontal cortex may be a substrate for sparsely encoding the significance of acoustic stimuli in a steady state, constant environment. Whether and how this type of representation is capable of supporting behavioral flexibility remains to be determined. The neuromodulator dopamine could contribute to encoding and updating PFC representations. Dopamine is released by neurons of the ventral tegmental area (VTA), a structure signaling reward prediction errors, in PFC and other brain structures. We trained rats on an adaptive behavioral task where they experience both positive and negative reward prediction errors. We find that during this adaptive task, rats consistently adopt

a four-epoch strategy, where they sequentially persevere, suppress, generalize and refine their responses to acoustic signals. Manipulating dopaminergic activity by patterned VTA stimulation alters the sequence of adaptive behavioral responses.

I-25. Activity in mouse pedunculo pontine tegmental nucleus reflects priors for action value

John Thompson
Gidon Felsen

JOHN.A.THOMPSON@UCDENVER.EDU
GIDON.FELSEN@UCDENVER.EDU

University of Colorado, Denver

Accrued neurophysiological evidence in rat, cat, and non-human primate has identified a network of brain areas that contribute to the process of decision making, from acquisition of sensory evidence to response selection (Gold and Shadlen, 2007). A nucleus recently situated within this network, about which little is known, is the pedunculo pontine tegmental nucleus (PPT), which is interconnected with several motor regions (Steckler et al., 1994; Jenkinson et al., 2009). We have performed tetrode recordings in the PPT of mice engaged in an odor-cued forced-choice spatial task. From initial findings we have observed units in the PPT that show direction preference (calculated by ROC analysis) during epochs associated with decision-making, locomotion, and reward prediction. Intriguingly, we have observed a subset of units that exhibit selectivity for the upcoming movement direction preceding delivery of the odor cue, suggesting that the animal is basing its response on information available before that trial's sensory evidence, such as the choices and outcomes of previous trials. This result is surprising given that, in contrast to a free choice task, in each trial of our forced-choice the task the only reliable evidence available to the mouse is the identity of the odor mixture. Nevertheless, applying a Wiener kernel analysis to our behavioral data revealed that choices are indeed significantly affected by the choice and outcome of the previous trial, particularly when the sensory evidence is weak. Further, we found that, preceding odor delivery, a subpopulation of PPT neurons represents the direction of the previous choice, similar to findings in primate cortex and striatum (Barraclough et al., 2004; Lau & Glimcher, 2007; Histed et al., 2009). These neurons may reflect the computed estimate of the prior value of the two reward ports that can be integrated with incoming sensory evidence to select the optimal choice.

I-26. Striatal optogenetic stimulation reveals distinct neural computation for value-based decision making

A. Moses Lee¹
Lunghao Tai¹
Antonello Bonci²
Linda Wilbrecht^{1,3}

MLEE@GALLO.UCSF.EDU
LTAI@GALLO.UCSF.EDU
ANTONELLO.BONCI@NIH.GOV
WILBRECHT@BERKELEY.EDU

¹University of California, San Francisco

²NIDA, NIH

³University of California, Berkeley

In constantly changing environments, animals adaptively select and evaluate actions to achieve their goals. Many theories have suggested a role for the striatum in goal-directed decision-making. However, the heterogeneity of neural correlates related to motivation, value, and motor responses found in these structures have given rise to conflicting accounts of striatal function. These theories are further complicated by the heterogeneity of cell types/subregions in the striatum, which may support different computational roles. Here we used optogenetics to causally demonstrate the impact of activity in distinct striatal populations and subregions on value-based decision making. We showed that transient cell-type specific stimulation in dorsal and ventral striatum make distinct contributions to action selection and outcome evaluation during separate epochs within a decision-making task.

Optogenetic stimulation was directed to D1- or D2-receptor expressing neurons of the dorsomedial striatum and nucleus accumbens core. Stimulation in the dorsal striatum during action selection biased choice behavior in a manner that was consistent with an additive change to the value of an action. Dorsal striatal stimulation was effective in a time window just prior to and during early movement initiation and in opposite directions for D1- and D2-expressing neurons. In contrast, stimulation of accumbens neurons in a time window during outcome evaluation affected choice behavior on the following trial. Stimulation of D1-expressing neurons promoted repetitions or 'staying' at the previous choice while activation of D2-expressing neurons promoted 'switching' away from the previous choice. Both the staying/switching bias was consistent with the predictions of an additive change to the value of the previously chosen action. Together these data support a model in which striatal neurons pool a wide variety of neural inputs across the brain and convert them into a common currency of value to bias current and future choices.

I-27. Encoding of yaw in the presence of roll or pitch: Studies in a fly motion sensitive neuron

Suva Roy¹
Shiva Sinha²
Rob de Ruyter van Steveninck²

SUVAROY@INDIANA.EDU
SRSINHA@INDIANA.EDU
DERUYTER@INDIANA.EDU

¹Indiana University
²Indiana University Bloomington

Movement in a 3D environment typically generates a visual flow field with translation components and rotations about the three principal axes. It is well established that wide field motion sensitive cells in the fly visual system selectively respond to these rotation components. Studies of these cells have generally used single rotation components. However, the encoding of motion about a given axis is likely affected by the presence of distracting motion about other axes. It is not understood how the encoding efficiency of these cells is affected by the presence of multiple rotations or whether the cells encode specific features buried in these more complex motion stimuli. We provide a quantitative analysis of this problem based on recordings of spikes from H1, a wide-field motion sensitive neuron in the fly lobula plate. In the experiment we presented 2D visual stimuli executing one dimensional and multi-dimensional wide field random motion. Since H1 is sensitive to regressive horizontal motion, we consider yaw the primary motion stimulus while either roll or pitch acts as the distractor. Analysis of H1 response patterns shows that the information in the spike train decreases monotonically with increasing distractor variance. However, total entropy remains relatively constant, decreasing only for very high variance. The reduction in information transmission is therefore mainly due to an increase in noise entropy. Reverse correlation of the stimuli captures first and second order features in the combined yaw-distractor stimulus space. The amplitude of the spike triggered average yaw stimulus diminishes with increasing distractor variance, suggesting decreased sensitivity to yaw in the presence of the distractor. Further, increasing distractor variance generates a larger set of significant eigenvectors of the spike triggered covariance matrix. This set splits into a subset with either pure yaw or pure distractor modes, and a subset with mixed modes.

I-28. Beyond Barlow: a Bayesian theory of efficient neural coding

Jonathan W Pillow
II M Park
University of Texas at Austin

PILLOW@MAIL.UTEXAS.EDU
MEMMING@AUSTIN.UTEXAS.EDU

Barlow's "efficient coding hypothesis" considers a neural code to be efficient if neural responses convey maximal information about sensory stimuli given the physiological constraints of the system. This idea has provided a guiding theoretical framework for systems neuroscience and inspired a great many experimental studies on

decorrelation and efficiency in early sensory areas. More recent work has focused on "Bayesian" theories of neural coding, which regard neural responses as encoding a posterior distribution over some stimulus variable of interest. However, there does not appear (as yet) to be any clear connection between these two paradigms. We aim to bridge this gap by introducing a Bayesian theory of efficient coding. The key additional ingredient is a loss function characterizing the desirability of various posterior distributions over stimuli. Barlow's information-theoretic efficient coding arises as a special case of this theory when the loss function is the posterior entropy. However, we will show that there is nothing privileged about information-maximizing codes; they are ideal for some tasks but suboptimal for many others. In particular, decorrelation of sensory inputs, which is optimal under Barlow's framework in the high-SNR regime, may be disadvantageous for loss functions involving squared or higher power errors, independent of the SNR. Bayesian efficient coding substantially enlarges the family of normatively optimal codes and provides a general framework for understanding the principles of sensory encoding. We derive Bayesian efficient codes for a few simple examples, show an application to neural data, and suggest several important avenues for future research.

I-29. Error statistics and error correction: Evidence for multiple coordinate encodings

Todd Hudson
Michael Landy
New York University

HUDSON@CNS.NYU.EDU
LADNY@NYU.EDU

Two movement codes appear to underlie reach planning: an endpoint code based on final desired hand position, and a vector code defining the desired movement distance and direction. Unlike previous work relying on either modified sensory information concerning one code or another, or different movements for studying each code, we examine predictions for these two coding systems while keeping the task, biomechanics, and sensory inputs constant. We have previously shown, based on a learning paradigm that manipulated only the recent history of previous movements, that the error statistics of natural, unperturbed reaches are consistent with these two encodings (Hudson & Landy, 2012a). Here, we extend those findings and ask whether a similar pair of encodings also underlies reach adaptation. We hypothesize that the encoding used to represent a movement error will dictate the nature of the corrective response that is computed. For example, rightward error made for a rightward reach could be encoded as an over-reach, or a rightward bias. This choice will determine whether a subsequent leftward reach will be shortened (in response to the previous 'over-reach') or shifted leftward (in response to the previous 'rightward endpoint error'). In our experiment, false feedback on a frontal computer monitor regarding reach endpoints (for point-to-point reaches on a horizontal tabletop) was used to induce adaptation. Endpoints were perturbed in two ways: either horizontally & vertically (Cartesian), or in distance & direction (polar). All perturbations followed a sinusoidal pattern over trials (Hudson & Landy, 2012b), which allowed us to frequency-tag perturbations along the two dimensions of each coordinate encoding. The resulting reach adaptation followed both types of perturbation, and also generalized to locations of the workspace where no feedback was given, suggesting that both coordinate encodings are used during motor adaptation.

I-30. Input dependence of local field potential spectra: experiment vs theory

Francesca Barbieri¹
Alberto Mazzoni²
Nikos K. Logothetis³
Stefano Panzeri⁴
Nicolas Brunel⁵

BARBIERIFRANCESCO@GMAIL.COM
ALBERTO.MAZZONI@IIT.IT
NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE
STEFANO.PANZERI@IIT.IT
NBRUNEL@GALTON.UCHICAGO.EDU

¹ISI Foundation

²Istituto Italiano di Tecnologia, Genoa

³Max Planck Institut for Biol. Cybernetics

⁴Istituto Italiano di Tecnologia

⁵University of Chicago

How sensory stimuli are encoded in neuronal activity is a major challenge for understanding perception. A prominent effect of sensory stimulation is to elicit oscillations in EEG and Local Field Potential (LFP) recordings over a broad range of frequencies. Belitski et al. recorded LFPs and spiking activity in the primary visual cortex of anaesthetized macaques presented with naturalistic movies and found that the power of the gamma and low-frequency bands of LFP carried largely independent information about visual stimuli, while the information carried by the spiking activity was largely redundant with that carried by the gamma-band LFPs. To understand better how different frequency bands of the LFP are controlled by sensory input, we computed analytically the power spectrum of the LFP of a theoretical model of V1 (a network composed of two populations of neurons - excitatory and inhibitory), subjected to time-dependent external inputs modelling inputs from the LGN, as a function of the parameters characterizing single neurons, synaptic connectivity, as well as parameters characterizing the statistics of external inputs. We then devised an algorithm to fit the data using these analytical results. The data consists in LFP recordings in the visual cortex of macaques, during presentation of a naturalistic movie. This fitting procedure permits to extract the temporal evolution, during the movie presentation, of both the network parameters, such as the excitatory and inhibitory firing rates, and the parameters of the input, such as for example its typical time scales. We found that the average firing rates extracted from the fits correlates significantly with the multi-unit activity. Furthermore we found a significant correlation between the parameters that describe the input and the features of the movie, such as for example the temporal contrast.

I-31. Object-based spectro-temporal analysis of auditory signals

Mikio Aoi^{1,2}

Yoonseob Lim³

Uri Eden³

Timothy Gardner³

MCAOI@BU.EDU

YSLIM@BU.EDU

TZVI@MATH.BU.EDU

TIMOTHYG@BU.EDU

¹Boston University, Department of Math & Stats

²Cognitive Rhythms Collaborative

³Boston University

Many signals found in nature are naturally described by continuous contours in the time-frequency (TF) plane, but standard TF methods disassociate continuous structures into discrete elements of locally stationary processes. This mismatch between natural signals and the representational elements used to describe them makes the construction of parsimonious representations of a signal from a given set of elements problematic. An alternative to frame-based TF representations arises from the a priori assumption that continuous contours in the TF plane best represent the signal. Earlier work [1] has demonstrated that the analytic geometry of the Gabor transform naturally leads to an invertible signal representation consisting of a sparse set of contours in the TF plane. From this observation, an 'object-based' decomposition of the signal can be defined. By combining contour information over analysis parameters (ex. window length), we can find the most structurally stable subset of 'consensus' contours to construct higher-order, but more parsimonious, representations of a non-stationary signal [2]. This object-based approach liberates TF analysis from discrete frames, and provides a computational analog to the central claim of Gestalt psychology: that sensory objects are perceived in their own simplest, most parsimonious forms. In this work, we generalize the contour method [2] to the continuous chirplet transform and show that these contours have an intimate relationship with the maximum likelihood estimators of the instantaneous frequency of the signal. The resulting representations provide higher resolution and more robust spike-triggered averages for birdsong than do conventional spectrogram-based methods or the prior contour method. We conclude with a discussion of the plausibility of similar object-based representations contributing to auditory processing.

I-32. Heterogeneity increases information transmission of neuronal populations

Eric Hunsberger
Matthew Scott
Chris Eliasmith
University of Waterloo

EHUNSBER@UWATERLOO.CA
MSCOTT@UWATERLOO.CA
CELIASMITH@UWATERLOO.CA

Noise, in the form of stochastic fluctuations added to the membrane voltages of neurons in a population, can have a beneficial effect on the information encoding ability of the population; this phenomenon is one type of stochastic resonance. We have found that heterogeneity, in the form of randomly varying firing thresholds among neurons in a population, can also improve the ability of the population to encode an input signal. Specifically, we performed numerical experiments using populations of FitzHugh-Nagumo neurons and leaky integrate-and-fire neurons, and measured the mutual information between the input signal and the decoded output signal. We found that heterogeneity exhibits a similar resonance effect to noise, where a non-zero amount of heterogeneity maximizes the mutual information between the input and output signals, for both neuron models. We also performed numerical experiments examining three common mechanisms that allow both noise and heterogeneity to increase information transmission in neuronal populations: 1) both temporally desynchronize neurons in the population, 2) both decrease the response time of a population to a sudden change in input signal, and 3) both linearize the response of the population to a stimulus. The main contribution of this research is that it demonstrates that heterogeneity can play an important role in neuronal population coding, and examines the mechanisms it uses to fill this role. Heterogeneity has frequently been overlooked in neuroscience; our study explains why it is an important feature of neural systems, and why it should not be overlooked when modeling such systems.

I-33. Constraining cortical population coding of visual motion in area MT by smooth pursuit behavior

Stephanie E Palmer
Leslie Osborne
University of Chicago

SEPALMER@UCHICAGO.EDU
OSBORNE@UCHICAGO.EDU

Visual sensation arises from the activity of large populations of neurons, however sampling that population activity can be difficult deep inside the brain. One approach to the analysis of population coding is to simulate ensemble data based on real neural responses in order to fill in the sampling gaps. To test which features of real neural responses are important for population coding, one needs a performance measure. We use the pursuit system as a model to test theories of cortical population coding of visual motion. Pursuit behavior offers a valuable performance metric for population models, since eye movement can be well-characterized, the neural circuitry is well-known, and pursuit initiation is tightly coupled to responses in area MT. The visual inputs that drive pursuit are formed over 100-200ms, a time period in which MT firing rates are both highly dynamic and diverse. MT neural thresholds for discriminating motion direction are about 10 times larger than those for pursuit and perception, indicating that target motion is estimated from the joint activity of the cortical population. We use data-driven simulations that preserve the heterogeneity of feature selectivity, dynamics, and temporal spike count correlations to compute the time course of population information about motion direction. Specifically, we compute the Cramer-Rao bound on the variance of target direction estimates in comparison to pursuit behavior. We find that the precision of motion estimates is influenced by the degree to which the simulation preserves the natural features of MT responses. For example, preserving the natural heterogeneity in neural response dynamics improves direction estimation by a factor of two. Simulated populations of size 200-300 reach direction discrimination thresholds that are consistent with behavioral data. We quantify the impact of MT response features as well as putative correlations on population information, informing our understanding of the sensory code.

I-34. The role of Inhibitory STDP in disynaptic feedforward circuits

Florence Kleberg¹

Matthieu Gilson²

Tomoki Fukai¹

KLEBERG@BRAIN.RIKEN.JP

GILSON@BRAIN.RIKEN.JP

TFUKAI@RIKEN.JP

¹RIKEN BSI

²RIKEN Brain Science Institute

Spike trains with temporal correlations are believed to be related to sensory and cognitive processing[1]. Therefore, extracting correlation information from input spike trains is an important computation for neurons. Learning rules such as Spike-Timing Dependent Plasticity (STDP) provide a mechanism by which a neuron may become selective to one source of correlation in spike times [2]. Recently, the functional role of STDP on inhibitory synapses has been discussed in single neuron models[3], embedded in a network [4]. There is no consistent evidence for inhibitory STDP (iSTDP) and considering the large number of inhibitory neuron types and connectivity patterns, theoretical analyses are necessary to provide possible paths and functions to investigate experimentally. We studied the role of iSTDP, in particular the learning window type and the statistical structure of spike trains, in correlation tuning of Leaky Integrate-and-Fire (LIF) neurons in a disynaptic feedforward circuit. This circuit is ubiquitous in the brain, such as in the cortex and hippocampus [5]. We go beyond the single neuron and also model the circuit explicitly, to our knowledge a novel approach. There is an optimal iSTDP learning window for inhibitory inputs onto the postsynaptic neuron when detailed balance is to be maintained between excitation and inhibition. In the disynaptic circuit, we show that the inhibitory neurons can be recruited by excitatory STDP to represent one of two equally correlated sources. Additionally, the correlation structure of the incoming correlated spike trains can select an optimal delay of inhibitory inputs. Our findings show the function of iSTDP in maintaining detailed E-I balance and selecting inhibitory pathways with different delays. [1]Riehle et al., (1997) *Science* 278:1950-1953 [2]Gilson and Fukai (2011) *Plos One* 6:e2533 [3]Vogels et al., (2011) *Science* 334:1569:1573 [4]Luz and Shamir (2012) *Plos Comp. Biol.* 8:e1002334 [5]Buszaki (1984) *Prog.Neurobiol.* 22:131-153

I-35. Rapid development of feed forward inhibition drives emergence of the alert cortical state

Matthew Colonnese

George Washington University

COLONNESE@GWU.EDU

The generation of appropriate cortical network states is a key regulator of perception and plasticity, but their role during development is poorly characterized. We have shown that human preterm infants and pre-eye opening rats undergo a rapid maturation of network states just before the onset of visual experience. This change results in a massive down regulation of visual responses, shifting them from all-or-none oscillatory bursts to graded responses capable of processing visual input. Here we test the hypothesis that this maturation of cortical network state is the rapid emergence of cortical activation, or 'active', states. We use in vivo current clamp and polytrode recordings in visual cortex of awake, head-fixed neonatal and infant rats. We find that cortical activation, defined by persistent membrane depolarization during waking, emerges suddenly 1-2 days before eye-opening. The amplitude of activation remained constant between emergence and adulthood, though stability and duration gradually increased. This switch in network properties was responsible for the down-regulation of visual responses as light began to evoke active states rather than supra-threshold plateau potentials observed before the switch. Reducing GABAA currents just after the switch eliminates activation and reverts activity to immature patterns. Measurement of the timing and amplitude of inhibitory and excitatory currents by voltage clamp showed the rapid development of fast feed-forward inhibition at this time. In total we have identified, for the first time, a specific role for changes in inhibitory circuitry in the developmental regulation of cortical activity. This change effectively divides visual cortex development into two clear phases—an early pre-visual period concurrent with spontaneous retinal waves and the establishment of retinal topography, and a late period linked to the onset of pattern vision, visual exploration and the onset of experience-dependent plasticity—that each have unique computational characteristics.

I-36. Laminar differences in receptive fields measured by reverse correlation of intracellular recordings.

Alejandro Ramirez
Eftychios A. Pnevmatikakis
Josh Merel
Ken Miller
Liam Paninski
Randy Bruno
Columbia University

AR2676@COLUMBIA.EDU
EFTYCHIOS@STAT.COLUMBIA.EDU
JSM2183@COLUMBIA.EDU
KEN@NEUROTHEORY.COLUMBIA.EDU
LIAM@STAT.COLUMBIA.EDU
RANDYBRUNO@COLUMBIA.EDU

Our knowledge of receptive fields and sensory transformations in rodent barrel cortex (S1) lags behind other sensory systems. Firing rates of neurons in S1 can be low, making reverse correlation of high-dimensional stimuli challenging. Additionally, most researchers rely on simple single-whisker laboratory stimuli for receptive field mapping, which are neither ethological nor capable of revealing spatiotemporal complexity. Here we use a novel multi-whisker stimulator system that moves 9 whiskers independently in arbitrary directions, exploring a vastly larger stimulus space than conventionally examined. By recording intracellularly rather than extracellularly, we can additionally access information available in the subthreshold response to calculate receptive fields even for neurons with little or no spiking activity. After exploring a number of stimulus-response models, including conventional Linear-Nonlinear models as well as quadratic models, we found that a filtered input nonlinearity model (of the form discussed in Ahrens et al, 2008) provided an effective and parsimonious representation of the responses. In this model, the whisker deflections are mapped through a static nonlinearity that re-represents the whisker movements binned into an 8-directional space, before being temporally filtered, weighted across whiskers, and summed to predict the voltage response. The static nonlinearity, temporal filters, and linear weights are all estimated simultaneously using rank-penalized regression methods. Our model is able to predict neural responses to novel stimuli with a correlation coefficient as high as 0.84. Furthermore, through repeated presentations of identical stimuli, we show that our model captures ~90% of the predictable variance (Sahani and Linden 2003), suggesting that the main nonlinearities are spike-threshold rather than network nonlinearities. Analysis of the spatiotemporal receptive fields across layers and cell-types reveals the emergence of unique spatial and temporal features encoded in the supra- and infra-granular layers, and serves as a useful comparison to similar studies from the visual and auditory systems.

I-37. T-type calcium channels promote predictive homeostasis in thalamocortical neurons of LGN

Su Hong
Haram Kim
Sora Yun
Christopher Fiorillo
KAIST

HONGSU@KAIST.AC.KR
YU_K519@KAIST.AC.KR
YSL0142@KAIST.AC.KR
CHRISTOPHER.D.FIORILLO@GMAIL.COM

The efficient coding hypothesis (Barlow, 1961) proposes that to maximize the information in its output, a neuron should respond only when it fails to predict its input (prediction error). Application of these principles to the cellular and molecular levels has suggested how a diverse set of synapses and voltage-gated ion channels could contribute to the prediction of a neuron's sensory input (Fiorillo, 2008). Theory suggests that mechanisms of 'predictive homeostasis' should work to maintain a nearly constant and linear input-output (I-O) relationship, as supported by experimental evidence from lateral geniculate nucleus (LGN) (e.g. Dan et al., 1996). However, the theory appears to be at odds with evidence that T-type calcium channels in LGN cause 'bursts' of two or more output spikes in response to a single retinal input spike ($I-O < 1:2$) when a neuron is sufficiently hyperpolarized to deactivate T-type channels. We formulated the hypothesis that T-type channels actually help to maintain a homeostatic I-O relation under natural conditions. To test this, we attempted to mimic in brain slices the synaptic

inhibition that has been observed to hyperpolarize LGN neurons for a few hundred milliseconds in vivo during naturalistic visual stimulation (Wang et al., 2007), in which two retinal input spikes typically sum temporally to cause one output spike (Sincich et al., 2007). We evoked real or artificial retinal EPSPs from a potential of -65 mV to recreate this natural 2:1 I-O relation. We found that following hyperpolarization to -80 mV for up to 800 ms by a real or artificial chloride conductance, the I-O relation gradually recovered from less than 4:1 to nearly 2:1. This recovery was not observed after T-type channels were blocked with nickel. Thus, under natural conditions, T-type channels appear not to generate bursts, but instead contribute to predictive homeostasis.

I-38. Neural implementations of motion detection in the fly medulla: insights from connectivity and theory

Arjun Bharioke^{1,2}
 Shin-ya Takemura^{3,4}
 Shiv Vitaladevuni^{3,4}
 Richard Fetter^{3,4}
 Zhiyuan Lu^{3,4}
 Stephen Plaza^{3,4}
 Louis Scheffer^{3,4}
 Ian Meinertzhagen⁵
 Dmitri Chklovskii^{3,4}

BHARIOKEA@JANELIA.HHMI.ORG
 TAKEMURAS@JANELIA.HHMI.ORG
 VITALADEVUNIS@JANELIA.HHMI.ORG
 FETTER@JANELIA.HHMI.ORG
 LUZ@JANELIA.HHMI.ORG
 PLAZAS@JANELIA.HHMI.ORG
 SCHEFFERL@JANELIA.HHMI.ORG
 IAM@DAL.CA
 MITYA@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus, HHMI

²University of Cambridge

³Janelia Farm Research Campus

⁴HHMI

⁵Dalhousie University

Visual motion detection is an important computation for most animals. In insects, the most widely accepted model of motion detection, the Elementary Motion Detector (EMD) (Hassenstein and Reichardt, 1956), uses multiplication to correlate light intensity between two points displaced in space and time. Yet, there has never been 'a smoking gun' identifying specific neurons with components of the EMD. Here, we combined newly obtained information about the wiring diagram with novel theoretical models and published physiological measurements, behavioral experiments, and genetic manipulations to synthesize a plausible model of motion detection. We took advantage of the stereotypy of the fly visual system, where most neuron types have been identified and named. Genetic silencing experiments (Rister et al, 2007; Joesch et al, 2010; Clark et al, 2011) determined that the cells, L1 and L2, T4 and T5, were necessary for motion detection. Utilizing the newly assembled connectivity matrix, we identified neurons linking these neurons, specifically Mi1 and Tm3, bridging L1 and T4. We found that multiple Mi1s and Tm3s provide input to each T4, thus defining two distinct anatomical components of the receptive field. The centers of mass of these receptive field components, in visual space, are distinct. This suggests that these cells may represent the two spatially distinct inputs within a correlation based motion detector. Because of the difficulty of implementing multiplication in neurons we modeled motion detection using the so-called Rectification Motion Detector (RMD) (Mizunami, 1990) which can be derived from the general sign-rules of motion detection. We could find only one implementation that satisfies all the constraints. This allows us to construct testable predictions: particular subsets of T4, T5 neurons should be activated by inputs of particular contrast. Also, when the assumptions are relaxed, we identified experiments useful in distinguishing models, e.g. looking for reverse phi illusions from single ommatidial stimulation.

I-39. Evidence for a nonlinear coupling between firing threshold and sub-threshold membrane potential

Skander Mensi^{1,2}
 Christian Pozzorini³
 Olivier Hagens^{1,4}
 Wulfram Gerstner³

SKANDER.MENSI@EPFL.CH
 CHRISTIAN.POZZORINI@EPFL.CH
 OLIVIER.HAGENS@EPFL.CH
 WULFRAM.GERSTNER@EPFL.CH

¹EPFL

²BMI

³EPFL, BMI

⁴SV BMI LNMC

In the field of computational neuroscience, it is of crucial importance to dispose of simplified spiking models that capture the behavior observed in single neurons. In the last years, many studies have demonstrated the ability of generalized integrate-and-fire models (GIF) to predict the occurrence of individual spikes with a millisecond precision. In a rate-based framework, individual neurons are often characterized by f-I curves: functions that map constant inputs onto steady state output rates. While on one hand the experimentally observed f-I curves of fast spiking interneurons are in good agreement with the ones predicted by GIF models, the same is not true for excitatory pyramidal neurons. In particular, the firing activity of pyramidal neurons typically saturates at relatively small rates (30-50 Hz) and, in presence of noise, pyramidal neurons maintain their sensitivity to fast fluctuations even at large baseline currents. To capture these features, we propose a model in which spike-dependent adaptation mechanism is complemented by a subthreshold one. This mechanism implements a nonlinear coupling between the firing threshold and the membrane potential. Importantly, all the model parameters, including the timescale and the functional shape of the nonlinear coupling, are not assumed a priori but are extracted from in-vitro recordings using a new convex optimization procedure. Our results demonstrate that the firing threshold and the subthreshold membrane potential are indeed nonlinearly coupled. This mechanism, consistent with subthreshold Na-channels inactivation, operates on a relatively short timescale (3-4 ms) and makes the firing threshold depend on the speed at which the threshold is approached. With this new model, the accuracy in spike-time prediction is improved. More importantly, the precise shape of the nonlinear coupling extracted from the experimental data, accounts for both the saturation and the noise sensitivity that characterize f-I curves of pyramidal neurons.

I-40. Spatial distribution and efficacy of thalamocortical synapses onto layer 4 excitatory neurons

Carl Schoonover
 Juan Tapia
 Verena Schilling
 Verena Wimmer
 Richard Blazeski
 Carol Mason
 Randy Bruno

CESCHOONOVER@GMAIL.COM
 JCTAPIA2007@GMAIL.COM
 RB2604@COLUMBIA.EDU
 VERENA.WIMMER@ME.COM
 RB21@COLUMBIA.EDU
 CAM4@COLUMBIA.EDU
 RANDYBRUNO@COLUMBIA.EDU

Columbia University

In the absence of synaptic depression, individual thalamocortical (TC) synaptic connections onto excitatory layer 4 (L4) neurons in rat somatosensory cortex are significantly more efficacious than corticocortical (CC) connections when measured in vitro. This discrepancy in relative synaptic strength been proposed as a reason why activity in thalamus governs activity in primary sensory cortex. One possible mechanism for this difference is a bias in how TC and CC synapses are spatially distributed along L4 dendrites. To test this hypothesis we developed a high-throughput light microscopy method to map the locations of putative TC synapses across a complete dendritic arbor, in conjunction with a correlative light and electron microscopy strategy to verify whether the putative

contacts represent actual synapses. As hypothesized, we found that TC synapses are on average more proximal to the soma than CC synapses. In order to investigate whether TC synapses are less attenuated by the passive cable properties of the dendrite, we built a compartmental model from a detailed volumetric reconstruction of a L4 excitatory neuron whose TC synapses were mapped. Model parameters were constrained by fitting physiological recordings during current injection. In the model, CC inputs were found to be indistinguishable from TC inputs on average, assuming equal conductance parameters. We verified this prediction by measuring putative unitary TC and CC EPSPs in vivo, and found no significant difference in efficacy between these two classes of inputs. Therefore, despite the relative proximity to the soma of TC versus CC synapses, their spatial distributions cannot explain the greater efficacy of TC synapses observed in vitro, given the passive membrane properties of the cell. An alternative mechanism must account for the strong influence of thalamic activity on primary sensory cortex.

I-41. A plasticity mechanism for the robustness of the oculomotor integrator

Pedro Goncalves
University College, London

PEDROG@GATSBY.UCL.AC.UK

The oculomotor integrator in the hindbrain transforms eye movement inputs into position signals to maintain stable eye fixations after saccades. Experimentally, it has been shown that neurons in the oculomotor integrator have firing rates that persist at a continuum of levels, where each level corresponds to a particular fixation. These results suggest that the oculomotor integrator features a continuum of stable states generated by a continuous attractor network (Seung1996). However, typical continuous attractor models require fine-tuning of the synaptic parameters with a precision below 1 %, and it remains elusive how the oculomotor integrator can fine-tune its synapses. In addition, these models are sensitive to noise, since even small amounts of noise are integrated by the system. Previous modeling work solved both issues by making neurons or dendrites bistable (Koulakov2002,Goldman2003), in agreement with the hysteresis found in the tuning curves of integrator neurons. However, bistability has never been observed in integrator neurons or dendrites. Here we propose to solve the fine-tuning problem by assuming the presence of synaptic plasticity in integrator neurons, a well-established biophysical process. The plasticity rule, based on (Moreau2003), leads to a network that is robust to perturbations in its parameters up to 10 %, making the overall model more robust than fine-tuned network models. In addition, the model is less sensitive to noise, since the plasticity rule filters out small amounts of noise. Interestingly, this integrator model shows hysteresis in the tuning curves as observed in the data. Therefore, this result supports our model and suggests that hysteresis is the footprint of the self-tuning mechanism in action. Finally, we propose an optogenetic experiment to test the existence of such mechanism and the extent of its flexibility.

I-42. Robust formation of continuous attractor networks

Daniel Robles-Llana
Ila Fiete
University of Texas at Austin

DANIEL.ROBLESLLANA@GMAIL.COM
ILAFIETE@MAIL.CLM.UTEXAS.EDU

Continuous attractor networks, or neural networks with a continuum of persistent activity states, are a theoretical cornerstone for understanding how the brain might temporarily store, manipulate, and integrate analog variables. Yet little is understood about how these complex networks could form. Existing models of their development share a theme: an unstructured network is driven by tuned feed-forward inputs, and an associative plasticity rule recurrently wires together neurons with similar input tuning, stabilizing bump-like activity. To work, these models rely critically on: 1) noise-free dynamics, 2) artificial suppression or normalization of recurrent weights during network formation, and 3) perfectly uniform exploration of the input training variable. These biologically untenable assumptions are made to tame weight instabilities that can occur when activity-dependent plasticity is applied to recurrent networks (1-2); and to generate an invariant weight profile across neurons in the mature

network, necessary for the network to maintain a continuous set of variables (3). Existing models therefore fail to explain how continuous attractor networks can form in the brain. We show, in recurrent networks of stochastically spiking neurons, that inclusion of specific single-neuron mechanisms can enable the brain to robustly assemble continuous attractor networks. First, spike-frequency adaptation nudges the system away from ‘sticking’ in local fixed points, a major cause of feedback instability from activity-dependent plasticity. Second, homeostatic scaling in neurons enhances the uniformity of activation across neurons, consistent with the goal of learning a continuous attractor. Third, a form of heterosynaptic competition within each neuron that penalizes all synapses by a fraction of their weight when the summed weight of synaptic inputs exceeds a threshold, counters the need for uniform visitation during development. Our results suggest that the interplay between specific adaptive single-neuron mechanisms and activity-dependent synaptic plasticity can be crucial for robustly organizing the complicated circuitry of recurrent cortical networks.

I-43. Connectomic constraints on computation in networks of spiking neurons

Venkatakishnan Ramaswamy¹
Arunava Banerjee²

VENKAT.RAMASWAMY@MAIL.HUJI.AC.IL
ARUNAVA@CISE.UFL.EDU

¹The Hebrew University of Jerusalem

²University of Florida

Several efforts are currently underway to decipher the connectome of a variety of organisms. Ascertaining the physiological properties of all the neurons in these connectomes, however, is out of the scope of such projects. It is therefore unclear to what extent knowledge of the connectome alone will advance our understanding of computation occurring in these neural circuits. We consider the question of how, if at all, the wiring diagram of neurons imposes constraints on computation when we cannot assume detailed information on the physiological response properties of the neurons. We call such constraints, that arise by virtue of the connectome, connectomic constraints on computation. For feedforward networks equipped with neurons that obey a deterministic spiking neuron model which satisfies a small number of properties, we ask how connectomic constraints restrict the computations they might be able to perform. One of our results shows that all networks whose architectures share a certain graph-theoretic property also share in their inability in effecting a particular class of computations. This suggests that connectomic constraints are crucial for computation; merely having a network with a large number of neurons may not endow it with the ability to effect a desired computation. In contrast, we have also proved that with the limited set of properties assumed for our single neurons, there are limits to the constraints imposed by network structure. More precisely, for certain classes of architectures, we must have more detailed information on single neuron properties, before we can prove that there exist computations that the networks cannot perform. Thus, our theory suggests that while connectomic constraints restrict the computational ability of certain classes of network architectures, we may require more information on the properties of neurons in the network, before we can prove such results for other classes of networks.

I-44. Dynamics of random clustered networks

Merav Stern^{1,2}
Haim Sompolinsky³
Larry Abbott¹

MERAV.STERN@MAIL.HUJI.AC.IL
HAIM@FIZ.HUJI.AC.IL
LFA2103@COLUMBIA.EDU

¹Columbia University

²Hebrew University

³The Hebrew University

Connections in many neural circuits span multiple spatial scales with abundant local interactions and sparser

interactions between distal populations. We study the effect of this organization on network dynamics using randomly connected networks of firing-rate units that include self-coupling terms. Individual units in these model networks are interpreted as clusters of neurons, and the self-coupling reflects the interactions among local sub-circuits. Connections between the N clusters of the networks we study are random, balanced between excitation and inhibition, and of order $1/\sqrt{N}$ in strength. The self-interactions are of order 1. We explore the dynamics of these networks for different values of s , the strength of the self-interaction representing local interactions, and g , the strength of the connections between clusters representing distal connectivity. When s is either negative (inhibitory), or positive (excitatory) but less than 1, two regimes arise. For $g < -s$, the network decays to zero activity. For larger g ($g > -s$), the network is chaotic. Using the dynamic mean-field approach, we compute the average autocorrelations in this model as a function of s and g . When s is slightly greater than 1 and exceeds a critical value ($s > sC(g) > 1$), chaos exists only transiently and, ultimately, the network reaches a non-zero fixed point. In this regime the time to settle to a fixed point grows exponentially with network size. Using a static mean-field approach, we compute the critical self-coupling $sC(g)$. With strong self-coupling ($s > 1$) this network combines computational features of fixed-point and chaotic dynamics. As a result, self-coupled networks could be used to combine aspects of short- and long-term memory, with fluctuations used as a source of internal dynamic activity over shorter time intervals and quasi-steady-state activity representing longer-lived 'behavioral' states.

I-45. NeuroElectro: a community database on the electrophysiological diversity of mammalian neuron types

Shreejoy J Tripathy¹
Richard Gerkin¹
Judy Savitskaya^{1,2}
Nathaniel Urban¹

STRIPATHY@CMU.EDU
RICKG@CMU.EDU
JUDYSAVIT@GMAIL.COM
NURBAN@CMU.EDU

¹Carnegie Mellon University

²University of Cambridge

Brains achieve efficient function through implementing a division of labor, in which different neurons serve distinct computational roles. One striking way in which neuron types differ is in their electrophysiology properties. These properties arise through expression of combinations of ion channels that collectively define the computations that a neuron performs on its inputs and its role within its larger circuit. Though the electrophysiology of many neuron types has been previously characterized, these data exist across thousands of journal articles, making cross-study neuron-to-neuron comparisons difficult. Here, we describe NeuroElectro, a public database where physiological properties for the majority of mammalian neuron types has been compiled through semi-automated literature text-mining. The corresponding web application, at neuroelectro.org, provides a rich dynamic interface for visualizing and comparing physiological information across neuron types, conveniently linking extracted data back to its primary reference. Mining the database content, we show that there exists but 3 or 4 major neuron classes in terms of electrophysiological properties, which separate largely based on neurotransmitter released and cell size. As an example of how this resource can help answer fundamental questions in neuroscience, we integrate NeuroElectro with neuronal gene expression data from the Allen Institute for Brain Sciences. We show that simple statistical models (penalized linear regression) can accurately predict features of a neuron's electrophysiological phenotype given information of its gene expression alone. We further investigate these models to ask which genes, of the 20K in the genome, are most predictive of neuron physiology. We find that while ion channel-related genes provide significant predictive power, the most predictive gene classes surprisingly correspond to G-proteins and transcription factors, suggesting the involvement of hundreds of diverse genes in regulating a neuron's computational function.

I-46. Traveling waves of the hippocampal theta rhythm encode rat position

Gautam Agarwal^{1,2}
 Ian Stevenson³
 Kenji Mizuseki⁴
 Gyorgy Buzsaki⁵
 Friedrich Sommer^{1,2}

GAGARWAL@BERKELEY.EDU
 I-STEVENSON@BERKELEY.EDU
 KENJI.MIZUSEKI@GMAIL.COM
 GYORGY.BUZSAKI@NYUMC.ORG
 FSOMMER@BERKELEY.EDU

¹Redwood Center for Theoretical Neuroscience

²University of California, Berkeley

³University of California Berkeley

⁴Allen Institute

⁵New York University

The theta rhythm is an ~8 Hz oscillation in the hippocampus that mirrors the timing and coordination of large groups of neurons. Its structure varies richly in both space and time. We seek to understand these variations in terms of the sequence of activity that unfolds as a rat runs along a linear track. Our data consists of multi-electrode recordings from layer CA1 of behaving rats. Exploratory analysis revealed that the theta rhythm exhibits a time-varying phase gradient along the axis defined by the apical dendrites of CA1 pyramidal cells. To identify putative sources responsible for this variation, we perform spatial ICA on the analytic (complex-valued) representation of the theta-band oscillation. This analysis reveals a population of sparse components, each with a characteristic spatial amplitude-phase relationship representing a traveling wave that propagates across the electrode array. We find that many of these components are activated in a place- and direction-selective manner; as a rat runs down the track, the components transiently activate in a specific sequence. Together, the set of 'place components' tiles the entire track. This observation is closely related to the known response properties of CA1 pyramidal cells, which also activate in a place-specific manner. However, unlike place cells, the sparse components in the theta band tile the track more uniformly, manifest across the entire electrode array, and are linked to unique cross-frequency dynamics, suggesting that they arise from a mechanistically distinct source. The LFP is commonly considered to be a relatively impoverished signal conveying only general information about behavioral state. In contrast, we find that the multi-electrode LFP can provide a rich behavioral readout. Our analysis approach may also be relevant for identifying the features encoded by other brain structures with prominent oscillations, such as the motor cortex and olfactory bulb.

I-47. Are aperiodic 1D grid cell responses consistent with low-dimensional continuous attractor dynamics?

KiJung Yoon¹
 Daniel Robles-Llana¹
 Amina Kinkhabwala²
 David Tank²
 Ila Fiete¹

KIJUNG.YOON@UTEXAS.EDU
 DANIEL.ROBLESLLANA@GMAIL.COM
 AKINKHAB@PRINCETON.EDU
 DWTANK@PRINCETON.EDU
 ILAFIETE@MAIL.CLM.UTEXAS.EDU

¹University of Texas at Austin

²Princeton University

Since the discovery of the striking activity of grid cells, the question of mechanism has received intense attention. One of two dominant models is based on 2D continuous attractor dynamics in recurrent networks. Briefly, lateral connections stabilize a state with a triangular lattice activity pattern in the population, and all its translations. This model is fully consistent with the rate dynamics of grid cells in 2D enclosures and has made many successful predictions. However, the response of cells along 1D tracks remains a confound and possible challenge. In 1D, grid cells fire at multiple locations, but the pattern is not periodic. Here we examine whether the 1D response patterns are consistent with continuous attractor dynamics, by analyzing multiple simultaneously recorded grid cells, with responses elicited in both 2D and 1D environments. First, we show that aperiodic responses are not

inconsistent with attractor dynamics: while attractor dynamics force cells to maintain fixed response relationships to each other, they do not dictate how network states are mapped to the external represented variable. This mapping may be continually varied or reset, e.g. by external landmarks. Second, we examine the stability of cell-cell relationships in 1D, even as individual cells exhibit drifts in the locations of fields over traversals of a track, showing that cell-cell response relationships are better preserved than the responses of individual cells. Third, we examine whether the 1D response is quasi-periodic, generated as a slice through a periodic 2D pattern, or whether resetting of fields by the environment is an important source of aperiodicity in the 1D response, and show evidence for the latter. Our results suggest that, independent of the spatial mapping between 1D and 2D, and despite the disparity in 1D and 2D-responses, the same low-dimensional dynamics observed in grid cells in 2D may underlie their 1D-responses.

I-48. A single computational mechanism for both stability and flexibility in vocal error correction

Samuel Sober
Conor Kelly
Emory University

SAMUEL.J.SOBER@EMORY.EDU
CWKELLY@EMORY.EDU

The brain uses sensory feedback to correct behavioral errors, and larger errors by definition require greater corrections. However, in several systems larger errors drive learning less effectively than smaller errors. Limiting motor changes in response to large sensory errors is reasonable, since sensory signals are inherently noisy. A successful control strategy must therefore use feedback to correct errors while disregarding aberrant sensory signals that would lead to maladaptive corrections. Our prior work has used online manipulations of auditory feedback to demonstrate that in adult songbirds, vocal error correction is inversely proportional to error size, with smaller errors corrected robustly but larger errors causing minimal vocal changes. These findings in adults, however, introduce an apparent paradox. If vocal learning is smaller in response to larger errors, and animals commit very large errors when initially learning to vocalize, then how does a young bird ever learn to sing? More generally, if error correction is an inverse function of error size, how can complex behaviors ever be acquired? Here, we propose a computational mechanism that can account for both the stability of adult behavior during large errors and the flexibility of behavior earlier in life. In adult songbirds, the extent of learning is well-predicted by the overlap between the prior distribution of sensory feedback and the distribution experienced during sensory perturbations. We therefore hypothesize that songbirds weight the probability of sensory feedback when computing how much to modify song during error correction. We test this hypothesis by quantifying error correction in younger birds, where vocal variability is greater than in older adults. We find that for the same error, learning is much greater in younger animals, but that this apparent difference can be explained by a single probabilistic weighting strategy. These findings suggest that throughout life, the statistics of prior experience constrain learning.

I-49. Optimal foraging and multiscale representation by hippocampal place cells

Jason Prentice¹
Vijay Balasubramanian²
¹Princeton University
²University of Pennsylvania

JASONSP@PRINCETON.EDU
VIJAY@PHYSICS.UPENN.EDU

The size of place fields of hippocampal place cells varies systematically along the dorsoventral axis, dorsal place cells having smaller fields than ventral (Kjelstrup et. al. 2008). Moreover, the phase of the theta oscillation varies coherently along the dorsoventral axis (Lubenov and Siapas, 2009), so that place cells representing different spatial scales may potentially coordinate their activity temporally. It is not clear of what benefit is this multiscale

organization – a coarse representation can only lose spatial information relative to a finer representation. We demonstrate that the multiscale structure arises naturally in a model of place-cell-based computation which generates optimal trajectories in foraging and goal-directed navigation tasks. The optimal path in the presence of a position-dependent cost or reward is calculated mathematically in a control theory framework, and we show that the exact solution to the equations can be computed by a simple neural network. The network involves two populations of neurons; a set of "value cells," computing expected reward, together with a set of place cells whose firing is modulated both by the animal's position and by the outputs of the value cells. The place cell activity is output to a simple population vector readout to generate the optimal step direction. This path-planning problem requires looking ahead to future positions and estimating associated rewards in order to guide the current step. In our model, larger-scale place fields facilitate this "looking ahead" to later times. The model thus requires place fields of many scales, giving a novel explanation for the topographic layout of place field scales in the hippocampus. The model also explains the observed theta phase offset between scales as a mechanism for ensuring that information about reward at each future time is coupled to the appropriate place field scale. The proposed mechanism can be subjected to specific experimental tests.

I-50. Slow time-scales in songbird neural sequence generation.?

Jeffrey Markowitz
 Gregory Guitchounts
 Timothy Gardner
 Boston University

JMARKOW@CNS.BU.EDU
 GG42@BU.EDU
 TIMOTHYG@BU.EDU

Synfire theories demonstrate that redundant, feedforward chains of neurons can stably produce temporal sequences with high precision. The fundamental time-scale of the synfire chain is the time-scale of an individual link in the chain. In songbirds, sparse, stereotyped bursts of projection neurons in nucleus HVC resemble a synfire chain built from 5ms units. Here we present evidence that firing times of projection neurons and interneurons in HVC are grouped together, but over time-scales significantly slower than the time-scale of individual neural units. This synchrony is evident both within HVC in a single hemisphere, and across hemispheres. In parallel, long term recordings with minimally invasive electrodes reveal that the precise firing patterns of interneurons are stable for as long as they have been observed—weeks to months. We consider theoretically how slow time-scale synchronous patterns in HVC may stabilize sequence production based on topological chains. In addition to this mesoscopic time-scale in HVC neural activity, we demonstrate that some bird songs are shaped by long-range correlations in syntax. Recent behavioral work in Bengalese finches suggests that song contains non adjacent dependencies between syllables. We apply prediction suffix trees to analyze the syntax of a complex singer, the canary. This analysis reveals that for canaries, the decision about what to sing next depends on choices made up to ten seconds earlier in song. Taken together, these results indicate that birdsong sequence generation is governed by stereotyped dynamics in wide a range of time-scales—timescales that must be integrated into any parsimonious dynamical or statistical model for song.

I-51. Criterion dynamics in the ferret optimise decision bias during an auditory detection task

Robert W Mill
 Christian Sumner
 MRC Institute of Hearing Research

ROBERTM@IHR.MRC.AC.UK
 CHRIS@IHR.MRC.AC.UK

The ability accurately to detect and discriminate stimuli despite uncertain sensory evidence is essential to an organism's survival. Signal detection theory (SDT) enables one to quantify this ability in terms of sensitivity (d') and bias. To account for bias, SDT posits a criterion, which is applied to a random decision variable

combining stimulus level and internal noise. Many classical psychoacoustic experiments assume for convenience that a static criterion is maintained during a trial sequence. However, experiments examining this assumption demonstrate that the criterion is influenced by the history of stimulus levels and responses. Here we report an animal behavioural experiment designed to highlight criterion dynamics and present a model that accounts for the results obtained. An auditory detection task required two trained ferrets to indicate whether a tone was present in broadband noise. Trials were arranged into blocks with alternating signal level statistics. Higher levels appeared in oddly-numbered blocks. Criterion dynamics were manifested on at least two time scales. The hit and false alarm probabilities, conditioned on the position of a trial within a block, revealed criterion variation with an exponential like rise and decay profile, which tracked the optimal decision criterion. Responses were also affected by the outcome of the immediately preceding trial. A SDT model with Markov criterion dynamics was fitted to the observed sequences of stimuli and decisions using maximum likelihood methods. The model comprises a trial-by-trial update rule applied to a single, hidden state variable, and captures behavioural changes on short and long time scales. This reveals how animals, when faced with changing stimulus statistics, dynamically optimise their behaviour to maximise reward. This simple theoretically-tractable formulation leads to a model capable of generating predictions that inform the search for neural correlates of perceptual decision behaviour.

I-52. A multiplicative learning rule for auditory discrimination learning in mice

Brice Bathellier
Sui Poh Tee
Christina Hrovat
Simon Rumpel

BATHELLIER@IMP.AC.AT
POHPOH1990@GMAIL.COM
CHROVAT@GMX.AT
RUMPEL@IMP.AC.AT

Research Institute of Molecular Pathology

Most learning models are based on additive learning rules in which the synaptic update does not explicitly depend on the current weight of the synapse. In contrast, recent observations of temporal fluctuations of dendritic spine volume (a proxy of synaptic efficacy) show that synaptic weight updates in the brain strongly depend on the current weight of the synapse and suggest that biological learning rules may in fact be multiplicative. However, behavioral signatures of this fundamental discrepancy are lacking so far. We recently observed that individual mice learning an auditory Go/NoGo discrimination task often showed sigmoid-like learning curves, with a long initial 'delay phase' where the performance stays at chance level followed by a steep increase of the performance. To understand this phenomenon, we used a reinforcement learning model equivalent to a circuit in which a set of excitatory 'sound' neurons projects with plastic synapses onto a decision unit. For this architecture, we could show analytically and confirm by simulations, that when the learning rule is additive no 'delay phase' can exist. In contrast, the same model endowed with a multiplicative learning rule can easily reproduce learning curves of individual mice, provided that the initial synaptic weights are low enough. Hence, we reasoned that the 'delay phase' should vanish if we could put the mice in a situation where the connections relevant for learning a certain task have already a high efficacy. Such a situation can be obtained in a reversal experiment where the Go and NoGo sounds are swapped after overtraining the mice on the initial task which should result in strengthening of the relevant connections. In line with the model, we observed no 'delay phase' for a reversal experiment. Altogether, these observations support the idea that the brain uses multiplicative learning rules affecting the dynamics of learning-related behaviors.

I-53. Strategies for encoding competing acoustic events by single neurons in primate auditory cortex

Yi Zhou¹Xiaoqin Wang²

YIZHOU@ASU.EDU

XIAOQIN.WANG@JHU.EDU

¹Arizona State University²Johns Hopkins University

In a natural environment, sensory signals are often mixed with distractors and background noises that cause difficulties in discerning the presence of a target stimulus. To minimize the masking effect, one strategy is to suppress noise by a stronger target in perceptual tasks. Yet many neurons in the auditory cortex do not respond more strongly to a louder sound. Rather, their discharge rates decrease, after reaching their maximum, with increasing sound level. It is not entirely clear how these level-tuned neurons contribute to the task of anti-masking. To address this question, we investigated tone-in-noise (T) detection by single neurons in the auditory cortex of awake marmoset monkeys. We found that the firing patterns of cortical neurons could reliably signal the presence of either T or N or both, depending on the T level. The pattern was dominated by noise-like responses at low T levels and tone-like responses at high T levels, while T and N patterns coexisted at moderate T levels. Importantly, target-like and noise-like firing patterns extended to both excitatory and inhibitory epochs in neural activity. Although the level-tuned neurons showed similar rate responses at low and high T levels, the related firing patterns had different levels of tolerance to noise. At high T levels, the noise-driven activity was greatly reduced through inhibitory epochs in the responses of level-tuned neurons. These results demonstrate that the role of level tuning in sensory coding could be better understood in a situation with competing sensory stimuli. We argue that single neurons in the auditory cortex could coordinate their firing patterns associated with different stimuli. By using both excitatory and inhibitory response epochs, multiple streams of sensory information can be transmitted either together or alternately by the auditory cortex.

I-54. View-invariance and mirror-symmetric tuning in a model of the macaque face-processing system

Joel Z Leibo¹Fabio Anselmi¹Jim Mutch¹Akinori Ebihara²Winrich Freiwald²Tomaso Poggio¹

JZLEIBO@MIT.EDU

NFABIO.ANSELMI@GMAIL.COM

JMUTCH@MIT.EDU

AEBIHARA@MAIL.ROCKEFELLER.EDU

WFREIWALD@MAIL.ROCKEFELLER.EDU

TP@AI.MIT.EDU

¹Massachusetts Institute of Technology²Rockefeller University

Recent experimental results characterizing the face-processing network in macaque visual cortex pose a major puzzle. View-tuned units (found in patches ML/MF) are a natural step to a view-tolerant representation (found in patch AM), as predicted by several models. However, the observation that cells in patch AL are tuned to faces and their mirror reflections remains unexplained (cf. Freiwald and Tsao 2010). We show that a model based on the hypothesis that the ventral stream implements a compressed version of a memory-based approach to transformation invariance predicts the main properties of ML/MF, AL and AM. In this view, a major computational goal of the ventral stream is to compute invariant signatures that can be used to recognize novel objects under previously-seen transformations of arbitrary "templates" (exemplar objects). These invariant signatures can be regarded as encodings of a novel object relative to the principal components of the transformation videos of familiar objects.

I-55. Ventral stream models that solve hard object recognition tasks naturally exhibit neural consistency.

Daniel Yamins
Ha Hong
Ethan Solomon
James DiCarlo

Massachusetts Institute of Technology

YAMINS@MIT.EDU
HAHONG@MIT.EDU
ESOLOMON@MIT.EDU
DICARLO@MIT.EDU

Humans recognize objects rapidly and accurately, a major computational challenge because low-level pixel data can undergo drastic changes in position, size, pose, lighting, occlusion, etc, while still containing the same high-level content. There is substantial evidence that the brain solves this challenge via a largely feedforward, hierarchical, network called the ventral visual stream. However, fundamental questions remain about the actual neural implementation, an understanding gap reflected in the difficulty computer models have had in equaling human performance. Here we describe models that perform substantially closer to human levels on a hard object recognition task, and in doing so, naturally discover representations consistent with experimentally observed high-level ventral stream neural populations. We first constructed a large parameter set of hierarchical feedforward computational models, encompassing a variety of mechanisms that have shown promise in describing ventral stream encoding. To search this vast space for high-performing models, we developed a "Principled" High-Throughput (PHT) approach that blends powerful computational techniques with a structured selection procedure. The PHT procedure solves multiple recognition subtasks simultaneously, identifying target subtasks by error pattern analysis. Complementary model components emerge naturally, forming a representational basis that supports non-screened tasks. This process is repeated hierarchically, producing deep networks that are nonlinear combinations of lower-level components. Models were constructed using this procedure with screening images containing objects on natural backgrounds, and then tested on neurophysiologically-measured images of entirely different objects in differing categories to rule out overfitting. The models showed major improvement in performance compared to existing computational models, even with the significant pose, scale, and position variation that typically hurt algorithm performance. They also exhibited feature representations strikingly similar to those observed in IT cortex, suggesting that the model's component substructures may predict identifiable functional motifs in higher-level ventral areas.

I-56. Similarity between spontaneous and sensory-evoked activity does suggest learning in the cortex

Cristina Savin¹
Pietro Berkes²
Chiayu Chiu³
Jozsef Fiser⁴
Mate Lengyel¹

¹University of Cambridge

²Brandeis University

³Yale University

⁴Central European University

CS664@CAM.AC.UK
PIETRO.BERKES@GOOGLEMAIL.COM
CHIAYU.CHIU@GMAIL.COM
FISER@BRANDEIS.EDU
M.LENGYEL@ENG.CAM.AC.UK

The developmental increase in similarity between spontaneous (SA) and average stimulus-evoked activity (EA) in the primary visual cortex has been suggested to reflect a progressive adaptation of the animal's internal model to the statistics of the environment (Berkes et al., Science 2011). However, it is unknown how much of this adaptation is due to learning or simple developmental programmes. If learning plays a role, it makes two predictions: changes in the functional connectivity between neurons should underlie the changes seen during development, and these developmental changes should be experience-dependent. Neither of the two has been satisfyingly tested, if at all, in previous work. Here we address the issue of functional coupling by novel analyses with maximum entropy

models (Schneidman et al., Nature 2006) that control not only for the effects of single unit firing rates, but also for the population firing rate distribution which could otherwise confound measures of functional connectivity (Okun et al., SfN, 2011). We show that functional connectivity plays an increasing role during development in shaping both SA and EA, and in particular that it significantly contributes to the similarity of SA and EA. Moreover, we directly assess the role of experience by comparing neural activities recorded in animals reared with their lids sutured (LS) to those recorded in normally developing controls. Neural activity in LS animals was qualitatively similar to that in controls, confirming that withholding natural visual experience does not abolish the general development of the visual system. However, there were some key differences: the match between SA and EA remained incomplete, and the specificity of this match for natural images was significantly reduced in LS animals. Taken together, these results strongly suggest that learning in the cortex crucially contributes to the similarity between SA and EA.

I-57. Uncertainty and the striatum: How tonically active neurons may aid learning in dynamic environments.

Nicholas Franklin
Michael Frank
Brown University

NICHOLAS_FRANKLIN@BROWN.EDU
MICHAEL_FRANK@BROWN.EDU

Computational reinforcement learning models of the basal ganglia often assume a fixed learning rate, making them suboptimal for flexibly adapting to non-stationarity in the environment. An optimal learner takes their own uncertainty into account to decide how much to update action values based on any given outcome. Here we consider how giant, cholinergic tonically active neurons (TANs), may provide a mechanism by which to modulate learning as a function of expected and unexpected uncertainty. Constitutively active TANs were added to a previously published neural model of the basal ganglia by Frank, 2006. Effects of M4-muscarinic receptors activation were simulated through direct inhibition of direct and indirect pathway medium spiny neurons (MSNs). Effects of M1-muscarinic receptor activation were simulated through a persistent increase in leak channel conductance in the indirect pathway. A stereotypical burst-pause TAN firing pattern of varying duration was simulated during reinforcement feedback. By modulating MSN activity and learning, TANs improved probabilistic reversal learning but with a tradeoff: long TAN pauses result in better asymptotic performance whereas short TAN pauses facilitate speeded learning following reversal. This tradeoff arises from TAN modulation of the degree to which indirect MSNs are active and thus eligible for learning during probabilistic negative outcomes. Longer pauses were also related to greater changes in entropy among MSN unit activity during learning. These findings suggest that TAN pause duration may be dynamically controlled by entropy of MSN activity signaling uncertainty in action values, promoting both stable and flexible learning regimes.

I-58. How neurogenesis and the scale of the dentate gyrus affect the resolution of memories

James Aimone
Craig Vineyard
Sandia National Laboratories

JBAIMON@SANDIA.GOV
CMVINEY@SANDIA.GOV

The dentate gyrus (DG) is one of two regions to receive new neurons throughout life in mammals. The existence of neurogenesis in the DG has been difficult to reconcile computationally with the regions long presumed function of pattern separation in memory formation. Specifically, if the DG's function was limited to the generation of a sparse, nearly orthogonal, set of neurons, should not the DG's large size be sufficient without neurogenesis? Recently we proposed an alternative view of DG function that accounts for neurogenesis as well as satisfies the DG's role in facilitating discrimination between memories (Aimone et al., Neuron 2011). We proposed that the DG has a mixed coding scheme: the 'classic' sparse code provided by mature neurons and a distributed code

comprised of more active immature neurons. From an information content perspective, this scheme would permit familiar features to be encoded by high-information mature neurons while assuring that all features, regardless of previous experience, would be encoded by at least the low-information immature neuron population. Here, we use a spiking neural network model of the DG with neurogenesis to examine whether this hypothesis is indeed a valid computational possibility for the real system. Our model is strictly constrained by the biological properties of the network with regard to neuronal dynamics, connectivity, input structure, neuronal maturation and plasticity. Notably, we find that while immature and mature neurons do differ in their information content and responses to novel and familiar features, this distinction is only clear in realistic scale models. The observation of a complex relationship of scale with neurogenesis function has implications for the relevance of neurogenesis in humans where rates are thought to be significantly lower. Further, the model demonstrates new possibilities for what the DG is doing in the global hippocampus circuit during memory formation.

I-59. A not so bad tradeoff: perception and generalization of aversive memories

Rony Paz

RONY.PAZ@WEIZMANN.AC.IL

Weizmann Institute

Dangerous stimuli should be better detected, but should they be better discriminated? We worked with the other hypothesis, that it is safer to have worse discrimination for stimuli that predict aversive consequences: if A predicts danger, and A' is similar; then it is likely that A' predicts similar danger. The most efficient way to activate a fight-or-flight response is to not discriminate the two stimuli to begin with. We first tested this using psychophysics, conditioning tones to odors, and found that in contrast to common improvement in sensitivity that is seen after training or mere-exposure, tones that were conditioned with an aversive odor have higher sensitivity thresholds, measured by just-noticeable-difference (JND). We then show that this lasts 24 hours, hence real perceptual learning; it is robust to the reinforcer modality (odors/sounds), pointing to a central mechanism; and it is an active-dynamic process formed during learning, hence not an attentional or general blocking of stimulus processing. Next, we tested secondary-reinforcers, and found that monetary-loss also results in perceptual deterioration and less sensitivity. Using imaging, we identify the Amygdala and prefrontal regions that modulate it as involved in the learning and its consequences. In additional experiments, we show that these perceptual changes can compromise decisions and induce 'irrational' choice-behavior, and might even underlie susceptibility to anxiety-disorders. To characterize the network architecture that can underlie this, we record neurons in the primate amygdala during aversive-conditioning. We show specific changes in tuning curves that reshape their width in relation to the conditioned tone. Using the real data and synthesized data with similar parameters regime, we find a tradeoff in representation: more information about the conditioned-stimulus, yet less discriminability surrounding it. We conclude that this representational tradeoff in stimulus space makes sense in light of evolution and normal behavior, yet can underlie some abnormal behaviors.

I-60. Spatial learning in *C. elegans*

Adam Calhoun^{1,2}

Srekanth Chalasani²

Tatyana O. Sharpee^{2,1}

A2CALHO@UCSD.EDU

SHALASANI@SALK.EDU

SHARPEE@SALK.EDU

¹University of California, San Diego

²Salk Institute for Biological Studies

In order to make optimal decisions in a changing environment, organisms must learn the structure of that environment. Here we characterize a previously unknown learning behavior in *C. elegans*. Upon removal from food, *C. elegans* makes many turns which decrease over time until reaching a basal level after ~15 minutes. We

have found that the number of turns and search strategy is dependent upon the animal's prior experience with food, with experience of a larger food patch resulting in more turns and thus a smaller area searched. Through a dimensional reduction technique (maximum noise entropy), we are able to extract the sensory filter that the animal is learning. We identify two D1-like dopamine receptor types and which neurons they are acting on, as well as identifying a single set of dopamine neurons required for the learning. The CREB homologue *crh-1* is also required for learning and the neuron required is identified. Finally, we use tetanus toxin (TeTx) to construct the full circuit required for the learning behavior. These neurons include ones that primarily drive off-food motor behavior, suggesting that learning in one environment requires the neurons that will later be used for behavior in a different environment. Loss of any of these neurons removes the ability to learn, stressing the importance of learning at the circuit level over the single synapse.

I-61. The effect of STDP temporal structure on the learning of single excitatory and inhibitory synapses

Maoz Shamir
Yotam Luz

SHMAOZ@BGU.AC.IL
YOTAM.LUZ@GMAIL.COM

Ben Gurion University of the Negev

Spike-Timing Dependent Plasticity (STDP) is characterized by a wide range of temporal learning patterns, depending on the studied system and experimental conditions. Long Term Potentiation (LTP) is marked by a positive sign of the characteristic function and Long Term Depression (LTD) by a negative sign. It is a common practice to define this function in segments of the time interval — typically in two segments, one for positive Δt (the causal branch) and the other for negative Δt (the acausal branch). Here we suggest a model in which this pattern is constructed from a superposition of two separate processes one for the LTP and the other for the LTD. We approximate these two functional branches using a continuous non-segmented 'probability like' function that captures the essential features of the STDP. We demonstrate how the various experimentally observed STDP temporal structures can be obtained by a gradual change of a single continuous parameter in our model. Analysis of the STDP dynamics reveals a critical point. Below this critical point the STDP dynamics is governed by a negative feedback and the synaptic weights are characterized by a unimodal distribution. Above this point, the stability of the STDP dynamics is governed by the synaptic weight dependence of the STDP rule. In the latter case there is a different parameter with a critical value, above which, a bimodal synaptic weight distribution exists. We show that the location of these critical points depends on general properties of the temporal structure of the STDP rule and not on its fine details. These results hold for both excitatory and inhibitory synapses. The symmetry in the learning dynamics of excitatory and inhibitory synapses is discussed.

I-62. Quantifying representational and dynamical structure in large neural datasets

Jeffrey Seely¹
 Matthew Kaufman²
 Adam Kohn³
 Matthew Smith⁴
 Anthony Movshon⁵
 Nicholas Priebe⁶
 Stephen Lisberger⁷
 Stephen Ryu⁸
 Krishna Shenoy⁸
 Larry Abbott¹

JSSEELY@GMAIL.COM
 MKAUFMAN@CSHL.EDU
 ADAM.KOHN@EINSTEIN.YU.EDU
 MASMITH@CNBC.CMU.EDU
 MOVSHON@NYU.EDU
 NICHOLAS@MAIL.UTEXAS.EDU
 SGL@PHY.UCSF.EDU
 SEOULMANMD@GMAIL.COM
 SHENOY@STANFORD.EDU

¹Columbia University

²Cold Spring Harbor Laboratory

³Albert Einstein College of Medicine

⁴University of Pittsburgh

⁵New York University

⁶University of Texas at Austin

⁷University of California, San Francisco

⁸Stanford University

Systems neuroscience often employs models that explain neural responses in terms of represented stimulus features or movement parameters. These models can be powerful, but may not apply equally well when neural activity is dominated by intrinsic dynamics. Here, we examine large datasets from a number of cortical areas and ask whether responses appear stimulus dominated (i.e., are most naturally described in terms of tuning functions for external parameters) or appear dominated by internal dynamics (i.e., where the future population response is a function of the past population response). We analyzed datasets (44 - 218 single and/or multi-unit isolations) from visual areas V1 and MT (recorded during the presentation of visual stimuli) and from primary motor and premotor cortex (recorded during a delayed reach task). Our analyses did not fit particular tuning or dynamical models, but instead asked whether basic features of the data tended to obey or violate expectations of representational and dynamical systems. Our datasets consist of firing rate values indexed by neuron, condition (stimulus), and time. Our analysis involves a higher-order generalization of SVD (a tensor decomposition) to expose two kinds of structure potentially present in the data. First, when the responses of different neurons reflect tuning for a set of 'represented' stimulus parameters, then structure should be best captured across neurons. Second, when the responses for different conditions reflect the evolution of fixed dynamics from a set of initial states, then structure should be best captured across conditions. We employed simulations to confirm that these expectations hold. We then applied this method to six neural datasets from four cortical areas. For datasets from V1 and MT, the dominant structure was across neurons, consistent with a representational framework. For motor and premotor cortex, the dominant structure was across conditions, consistent with a dynamical framework.

I-63. Motor coding in the supplementary motor area of humans and monkeys

Gal Chechik¹
 Hadas Taubman²
 Ariel Tankus³
 Eilon Vaadia⁴
 Itzhak Fried⁵
 Rony Paz⁶

GAL.CHECHIK@GMAIL.COM
 TAUBMANH@GMAIL.COM
 ARIELTA@GMAIL.COM
 EILON.VAADIA@ELSC.HUJI.AC.IL
 IFRIED@MEDNET.UCLA.EDU
 RONY.PAZ@WEIZMANN.AC.IL

¹Gonda Brain Research Center

²Bar Ilan University

³Department of Neurology and Functional N

⁴The Hebrew University of Jerusalem

⁵University of California, Los Angeles

⁶Weizmann Institute

Neural coding of movement planning has been studied intensively in trained primates, but little is known about its implementation in naive human subjects. Using depth electrodes in two neurosurgical patients and in three trained monkeys, we characterized the activity of small groups of neurons in the supplementary motor area (SMA), a brain region known to participate in movement planning and execution. We find that many human cells exhibit very narrow sensitivity to movement direction, as compared to parallel analysis of neurons in the SMA and primary motor area (M1) of rhesus monkeys. These narrow directional tuning curves in turn yield significantly lower signal correlations between pairs of cells with adjacent preferred directions. However, due to higher variability across trials in human SMA, decoding movement direction from individual and groups of human SMA neurons achieves significantly lower accuracy and conveys lower information than obtained with monkey neurons. Interestingly, coding the movement direction in human SMA neurons decays strongly around movement onset, while monkey SMA and M1 neurons continue to participate in direction coding throughout the movement. In comparison, the accuracy of decoding the movement speed in human neurons peaks shortly after the movement onset. This suggests that the human SMA may be rapidly switching from coding the direction of movement to coding the speed or other aspects of movement.

I-64. Transient collective dynamics in inhibition-stabilized motor circuits

Guillaume Hennequin^{1,2}

Tim P Vogels³

Wulfram Gerstner⁴

GJE.HENNEQUIN@GMAIL.COM

TIMVOGELS@GMAIL.COM

WULFRAM.GERSTNER@EPFL.CH

¹School of Computer and Communication Sciences

²Brain Mind Institute

³Ecole Polytechnique Federale de Lausanne

⁴EPFL, BMI

The generation of motor patterns has been the focus of several recent experimental studies. Recordings have shown that populations of neurons in motor cortex transition into a preparatory state while a movement is being planned, and engage in transient collective dynamics of large amplitude during its execution. We investigate this phenomenon in rate models of cortical dynamics. Weakly coupled networks cannot produce the substantial transient departure from background activity observed in the experiments. Strongly coupled random networks with their inherent chaotic dynamics, on the other hand, do not capture the transient nature of movement-related activity. Here we introduce a new class of models with strong and complex excitatory recurrence, and inhibitory feedback of matching complexity to stabilize the dynamics. We show that such inhibition-stabilized networks transiently amplify certain network states. The network activity can be forced to arrive at one of those states by the end of the preparatory period through the delivery of an appropriate external input. Upon a go-signal, the input is withdrawn and the network is released to elicit transient single-neuron and collective dynamics that match the data well. In particular, we reproduce the recently uncovered phenomenon of rotational dynamics during the execution of movement [Churchland et al (2012)]. Additionally, muscle activity may be read out from these noisy transients to produce complex movements. Surprisingly, inhibition-stabilized circuits connect several previously disparate aspects of balanced cortical dynamics. The mechanism that underlies the generation of large transients here is a more general form of “Balanced Amplification” [Murphy and Miller (2009)], which was previously discovered in the context of visual cortical dynamics. Furthermore, during spontaneous activity in inhibition-stabilized networks, a detailed balance of excitatory and inhibitory inputs to single cell exists that is much finer than expected from shared population fluctuations.

I-65. Eye movements depend on the intrinsic reward structure in a natural navigation task

Constantin A Rothkopf^{1,2}
Dana Ballard³

ROTHKOPF@FIAS.UNI-FRANKFURT.DE
DANA@CS.UTEXAS.EDU

¹Frankfurt Institute for Advanced Studies

²University of Osnabrueck

³University of Texas at Austin

Empirical studies in active vision have shown that human gaze in extended sequential behavior such as navigation is not well predicted by low-level salience of visual stimuli such as luminance, contrast, or color. It has been suggested, that this is because in such naturalistic tasks vision primarily subserves behavioral task goals [1]. The difficulty is to quantitatively determine what the actual task goals are. Here we use an optimal control model of human navigation to infer the intrinsic costs and benefits implicit in observed navigation behavior using inverse reinforcement learning and show that these costs are strongly correlated with the associated gaze distributions. We present a computational model of human navigation to a target while avoiding obstacles and walking along via-points based on the optimal control framework. This model reproduces previous empirical data on average human walking trajectories [2]. We recorded navigation behavior and eye movements of participants in a set of virtual navigation tasks. A recently developed Bayesian inverse reinforcement learning algorithm [3] was applied to recover the rewards underlying the observed behavior and used to infer the relative contributions of avoidance, approach, and target reaching per trial and per subject. This quantifies the task goals of participants during naturalistic sequential behavior. We then related parameters of the eye movements to the inferred relative costs. E.g., the proportion of time spent on the obstacles and via-point targets is highly correlated with the inferred costs across all conditions. We conclude that intrinsic costs and benefits in human navigation can be inferred from walking behavior with inverse reinforcement learning and that eye movement properties are highly correlated with the inferred rewards. This provides evidence for reward based gaze allocation in natural tasks.

I-66. Behavior-modulated correlation of cerebellar Purkinje neuron and network activity

Sungho Hong¹
Mario Negrello²
Erik De Schutter¹

SHHONG@OIST.JP
MNEGRELLO@GMAIL.COM
ERIK@OIST.JP

¹Okinawa Institute of Science and Technology

²Erasmus MC

Correlations are widely observed in many neural systems not only between similar neurons but also across different components of the circuits, such as excitatory/inhibitory neurons (Okun and Lampl, 2008; Cafaro and Rieke, 2010). Importantly, the correlations can be dynamically modulated by various mechanisms such as adaptation (Gutnisky and Dragoi, 2008), sensory stimulus (Middleton et al., 2012), attention (Cohen and Newsome, 2008), etc. Correlations between individual neurons and the overall activity of the network, often represented by local field potential (LFP), have been widely observed and proposed as a mechanism for multiplexed coding (Huxter et al., 2003; Panzeri et al., 2010). Here we show that correlations between the activity of cerebellar Purkinje neurons (PC) and cerebellar LFP can be dynamically modulated. We analyzed the extracellular recording data from the vermal cortices of rhesus (*Macaca mulatta*) monkeys during spontaneous and guided saccades (provided by Marc Junker and Peter Thier, Universität Tübingen), and found that the PC simple spikes are weakly correlated with the (low-pass filtered) LFPs when averaged over all the spikes. However, the correlation, both in magnitude and time scale, can significantly vary with the interspike intervals (ISI) before/after each spike. We also show that the LFP-eye speed correlation is almost irrespective of the saccade angle while the spike train-eye velocity correlation is often angle-dependent, which implies that the saccade angle modulates the LFP-spike correlation. PC spike trains often show characteristically long ISIs, so-called 'pauses', and it has been debated whether such

temporal information, other than the rate, can be used for information coding (De Schutter and Steuber, 2009). Our results suggest that the cerebellar cortex can use multiplexed coding where the population rate (reflected in the LFP) and spike times of individual PCs can transfer different information simultaneously.

I-67. Neuronal avalanches in the resting meg of the human brain

Oren Shriki¹

Jeffrey Alstott¹

Frederick Carver¹

Tom Holroyd¹

Richard Henson²

Marie Smith³

Edward Bullmore²

Richard Coppola¹

Dietmar Plenz¹

OREN70@GMAIL.COM

JEFFREY.ALSTOTT@NIH.GOV

CARVERF@MAIL.NIH.GOV

TOM.HOLROYD@NIH.GOV

RIK.HENSON@MRC-CBU.CAM.AC.UK

MARIE.SMITH@BBK.AC.UK

ETB23@CAM.AC.UK

COPPOLAR@MAIL.NIH.GOV

PLENZD@MAIL.NIH.GOV

¹NIMH/NIH

²University of Cambridge

³University of London

What constitutes normal cortical dynamics in healthy human subjects is a major question in systems neuroscience. Numerous *in vitro* and *in vivo* animal studies have shown that ongoing or resting cortical dynamics are characterized by cascades of activity across many spatial scales, termed neuronal avalanches. In experiment and theory, avalanche dynamics are identified by two measures (1) a power law in the size distribution of activity cascades, with an exponent of $-3/2$ and (2) a branching parameter of the critical value of 1, reflecting balanced propagation of activity at the border of premature termination and potential blow up. Here we analyzed resting-state brain activity recorded using non-invasive magnetoencephalography (MEG) from 124 healthy human subjects and two different MEG facilities using different sensor technologies. We identified significant events at single MEG sensors and combined them into spatiotemporal cascades on the sensor arrays, using multiple timescales. Cascade-size distributions obeyed power laws (Figure 1). For the timescale at which the branching parameter was close to 1, the power law exponent was $-3/2$. This behavior was robust to scaling and coarse-graining of the sensor arrays. It was absent in phase-shuffled controls with the same power spectrum or empty-scanner data. Our results demonstrate that normal cortical activity in healthy human subjects at rest organizes as neuronal avalanches and is well described by a critical branching process. Theory and experiment have shown that such critical, scale-free dynamics optimize information processing. Thus, our findings imply that the human brain attains an optimal dynamical regime for information processing. Neuronal avalanches could also provide a biomarker for disorders in information processing, paving the way for novel quantification of normal and pathological cortical states.

I-68. Network dynamics amplifies the effects of weak electric fields on gamma and slow-wave activity

Lucas C Parra

Davide Reato

The City College of CUNY

PARRA@CCNY.CUNY.EDU

DAVIDE.REATO@GMAIL.COM

The effects of transcranial electrical stimulation on neural activity and human brain functions have been the subject of intense investigation [1-4]. Despite these efforts, the mechanisms by which weak electrical stimulation can result in cognitive effects remain unclear. We showed that weak electric fields can modulate *in-vitro* hippocampal gamma oscillations in power and frequency, and that spiking activity can be entrained with fields as low as 0.2V/m

[5]. Equally striking is the observation that slow-wave oscillation during human sleep can be entrained by oscillating currents that may polarize cells by no more 200uV [6-7]. How can such small polarization affect network oscillations to such an extent? To explain the effects on gamma oscillations we build a computational network model with excitatory and inhibitory neurons, coupled to the applied fields through incremental polarization of cell somata. The sensitivity of excitatory/inhibitory balance predicted outsized effects on firing rate and spike timing, which we subsequently confirmed with intracellular recordings. The model was then refined to reproduce UP/DOWN state transitions that underlie slow-wave oscillations. When weak oscillatory fields are applied, these oscillations entrain to the applied fields by virtue of the sensitivity of the DOWN/UP state transition, thus explaining the findings of animal and human experiments. We further hypothesized that altered firing rate could alter homeostatic plasticity, which is reflected in the homeostatic decay of slow-wave power during the night. To test this prediction we analyzed human sleep EEG after transcranial stimulation. We found a correlation of slow-wave decay with values predicted by the model in combination with detailed anatomical simulations of current flow in the brain. Summarizing, we provide a detailed quantitative description of how weak electric field stimulation can lead to significant network effects, including the acceleration of sleep homeostasis.

I-69. Correspondence between perceptual salience of 4th-order visual textures and natural scene statistics

John Briguglio¹
 Ann Hermundstad¹
 Mary M. Conte²
 Jonathan Victor²
 Gasper Tkacik³
 Vijay Balasubramanian¹

JOHNBRI@SAS.UPENN.EDU
 ANNHERM@PHYSICS.UPENN.EDU
 MMCONTE@MED.CORNELL.EDU
 JDVICTO@MED.CORNELL.EDU
 GASPER.TKACIK@IST.AC.AT
 VIJAY@PHYSICS.UPENN.EDU

¹University of Pennsylvania

²Weill Cornell Medical College

³IST Austria

The regularities of natural signals are a starting point for understanding characteristics of early visual processing, e.g. the center-surround receptive fields of retinal ganglion cells. Can this matching between natural signal statistics and neural processing mechanisms be extended beyond the sensory periphery? Our recent work (Tkacik et al., 2010) showed that human sensitivity to isodipole (fourth-order correlated) synthetic textures, known to arise beyond V1, is closely related to the structure of fourth-order spatial correlations in natural scenes. This thus propose an organizing principle: The perceptual salience of visual textures increases with the variance (i.e. unpredictability) of the corresponding correlations over the ensemble of natural scenes. To test this idea we focused on local image statistics: correlations between two, three, and four adjacent pixels within a 2x2 square. For binarized images, there are four pairwise correlations – vertical (β_v), horizontal (β_h) and diagonal (β_d) - four third order correlations ($\theta_{1,2,3,4}$) and one fourth-order correlation (α). We measured these correlations in each image patch in a large ensemble taken from the UPenn Image Database. The variance in the correlations over the ensemble was robustly ordered as: $\text{Var}(\theta_{1,2,3,4}) < \text{Var}(\alpha) < \text{Var}(\beta_v) < \text{Var}(\beta_h, \beta_d)$. Thus our broad hypothesis predicted the same ordering of perceptual salience of artificial textures with correlations of different types. This prediction was confirmed in psychophysical experiments: observers' ability to use image statistics to segment artificial visual textures conformed to the ordering of their variances in natural images. Ongoing work tests whether the co-variance between different image statistics correctly predicts the co-variance in detection thresholds for artificial textures where two kinds of correlations (e.g. β_v , $\theta_{1,2,3,4}$) are turned on at the same time. Our results suggest even central neural mechanisms are efficiently tuned to the statistics of natural scenes.

I-70. Multiple temporally-tuned mechanisms control visual adaptation to contrast

Stephen Engel¹
Min Bao²

ENGEL@UMN.EDU
USTCBM@YAHOO.COM.CN

¹University of Minnesota

²Chinese Academy of Sciences

Adaptation optimizes vision to a changing world. Alterations to the environment can happen at many timescales, from very transient to semi-permanent. To adapt optimally, the visual system also adjusts at different timescales, with longer-lasting environmental changes producing longer-lasting effects. But how the visual system adapts in this way remains unknown. Here, we show that contrast adaptation—the most-studied form of visual adaptation—has multiple controllers, each operating over a different time scale. In three experiments subjects performed one of three tasks, either a contrast matching, contrast detection, or tilt adjustment, while adapting to contrast a specified orientation. Following a relatively long period (5 minutes) of adaptation to high contrast, subjects were ‘deadadapted’ for a short while (40 seconds) to a low contrast. Deadadaptation eliminated perceptual aftereffects of adaptation, but continued testing revealed their striking reemergence, a phenomenon known as spontaneous recovery. To model these results, we assumed behavior depended upon neural gain set by mechanisms that minimized error between the current setting and an optimal one for the current environment. Adapting to a high contrast pattern, for example, caused a reduction in gain, decreasing matched contrast and increasing detection thresholds. A single controlling mechanism could not account for spontaneous recovery. Instead, our data were well fit by a model containing two mechanisms, each operating on a different timescale, whose output summed to set the gain. Initial adaptation to high contrast caused both to signal for a gain decrease. Following rapid deadadaptation, the long-term mechanism continued to signal a decrease, but this was cancelled by the short-term mechanism’s signaling for an increase. During continued testing, the short-term mechanism quickly returned to baseline, but the long-term mechanism still signaled a decrease, producing spontaneous recovery. Contrast adaptation is likely controlled by multiple neural mechanisms, whose time constants span a broad range timescales.

I-71. When is sensory precision variable?

Shan Shen
Ronald Van den Berg
Wei Ji Ma

SSHAN@CNS.BCM.EDU
RB2@BCM.EDU
WJMA@BCM.EDU

Baylor College of Medicine

Recent physiological studies have found that neuronal firing is best modeled as a doubly stochastic process, in which not only spike counts, but also the underlying firing rates vary from trial to trial, even in response to the same stimulus (1-2). This finding might have implications for perception. Specifically, given the close relationship between firing rate and sensory precision, one might predict sensory precision to vary from trial to trial. Models of perception typically assume that precision is constant, and this appears adequate in simple discrimination tasks. However, evidence for variable precision has been found recently in several visual attention and visual working memory studies (3-5). This suggests that sensory precision is variable in some circumstances, but not in others. Here, we conducted 8 psychophysical experiments in humans to determine what factors might be responsible for variability in sensory precision. In each task, subjects determined whether a target stimulus was oriented clockwise or counterclockwise with respect to a reference orientation. In some experiments, visual context was provided by distractor stimuli. We tested whether the following factors made sensory precision variable: fluctuations in attention over time, the stimulus range, the presence of multiple items in a display, variability in the number of items across trials, the presence of distractors, and variability of the distractors. In each experiment, we compared an optimal-observer model with constant precision (CP) to one with variable precision (VP). We found that the VP model outperforms the CP model by an average log likelihood difference of 44.7 in the experiments with variability in the distractors, and the models perform equally in all other experiments (with an average difference

of -1.6). These results suggest that variability in precision is a consequence of the processing of variable sensory context.

I-72. Adaptive shaping of feature selectivity in the rodent vibrissa system

He Zheng¹
Douglas Ollerenshaw^{1,2}
Qi Wang¹
Garrett Stanley¹

NEURON@GATECH.EDU
D.OLLERENSHAW@GATECH.EDU
QI.WANG@BME.GATECH.EDU
GARRETT.STANLEY@BME.GATECH.EDU

¹Georgia Institute of Technology

²Emory University

The brain adapts to highly dynamic stimuli in the natural environment as it extracts information and forms sensory percepts. How the adapting stimuli may shape cortical responses and optimize information for different kinds of tasks is unknown. Using voltage-sensitive dye imaging in the vibrissa pathway of anesthetized rats, we previously found that adaptation decreases the magnitude and spatial overlap between the cortical responses to adjacent single-whisker stimulations. From the perspective of an ideal observer of the cortex, adaptation enhances the spatial localization of a whisker deflection at the expense of degraded detection. We have also shown that this phenomenon manifests behaviorally in awake, behaving rats. Here, we investigate how properties of adapting stimuli shape cortical responses and the ideal observer's ability to detect and discriminate a stimulus. We used adapting whisker stimuli from physiologically-relevant frequencies and velocities because they are prominent features encoded by the pathway. The cortical response is differentially modulated by the adapting stimulus features. Furthermore, the cortical response is not unique to each frequency and velocity combination, but rather, the extent of adaptation is proportional to adapting stimulus power. We find that the extent of adaptation shapes some key characteristics of the cortical responses, such as response magnitude, area of cortical activation, and their covariance, which are direct predictors of detection and spatial discrimination performances. We show that a stimulus becomes less detectable to the ideal observer with more profound adaptation, while discrimination performance increases. However, when the cortical response falls below a detectable threshold, the ideal observer cannot discern the features of the stimulus without being able to detect it first. These results suggest there may be an optimal adaptation where the responses to adjacent whisker stimulations are minimally overlapped to enhance spatial localization of the stimuli, but remain detectable.

I-73. Neural coding of perceived odor intensity

Yevgeniy Sirotnin¹
Roman Shusterman^{2,3}
Dmitry Rinberg⁴

YSIROTNIN@ROCKEFELLER.EDU
ROMA.SHUSTERMAN@GMAIL.COM
DMITRY.RINBERG@NYUMC.ORG

¹The Rockefeller University

²Janelia Farm Research Campus

³HHMI

⁴New York University

In olfaction, as in other sensory systems, intensity is a fundamental perceptual feature. However, we do not know how intensity is coded by neurons in the olfactory system. Strong links with perception have been established in the visual and somatosensory systems, however, such perceptual links are dramatically absent in olfaction, which has become an established system for studying neural coding due to its relatively simple, accessible, and evolutionarily conserved organization. To examine the link between intensity perception and neural activity, we compared odor intensity ratings in humans to odor evoked neural activity across a population of cells in the olfactory bulb of awake mice. For humans, numerical ratings of perceived intensity decreased both with odor

dilution and across successive sniffs of a constant odor source. In addition they became less reliable across trials after the first sniff. In mice, neurons responded to threefold dilutions of odor by changing their mean firing rates, response shapes, and increasing response latency. For the same odor concentration, the evolution of the temporal features but not average firing rates over consecutive sniffs resembled the changes with odor dilution on the first sniff. The temporal properties of responses were in agreement with perceptual results, becoming less distinct and more similar to responses evoked by lower odor concentrations. We conclude that temporal patterns of neural activity in the olfactory bulb likely represent fine changes in the perceived intensity of an odor.

I-74. How is duration information from multiple sensory sources combined?

Mingbo Cai
David Eagleman
Baylor College of Medicine

MINGBO.CAI@GMAIL.COM
DAVID@EAGLEMANLAB.NET

The perception of duration can be biased by the physical properties of a sensory stimulus. For example, visual stimuli with higher temporal frequency are perceived as longer (Kanai et al., 2006). Objects of different temporal frequencies often appear simultaneously in the environment, providing conflicting information about duration. Does the brain keep separate duration representations for each object, or form a single representation? If a single duration representation is kept, how is it formed? One possibility is by Bayesian cue integration (Ahrens & Sahani, 2011); another is by reading out the total neural energy for encoding all the stimuli (Eagleman & Pariyadath 2009, 2012). Human participants estimated the duration of Gabor patterns drifting at 1Hz and 6Hz (denoted by L for low and H for high frequency, and LH when the two were simultaneously presented). In Experiment 1, participants compared the duration of LH against H. Psychometric functions revealed no bias between them. This suggests observers might overweight the dominant frequency channel (every stimulus includes an H), or were able to keep separate duration representations for each frequency channel and only use the H channel for judgments. In Experiment 2, LH was always presented first, followed by LH, H, or L. Duration of H was perceived longer than LH, consistent with a Bayesian cue integration model. Relative to LH, the judgments to H and L were significantly different, ruling out the model of separate duration representations. The precision of judging LH was better than H and L for the majority of participants. Experiment 3 used a static Gabor pattern (S) as the standard stimulus, and showed a compatible result. These data suggest observers weight duration information from multiple stimuli to form a single estimate.

I-75. Controlling the trade-off between categorization and separation via resonance in the olfactory bulb

Maximilian Puelma Touzel^{1,2}
Michael Monteforte³
Fred Wolf³

MPTOUZEL@NLD.DS.MPG.DE
MONTE@NLD.DS.MPG.DE
FRED@NLD.DS.MPG.DE

¹Max Planck Institute for Dynamics and Self-Or
²Bernstein Center for Computational Neuroscien
³MPI for Dynamics and Self-Organisation

A brain must generalize in some cases and scrutinize in others. Finding a plausibly modulated mechanism that sets the degree to which sensory systems resolve their input spaces has been difficult. During learning, the olfactory cortex can nevertheless categorize and separate stimuli while LFP oscillations in the olfactory bulb are high. The phasic spiking generated by the oscillation is unlikely to play a role downstream whereas there is mounting evidence in the bulb for a code based on spiking trajectories. Trajectories with synchronous spikes, however, make up the basin boundaries that enclose unique, stable trajectories in the recent explanation of stable chaos observed in balanced, inhibitory network models. Thus, the higher the oscillation power, the shorter the distance

to the basin boundaries, implying smaller basins in the compact phase space and a deeper pattern reservoir that more finely partitions the input space, and vice versa. The intracellular resonator properties exhibited by mitral cells could be readily modulated by top down projections and likely contribute to the population oscillation. Thus, here we extend the research of stable chaos and the balanced state to cover this single neuron dynamical mode of operation. We built a spiking network modelled on the predominately inhibitory mitral-granule cell network. By varying a single parameter that controls the intrinsic frequency, we explore the network dynamics across integrator to resonator units, ending with a network that replicates the zebrafish bulb's first and second order spiking statistics and its network oscillation. Our analysis shows that the balanced state and stable chaos are qualitatively unchanged in resonator networks, except that the distance to basin boundary changes with the intrinsic frequency. This finding suggests a mechanism for controlling the operating point of the system along the trade-off between categorizing and separating inputs to facilitate learning by the olfactory cortex.

I-76. Multisensory integration of vision and intracortical microstimulation for sensory feedback.

Maria Dadarlat
Joseph O'Doherty
Philip N. Sabes

MDADARLA@PHY.UCSF.EDU
JOEYO@PHY.UCSF.EDU
SABES@PHY.UCSF.EDU

University of California, San Francisco

Efforts to sensorize neural prostheses have created an interest in delivering artificial sensory information to the brain by directly evoking functional patterns of neural activity. Ideally, this information could substitute for or be integrated with natural sensory input to improve prosthetic control and increase the sense of prosthetic embodiment. Previous psychophysical and computational work, from our lab and others, suggests that human sensory integration is a statistically efficient and highly adaptive process, driven by temporal and spatial correlations between sensory signals. We have exploited this process to develop a novel, learning-based approach to artificial sensory feedback. A rhesus macaque monkey was exposed to a multichannel, intracortical microstimulation (ICMS) signal delivered to somatosensory cortex. The signal was determined by a fixed but arbitrary mapping to natural visual feedback. We hypothesized that strict temporal congruency between these two signals would be sufficient for the monkey to learn to interpret and use the ICMS signal. First, the animal was trained to reach to a hidden target using continuous visual feedback of its current error (difference between target and hand) given in the form of a random dot kinetogram of variable coherence. After learning this task, we trained the animal with paired visual and ICMS feedback. The animal was then able to perform the task with the ICMS signal alone (sensory substitution). Furthermore, his performance with both vision and ICMS surpassed his performance with either signal alone (sensory augmentation). Notably, movement variability in the multisensory condition matched the optimal, minimum variance, combination, indicating that the animal was integrating the visual and ICMS feedback signals as he would for two natural signals. Together, these results demonstrate the potential power of a learning-based approach to artificial sensory feedback for neural prostheses.

I-77. Predicting V1 neural responses to natural movies using the shift-invariant bispectrum

Mayur Mudigonda^{1,2}

Ian Stevenson³

Urs Koster²

Christopher Hillar⁴

Charles Gray⁵

Bruno Olshausen²

MUDIGONDA@BERKELEY.EDU

I-STEVENSON@BERKELEY.EDU

URS@BERKELEY.EDU

CHILLAR@MSRI.ORG

CMGRAY@CNS.MONTANA.EDU

BAOLSHAUSEN@BERKELEY.EDU

¹Redwood Center For Theoretical Neuroscience

²University of California, Berkeley

³University of California Berkeley

⁴Redwood Center for Theoretical Neuroscience

⁵Montana State University

Evidence from electrophysiology [Purpura1994, Felsen2005] suggests that the visual system is highly sensitive to higher-order stimulus statistics. However, most models for the stimulus response of V1 neurons are limited to first- and second-order statistics, i.e. features defined on raw pixels or the power spectrum of the stimulus [David2005]. We explore the image bispectrum as a way to capture higher order features in the stimulus. We show that the performance of spiking response models can be improved by including these higher order features compared to only first and second order features. The bispectrum, which consists of the products of pairs of complex Fourier coefficients, has been used by researchers in machine learning and image coding to produce invariant representations and characterize higher order image features such as curvature [Krieger1997]. The elements of the bispectrum are translation-invariant like the power spectrum, yet retain relative phase information. This allows the bispectrum to capture features such as sharp edges, corners and T-junctions, which may underlie response properties of V1 cells. We test this hypothesis by fitting models to 128 cells recorded from cat and primate primary visual cortex. Three different models were fit to each cell: 1) raw pixels, 2) power spectrum and 3) the bispectrum of the stimulus movies. For 27/128 cells, the bispectrum model outperforms the pixel model and the power spectrum model. Thus, while the majority of cells can be better described as either simple cells with the pixel model or complex cells with the power spectrum model, a significant fraction (21%) of cells have more complex receptive fields and can be better modeled in terms of bispectrum features.

I-78. Spatial structure and organization of nonlinear subunits in primate retina

Jeremy Freeman¹

Greg Field²

Peter Li³

Martin Greschner³

Lauren Jepson³

Neil Rabinowitz¹

Eftychios A. Pnevmatikakis⁴

Deborah Gunning⁵

Keith Mathieson⁵

Alan Litke⁶

FREEMAN@CNS.NYU.EDU

GDFIELD@USC.EDU

PETERLI@SALK.EDU

GRESCHNER@SALK.EDU

LHJEPSON@GMAIL.COM

NC.RABINOWITZ@GMAIL.COM

EFTYCHIOS@STAT.COLUMBIA.EDU

DEBORAH.GUNNING@STRATH.AC.UK

KEITH.MATHIESON@STRATH.AC.UK

¹New York University

²Zilkha Neurogenetic Institute, USC

³Salk Institute for Biological Studies

⁴Columbia University

⁵University of Strathclyde

⁶Physics Department, USC Santa Cruz

Sensory processing is commonly described using hierarchical cascades of linear and nonlinear operations. For example, in the primate retina, several types of retinal ganglion cells (RGCs) exhibit nonlinear responses to spatially-structured stimuli that can be explained by 'subunits' within the receptive field – localized filters with rectified outputs (Victor and Shapley, 1979). These subunits are hypothesized to reflect the function of bipolar cells that convey cone photoreceptor signals to RGCs, but their structure and function remain incompletely understood. We developed a novel approach to understand subunit computations in the retinal circuitry at single-cell resolution. Multi-electrode recordings and high-resolution stimuli were used to record from populations of identified RGCs in isolated primate retina while stimulating individual cones. Responses were fitted with a model consisting of two linear-nonlinear stages. The first stage consists of subunits that linearly combine signals from groups of cones followed by a nonlinearity. The second stage is a weighted sum of subunit responses followed by a final output nonlinearity. The assignment of cones to subunits was inferred using a greedy search for assignments that maximized response likelihood. Estimates of weights at both stages, as well as a smooth parameterization of the subunit nonlinearity, were obtained using block coordinate ascent on likelihood. Fitted subunits for ON and OFF midrange RGCs revealed varying degrees of rectification. Subunits typically included 1-3 cones, and convergence varied with eccentricity as predicted from anatomical data. The improvement in explained variance of RGC responses was typically 10-20% over a standard linear-nonlinear model for white noise stimuli, but much larger for noise segments that maximally differentiated the models. Additional validation was performed with repeated white noise, sinusoidal gratings, and targeted stimulation of selected pairs of cones. The results provide a picture of nonlinear signaling and circuitry in RGC populations at cellular resolution.

I-79. Modulatory signals from V1 extra-classical receptive fields with distinct spatio-temporal dynamics

Christopher Henry
Michael Hawken
New York University

HENRY@CNS.NYU.EDU
MJH2@NYU.EDU

Stimuli that do not directly elicit spiking responses in V1 neurons have been shown to strongly modulate responses to stimuli that drive their classical receptive fields (CRF). This modulatory influence from the extra-classical receptive field (eCRF) has been reported to be either facilitative or suppressive, which leads to different implications for the role that spatial context signals have within the cortical visual pathway. Studies of eCRF properties have largely focused on the average modulation to a given stimulus. We have used different paradigms (subspace reverse-correlation and isolated drifting gratings) to map the spatio-temporal dynamics of eCRFs in neurons from anaesthetized macaque V1. We found multiple component eCRF mechanisms with distinct temporal response profiles are present to varying extents in individual neurons across different cortical layers. Facilitation exhibited a relatively short latency, followed by orientation-untuned suppression that was slightly delayed and an orientation-tuned suppression that had a longer onset latency. Further, we found that the time course of eCRF suppression was also influenced by stimulus contrast; lower contrast evoked weaker suppression with an increased onset latency. Our working hypothesis is that the average modulation from a given stimulus will be a combination of the underlying component eCRF mechanisms with different time courses. The average modulation to a prolonged stimulus (of multiple seconds duration) will be influenced by both short- and long-latency mechanisms approximately equally; the average modulation to a stimulus of shorter duration (a few hundred milliseconds) will reflect a bias for mechanisms with shorter latency. Thus, quantitatively and qualitatively different contextual modulation effects can be produced by varying the stimulus duration. We model neural responses to stimuli of varying duration and show that the sign of modulation, strength of modulation, orientation tuning, and contrast sensitivity can vary systematically over time, influencing how spatial context signals are used for computation.

I-80. Silencing V2/V3 reduces spiking variability in MT: implications for excitatory/inhibitory balance

Camille Gomez
 Alexandra Smolyanskaya
 Gabriel Kreiman
 Richard Born
 Harvard Medical School

CAMILLE.GOMEZ@TCH.HARVARD.EDU
 ALEXANDRA_SMOLYANSKAYA@HMS.HARVARD.EDU
 GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU
 RICHARD_BORN@HMS.HARVARD.EDU

How does the convergence of synaptic input from different cortical areas contribute to the large trial-to-trial spiking variability observed in a cortical neuron? To explore this question, we studied how the reversible inactivation of visual areas V2 and V3 by cortical cooling affects the spiking variability of neurons in visual area MT in two monkeys performing a detection task. During visual stimulation under control conditions, MT neurons exhibited a Poisson spiking pattern (Fano factor, $F = 1.01$). With inactivation of V2/V3, mean spike counts were reduced, but spike count variance was disproportionately reduced further, resulting in sub-Poisson variability ($F = 0.72$). The reduction in Fano factor persisted in subsets of the data that were mean-matched for spike count, and it was present across different behavioral tasks, visual stimuli, and monkeys. We investigated how inactivation-related changes in the input to MT might alter spiking variability by simulating the partial inactivation of homogeneous synaptic inputs projecting onto an integrate-and-fire model neuron. The simulations also varied the fraction of inactivated excitatory and inhibitory inputs, which were pairwise correlated. We found that two conditions were necessary for reducing spiking variability during input inactivation: a high input regime (60–80 spikes / s) and a bias towards a predominant effect of inactivation of inhibitory input (5–10% more inactivation of inhibition over excitation). Weak pairwise correlation ($0.1 \leq r \leq 0.3$) further improved agreement with the distribution of spike count mean and variance observed experimentally. In summary, our model shows how an imbalanced inactivation of excitation and inhibition can alter renewal process spiking under a high input regime. The model could not, however, reproduce the magnitude of the observed reductions in mean count and Fano factor. This limitation suggests that population heterogeneity of MT input may also be affected by V2/V3 inactivation.

I-81. Distinct neural selectivity for 3D directions of visual motion.

Thaddeus B Czuba¹
 Lawrence Cormack²
 Alexander Huk²
 Adam Kohn¹

CZUBA@UTEXAS.EDU
 CORMACK@MAIL.UTEXAS.EDU
 HUK@UTEXAS.EDU
 ADAM.KOHN@EINSTEIN.YU.EDU

¹Albert Einstein College of Medicine
²University of Texas at Austin

Real world motion occurs in three dimensions. The processing of frontoparallel motion has been studied extensively and is relatively well understood, but how the visual system encodes motion towards or away from the observer remains unclear. Traditional accounts suggest such motion involves a changing disparity cue, but recent work argues that in many situations the most relevant cue might be interocular velocity differences (IOVD): an object moving directly toward or away from the observer will generate equal and opposite horizontal motion in the two eyes. We tested how the motion processing pathway encodes 3D motion by performing extracellular recordings in macaque area MT. We measured responses to a full matrix of monocular and binocular motion conditions using drifting sinusoidal gratings. The binocular conditions corresponded to a broad range of 3D motion trajectories. Many MT cells showed similar preferences for monocular motion in each eye and straightforward summation of these signals for binocular stimuli—these cells did not encode IOVD information. However, an interesting subset of cells exhibited robust IOVD information, evident either as opposite direction preferences for motion shown in each eye or strong nonlinear interactions for binocular motion. We also performed detailed measurements of disparity selectivity, and found cells encoding IOVD information could be either sensitive or insensitive to disparity. Together, our results suggest MT contains robust signals encoding of 3D motion through IOVD sensitivity.

Our data also provide a promising framework for exploring how cortical neurons combine well-defined signals for performing a distinct computation needed for the processing of real world visual input.

I-82. Maximum-entropy modeling of population calcium activity in the visual cortex

Dimitri Yatsenko¹
Kresimir Josic²
Andreas Tolias¹

YATSENKO@CNS.BCM.EDU
JOSIC@MATH.UH.EDU
ATOLIAS@CNS.BCM.EDU

¹Baylor College of Medicine
²University of Houston

The nervous system is more than the sum of its parts: sensory perception emerges from the collective interactions of large populations of cells. Deeper understanding of neural computation will require detailed characterization of regularities in the joint population activity. However, even as multineuronal recordings continue to improve, few analyses extend beyond pairwise correlations. New statistical methods are needed to isolate significant, physiologically relevant joint features in the broader population activity. We demonstrate a novel statistical method to uncover significant features of multineuronal activity. Based on the maximum-entropy principle, the method approximates the joint probability distribution of population activity using a minimal number of parameters. We reason that the smallest subset of activity features that constrains the maximum-entropy distribution to accurately fit the empirical data is more likely to be physiologically relevant than other subsets. While the feature dictionary remains at the investigator's discretion, the algorithm draws from the dictionary to devise the parametric form of the distribution that optimally fits the data. In comparison to prior maximum-entropy applications in neuroscience, we extend the model to real-valued signals and employ a principled model selection criterion to make the model more informative about the underlying physiological mechanisms. We fitted recordings of two-photon calcium activity from groups of 70–150 cells in mouse visual cortex under fentanyl anesthesia. With the feature dictionary comprising stimulus, pairwise, and third-order interactions, fitted distributions required

<

500 terms to adequately explain the empirical patterns. The fitted models revealed new information: cells' visual tuning that is not attributable to interactions with the other recorded cells, graphs of pairwise interactions, and the clustering of high-order interactions. Principled and incisive models of population activity in two-photon imaging will become instrumental in establishing the relationship between the cytoarchitecture of neocortical circuits and their computational function.

I-83. A spatio-temporal lattice filter model for the visual pathway

Hervé Rouault^{1,2}
Karol Gregor^{1,2}
Dmitri Chklovskii^{1,2}

ROUALTH@JANELIA.HHMI.ORG
GREGORK@JANELIA.HHMI.ORG
MITYA@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus
²HHMI

In the mammalian visual pathway, information is processed at several stages before reaching the visual cortex. While transiting through the retina, lateral geniculate nucleus (LGN) and finally the cortex, information is compressed by progressively removing local spatio-temporal redundancies. Previously such compression, also known as predictive coding, was modeled by a single static filter that reflects the average correlations that exist among a wide diversity of natural visual stimuli. However, a single static filter cannot account for the stage-wise

structure of the visual pathway and the adaptation recorded in the LGN in response to variations in natural environments experienced by animals. Here, we propose to model the visual pathway by an adaptive filter known as Laguerre Lattice Filter (LLF), a circuit designed for adaptively removing correlations in time-varying signals. The non-trivial structure of the LLF is similar to neuronal connections observed in the visual pathway. LLFs are composed of two parallel branches, so-called backward and forward, the temporal responses of which are similar to temporal receptive fields of (broadly defined) lagged and non-lagged cells of the LGN respectively. Inter-stage transfer functions in the backward pathway of LLFs are all-pass filters, which closely match spike cross-correlation between retinal ganglion cells and LGN lagged cells. Most interestingly, the connection weights of LLFs can be learned using Hebbian learning rules. By training the weights of multichannel LLFs on natural stimuli we obtain a center-surround receptive field observed in the retina and the LGN. Furthermore, in response to changes in the stimulus ensemble we predict different receptive field changes for each processing stage and branches. Hence, our study provides a theoretical framework to understand signal processing and adaptation in the visual pathway and helps design further experiments.

I-84. 3D Random access ultrafast two-photon imaging reveals the structure of network activity.

R. James Cotton
Emmanouil Froudarakis
Peter Saggau
Andreas Tolias

RCOTTON@CNS.BCM.EDU
EFROUD@CNS.BCM.EDU
PETER@CNS.BCM.EDU
ATOLIAS@CNS.BCM.EDU

Baylor College of Medicine

Deciphering the connectivity structure of local networks and the principles of organization of their activity are fundamental questions in neuroscience. To date, the efforts to characterize the structure of population activity in vivo in the neocortex have primarily focused on measuring the distribution of pairwise correlations. However, knowing this distribution does not preserve the full correlation structure, which is important for characterizing network activity. Measuring the correlation structure in the neocortex has been hampered by the inability to get high quality data from a large number of neurons in a local microcircuit. We have developed a novel three-dimensional random access multi photon in vivo microscope, which can record the activity of hundreds of neurons in a 3D volume with high temporal resolution (e.g. 400 adjacent neurons sampled at 125hz). This method offers significant advantages over galvanometric two-photon imaging, which is too slow and typically restricted to a single plane and multi-electrode cortical recording methods, which lack dense coverage. We recorded the activity of hundreds of neurons from a small microcircuit while simultaneously patching a single neuron in the primary visual cortex of the mouse. We fitted the activity of patched cells as a function of both the stimulus and the neurons in the nearby microcircuit using generalized linear models. As expected, we found that the local microcircuit can significantly predict the trial-to-trial variability of the response beyond what can be accounted by the stimulus alone. Interestingly, the cells that were predictive made up only a small subset of the total population, rather than small contributions from many neurons. Moreover, cells that were more predictive of the patched cells were also more correlated amongst themselves compared to the rest of the population. This correlation structure is consistent with a clique like connectivity architecture shown previously in slice multi-patching studies.

I-85. Energy and Information in insect photoreceptors

Nikon Rasumov¹
Jeremy E. Niven²
Simon Laughlin¹

NIKALRAS@YAHOO.DE
JEN23@SUSSEX.AC.UK
SL104@CAM.AC.UK

¹University of Cambridge

²University of Sussex

We study noise-limited communication under energy constraints. To this end, we combine intracellular recordings with an RC-circuit model of the blowfly and desert locust photoreceptor. We measure light and current response to determine a complete set of biophysical properties in each photoreceptor and build individual models to avoid averaging errors. We use independent methods to obtain measurements of the light- and voltage-gated potassium conductance, which determines the gain and combine them to estimate the energy consumption used by the Na ATPase. Photoreceptor gain is set by its conductances to protect the incoming information from the synaptic noise. We find that the blowfly reduces its gain and sacrifices up to 20% of the maximum achievable information from the natural stimulus to save 90% of energy consumption. This finding questions the main hypothesis for sensory communication of maximizing information and sheds new light on gain amplification. This energy-aware trade-off can be further interpreted with regards to the insect visual ecology. As the locust has a slower angular velocity than the blowfly, it reduces its energy consumption in dim light intensities by 23% as compared to the blowfly. Additionally, we find that this energy-aware adaptation has important consequences for the circadian rhythm. The locust reduces its energy consumption by 38% from day to night as compared to the night state. Summarizing, we show how gain in the photoreceptor is adjusted to trade-off information processing and energy consumption and how this trade-off depends on the incoming stimulus. However, communication in noise-limited systems is not restricted to the photoreceptor. Analogue systems use gain to amplify the signal in dendrites, molecular and neuronal networks across sensory modalities.

I-86. Kernel regression for receptive field estimation

Urs Koster¹
 Bruno Olshausen¹
 Christopher Hillar²

URS@BERKELEY.EDU
 BAOLSHAUSEN@BERKELEY.EDU
 CHILLAR@MSRI.ORG

¹University of California, Berkeley

²Redwood Center for Theoretical Neuroscience

Linear receptive field models of V1 such as the GLM have limited predictive power, but are challenging to generalize with nonlinearities such as complex cell pooling. This is on one hand due to local minima in non-convex models, necessitating clever optimization schemes or hard-coded nonlinear features that need to be chosen ad hoc. The second difficulty is that these usually high-dimensional models require strong priors for regularization. Therefore, much of the explainable variance of the response, i.e. a stimulus-agnostic model fit to just the PSTH in response to a repeated stimulus, still eludes our models. We explore the use of kernel methods to extend the GLM in a flexible nonlinear way. Applying the kernel trick, we replace the dot product between the linear filter and stimulus in the GLM with a dot product in a high dimensional kernel space. The resulting model is equivalent to kernel logistic regression. The problem of estimating the linear receptive field is replaced by finding a linear combination of data samples that describe the space of stimuli that result in a spiking response. To model the activity of single units recorded from anesthetized cat visual cortex in response to natural movie stimuli, we compare the following methods: the standard GLM for phase-separated Fourier coefficients [David 04], the kernelized GLM, and a standard SVM classifier (LIBSVM [Chang 09] wrapped by scikit-learn) where classification outputs (spiking / no spiking) are converted to spike probabilities using Platt's method [Platt 99]. Performance is evaluated using receiving operator characteristics (ROC curves). For the the linear GLM we obtain $ROC_{GLM} = 0.72$, for the kernelized model $ROC_{kGLM} = 0.73$ and for the probabilistic SVM model $ROC_{SVM} = 0.76$ for logistic and RBF kernels.

I-87. Trade-off between curvature tuning and position invariance in visual area V4

Tatyana O. Sharpee^{1,2}
 Minjoon Kouh¹
 John Reynolds¹

SHARPEE@SALK.EDU
 MKOUH@DREW.EDU
 REYNOLDS@SALK.EDU

¹Salk Institute for Biological Studies

²University of California, San Diego

Humans can rapidly recognize a multitude of objects despite differences in their appearance. The neural mechanisms that endow high-level sensory neurons with both selectivity to complex stimulus features and 'tolerance' or invariance to identity-preserving transformations, such as spatial translation, remain poorly understood. Previous studies have demonstrated that both tolerance and sensitivity to conjunctions of features are increased at successive stages of the ventral visual stream that mediates visual recognition. Within a given area, such as V4 or the inferotemporal cortex (IT), tolerance has been found to be inversely related to the sparseness of neural responses, which in turn was positively correlated with conjunction selectivity. However, the direct relationship between tolerance and conjunction selectivity has been difficult to establish, with different studies reporting either an inverse or no significant relationship. To resolve this, we measured V4 responses to natural scenes, and using recently developed statistical techniques, estimated both the relevant stimulus features and the range of translation invariance for each neuron. Focusing the analysis on tuning to curvature, a tractable example of conjunction selectivity, we found that neurons that were tuned to more curved contours had smaller ranges of position invariance and produced sparser responses to natural stimuli. These trade-offs provide empirical support for recent theories of how the visual system estimates three-dimensional shapes from shading flows, as well as the tiling hypothesis of the visual space for different curvature values while allowing for increased pooling across spatial positions for shallower curvatures.

I-88. Multi-stability in motion perception combines multiple underlying neural mechanisms

Andrew Meso^{1,2}
 James Rankin³
 Pierre Kornprobst³
 Guillaume S Masson⁴

ANDREW.MESO@UNIV-AMU.FR
 JAMES.RANKIN@INRIA.FR
 PIERRE.KORNPROBST@INRIA.FR
 GUILLAUME.MASSON@UNIV-AMU.FR

¹CNRS/Aix-Marseille Université

²Institut de Neurosciences de la Timone

³INRIA Sophia Antipolis

⁴INT. CNRS/Aix-Marseille Université

Multi-stable perception is an interesting phenomena in which a constant but inherently ambiguous sensory input drives an observer's percept dynamically between mutually exclusive alternative interpretations. It provides an insight into the interaction between conscious perception and dynamic input transformation into multichannel neural representations. Previous work has proposed that dynamic shifts in perception can be driven by either stochastic processes or asynchronous neural adaptation in channels encoding the alternatives. We analysed the results of visual motion perception experiments in which human observers were presented a moving grating stimulus over 15s while eye movements and reports of perceptual switches were recorded. Two orthogonal directions — horizontal (H) and vertical (V) and the intermediate diagonal (D) remained in competition during the task. We varied the input signal strength using luminance contrast as a dependent variable and tested the data for three characteristic signatures to categorise underlying mechanisms according to predictions of noise and adaptation processes. We computed (1) the regularity of the duration of perceptual states by calculating the autocorrelation of a matrix of durations. These showed a consistent increase in the coefficient with contrast, reaching over 0.5 at the highest contrast. (2) The stability of perceptual states was estimated from the variance

of the computed direction of smooth eye movements during the task. Estimates increased with contrast. (3) The distribution of eye directions over multiple trials. These showed a slight shift from a tri-stable (H-D-V) to a bi-stable (H-V) solution structure with increasing contrast. We conclude from the current experiments that multiple mechanisms sub-serve perception, gradually shifting from noise dominating at low contrast to adaptation dominating at higher contrast. Other multi-stable phenomena may similarly be driven by multiple mechanisms. Categorical signatures like fitting gamma vs log-normal distributions to switching time data might not sufficiently capture this complexity over limited trials.

I-89. Changes in laminar synchrony in V1 reflect perceptual decisions.

Neda Shahidi¹

Ariana Andrei¹

Ming Hu²

Valentin Dragoi²

NEDA.SHAHIDI@UTH.TMC.EDU

ARIANA.R.ANDREI@UTH.TMC.EDU

MING.HU@UTH.TMC.EDU

VALENTIN.DRAGOI@UTH.TMC.EDU

¹University of Texas

²University of Texas, Houston

The functional implications of correlated and synchronous firing across cortical layers are still unknown. To examine the effects on behavior we looked at synchronized firing across and within layers in V1, while monkeys performed a delayed match to sample task. We report here that the laminar pattern of synchronous firing in V1 was reflective of the behavioral outcome on the task. In our experiments, the task was to detect a difference in orientation (3-5deg) two natural image stimuli that were presented 500-1200 milliseconds apart. We recorded more than 150 single units and local field potentials with laminar probes from two monkeys (*Macaca mulatta*). To identify the granular, super-granular and infra-granular layers, we looked at current source density of local field potentials. To measure the synchrony we looked at cross-correlogram (CCG) of spike trains. To reveal the flow of information between neurons, first, the effect of co-fluctuations of averaged spiking activity due to the visual stimulus and background activity were removed by using an instantaneous firing rates (IFR) predictor that was subtracted from the raw CCG. Second, co-firing is only meaningful for spikes that can be causally related i.e. not more than 10 milliseconds apart. Therefore we computed the area under CCG for time shifts of -10 to 10 milliseconds. This analysis was performed using a sliding time window of 200 ms, sliding in 20 ms steps across the length of the trial. Increased synchrony between granular layer neurons and either supra- or infragranular neurons was observed mainly on incorrect trials, whereas increased synchronization between supra and infragranular layer neurons was seen on correct trials both during the delay period. Our results suggest that changes in synchronous firing even in primary sensory areas can either contribute to, or reflect higher order perceptual processes.

I-90. Lag normalization in an electrically-coupled neural network

Gautam Awatramani¹

David Schwab²

Stuart Trenholm¹

Vijay Balasubramanian³

GAUTAM@UVIC.CA

DSCHWAB@PRINCETON.EDU

STUT@DAL.CA

VIJAY@PHYSICS.UPENN.EDU

¹University of Victoria

²Princeton University

³University of Pennsylvania

The slow transduction of light by rods and cones in the retina introduces unavoidable spatial lags as images move across the visual field. During the 30-150 ms response latency of ganglion cells (retinal projection neurons), a moving object can cover a substantial distance proportional to its velocity. Since the visual system can localize moving objects with exquisite accuracy, it must possess mechanisms to rapidly compensate for such velocity-

dependent spatial lags. Here, we describe a neural circuit in the retina with such capabilities. We demonstrated that in a gap-junction coupled population of directionally selective ganglion cells (DSGCs) which code superior motion, the leading edges of moving objects are registered at a nearly constant location in space over a 10-fold range of stimulus velocities, i.e. responses are “lag normalized”. In contrast, responses in uncoupled populations of DSGCs exhibited the expected lag. Paired recordings established that coupled DSGCs provide direct electrical input to downstream cells. We developed a model to show how priming signals transmitted over gap junctions can compound as a moving object sequentially stimulates coupled DSGCs. An analytical treatment and numerical simulations demonstrated that lag normalization arises as a collective effect. By masking the stimulus, we confirmed that lag normalization in the response onset develops over an extended region, several fold larger than the receptive field size of individual DSGCs, corroborating the cooperative nature of the effect. Drugs that block gap junctions abolished lag normalization, further establishing the role of lateral electrical coupling assumed in the model. To our knowledge, this is the first report describing how a neural population acts collectively to provide a faithful spatial representation of the location of moving edges. This surprising new collective phenomenon could be implemented in other gap-junction coupled neural populations in the brain to compensate for circuit delays.

I-91. Generalized reward-modulated STDP: a model of operant conditioning of cortical neurons

Robert Kerr¹
 David B. Grayden¹
 Doreen A. Thomas¹
 Matthieu Gilson²
 Anthony N. Burkitt¹

ROBKERR87@GMAIL.COM
 GRAYDEN@UNIMELB.EDU.AU
 DOREEN.THOMAS@UNIMELB.EDU.AU
 GILSON@BRAIN.RIKEN.JP
 ABURKITT@UNIMELB.EDU.AU

¹University of Melbourne

²RIKEN Brain Science Institute

Operant conditioning refers to an individual modifying its behavior based on some consequence of that behavior, such as reward or punishment. Experiments have shown that changes in the firing rates of individual neurons in the motor cortex of rhesus monkeys can be elicited through operant conditioning. In these experiments, the monkeys were presented with feedback based on the firing rate of a neuron (measured from an implanted electrode) and rewarded for increasing that rate. Underlying this behavioral learning is plasticity at the synaptic level. Reward-modulated spike-timing-dependent plasticity (RSTDP) has been proposed as such a model of synaptic learning and has previously been used to explore analytically the results of these biofeedback experiments. In RSTDP, neuromodulatory signals (such as dopamine) modulate the amplitude of the learning window. We introduce a generalization of RSTDP where, unlike classical RSTDP, the long-term potentiation and depression parts of the learning window (LTP and LTD) are modulated separately by a neuromodulatory signal. Our model is based upon the way that neuromodulators have recently been experimentally observed to modify STDP. Using the Poisson neuron model, we analytically investigate the conditions under which generalized RSTDP generates the results seen in the biofeedback experiments. We compare it to classical RSTDP and use numerical simulations with leaky integrate-and-fire (LIF) neuron models to support our findings. We show that generalized RSTDP is able to account for the change in the firing rate of a neuron and, contrary to previous studies, classical RSTDP is not. We also show that the reinforcement is only possible when the reinforced neuron is in a fluctuation-driven regime where it receives a balance of excitatory and inhibitory input.

I-92. Network motifs and collective spiking in neuronal networks

Yu Hu¹
James Trousdale²
Kresimir Josic²
Eric Shea-Brown¹

HUYU@UW.EDU
JRTRROUSD@MATH.UH.EDU
JOSIC@MATH.UH.EDU
ETSB@WASHINGTON.EDU

¹University of Washington

²University of Houston

One approach employed in connectomics is to characterize the frequencies of small connection patterns, or motifs, repeated multineuron sampling in vitro. Intriguingly, recent experiments found that certain motifs in biological neural networks occur at markedly divergent frequencies than what would be expected if the networks were randomly connected (as Erdos-Renyi (E-R) networks) (1; 2; 3). We aim to understand what – if anything – we can infer about the dynamics of biological neural networks based on these limited, empirically sampled aspects of connectivity (4). In particular, we identify the relationship between motif frequencies and the level of correlated, or synchronized, spiking activity among pairs of cells in a recurrent spiking networks. We choose this measure because correlations in spiking activity have been shown to impact population coding and signal transmission (5; 6; 7), cf. (8). We show that network-averaged spike correlations are determined by a novel set of network statistics, which we call motif cumulants. For a range of complex network architectures, we find that mean correlations can be accurately predicted using only a few low-order motif cumulants – that is, those involving small numbers of neurons and connections. This observation is useful, as experimental sampling of large motifs can be difficult; it is also nontrivial, as interactions over long paths still contribute significantly to network correlations. Up to second order, these required motif statistics are the overall network connection probability and frequencies of chain and diverging motifs – the prevalence of which increases spike correlations. We also demonstrate extensions to multi-population networks, which predict correlations in networks of excitatory and inhibitory cells, a topic of much recent interest (9; 10).

I-93. Quadratic programming of tuning curves: from spiking dynamics to function

David GT Barrett^{1,2}
Sophie Deneve
Christian Machens³

DAVID.BARRETT@ENS.FR
SOPHIE.DENEVE@ENS.FR
CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

¹École Normale Supérieure

²Champalimaud Centre for the Unknown, Lisbon

³Champalimaud Foundation

Firing rate tuning curves come in many shapes and sizes, from bump-shaped to sigmoidal, from sharply peaked to broad. What are the functional roles of these many shapes, and how do they obtain their shape from their underlying neural circuitry? These questions have been central to neuroscience since the first firing rate recordings of Adrian in 1928. Our approach is to work backwards, and ask: how should tuning curves be shaped for a given function, such as eye position representation or orientation representation? We assume that representation performance can be quantified with a quadratic cost function, and that this representation can be decoded with a linear decoder. We calculate the firing rates that optimise this cost function under the constraint that firing rates must be positive. This is known as a quadratic programming problem (a method for optimising a quadratic cost function under an inequality constraint). This framework leads to some surprising new insights. Neural populations optimised to represent angular variables such as orientation have bump-shaped tuning curves and those optimised to represent linear variables such as eye-position have oculomotor-like sigmoidal tuning curves. Beyond this, there is a huge diversity of optimal tuning curve shapes, reminiscent of experimentally observed heterogeneity. These curves are highly non-linear, despite our assumption that signals are linearly decoded. Additionally, tuning curves display rapid plasticity following neuron ablation. All of these results are a consequence of a net-

works attempt to optimise representation performance. As such, tuning curves should be considered network properties, not single neuron properties. These intriguing predictions require an underlying spiking model. To that end, we find that a network of tightly balanced leaky integrate-and-fire neurons can produce spike trains that optimise our quadratic cost function. Therefore, we can think of spiking dynamics as a quadratic programming algorithm, and tuning curves as the solution.

I-94. Learning activity patterns in recurrent networks of visible and hidden spiking neurons

Johanni Brea
Walter Senn
Jean-Pascal Pfister
University of Bern

BREA@PYL.UNIBE.CH
SENN@PYL.UNIBE.CH
PFISTER@PYL.UNIBE.CH

The brain is able to learn and reproduce activity patterns. It is however unclear what kind of dynamical system and learning rule best explains this phenomenon. Here we consider recurrent networks of stochastic spiking neurons. We divide the set of neurons into "visible" and "hidden" neurons. The goal of learning is that the distribution of activity patterns of the visible neurons approaches some target distribution, i.e. the learning rule should minimize the Kullback-Leibler divergence from the target distribution to the model distribution of visible activity patterns. The hidden neurons can help to increase the class of learnable distributions by adding additional degrees of freedom. However, training the hidden neurons is a notoriously difficult problem. The novelty of this study is that we derive a learning rule which performs stochastic gradient descent on an upper-bound of the Kullback-Leibler divergence from target to model distribution of visible activity patterns. This leads to a Perceptron-like learning rule for the visible neurons, which is compatible with Spike Timing Dependent Plasticity (STDP) and for the hidden neurons contains additionally a modulating factor, like in reinforcement learning, but here, the "reward" signal is not provided by the environment but by the neural network.

I-95. Associative memory encoding in bump attractor networks

Vladimir Itskov
Carina Curto
University of Nebraska - Lincoln

VLADIMIR.ITSKOV@UNL.EDU
CCURTO2@MATH.UNL.EDU

The hippocampus can be thought of as a "Swiss knife" of the brain. It is implicated in learning and memory, but it is also critical for spatial navigation. The spatial navigation function, as manifested by place cell activity, has been successfully modeled using bump attractor networks – i.e., networks of neurons where the synaptic efficacies vary according to the cells' relative positions in a "feature space." Here the feature space reflects the coding properties of neurons, rather than their physical locations within the brain. Interestingly, bump attractor networks are characterized by a highly structured pattern of synaptic connections, whereas functions such as associative memory encoding appear to require a different synaptic organization. How can the varied functions of the hippocampus be accomplished in the same network? We investigate this question in the context of a recurrent network that can function both as a bump attractor network and as a network capable of auto-associative pattern completion. Remarkably, we find that both functions can be realized on a network obtained as a perturbation of a bump attractor network. In this context, the level of spatially homogeneous excitation is the "switch" that determines which of the two functions the network is implementing. We suggest that sparse perturbations of bump attractor networks might be a generic mechanism that allows the same neuronal network to implement both associative memory and spatial navigation, with the structure of feedforward inputs determining the "switch" between these two operating regimes.

I-96. Recursive conditional means (RCM): A powerful mechanism for combining sensory estimates

Wilson Geisler
 Jeffrey Perry
 University of Texas at Austin

GEISLER@PSY.UTEXAS.EDU
 JSP@MAIL.UTEXAS.EDU

Many tasks performed by real and artificial sensory systems involve combining multiple noisy estimates of a given variable (e.g., depth) to obtain a more accurate estimate of the variable. The classic approach is to linearly combine the estimates in parallel, where the weight on each estimate is its relative reliability. Here we describe a more powerful and computationally efficient approach that is based on combining estimates recursively (e.g., the second estimate is based at least in part on the first estimate). We demonstrate this approach for the task of denoising images that are corrupted by multiplicative (Poisson-like) noise, but it is applicable to many sensory processing tasks. We show that this approach yields denoising performance that matches or exceeds the state-of-the-art algorithms in the image processing literature. And, importantly, the approach is simple and could be implemented hierarchically with plausible neural circuits. Thus, it represents a viable alternative for how sensory estimates are combined in neural systems.

I-97. Energy constraints link structure and function in thin axons in the brain's wiring

Ali Neishabouri
 Aldo Faisal
 Imperial College London

M.NEISHABOURI10@IMPERIAL.AC.UK
 ALDO.FAISAL@IMPERIAL.AC.UK

The energy necessary for propagating action potentials (APs) in axons [1] is stored in the form of ionic concentration gradients across the membrane. It is commonly assumed that the number of ions crossing the membrane is very small compared to the total number of ions involved, as this is the case in classically studied axons e.g. the squid giant axon [2] with diameters of hundreds of micrometers. However, the mammalian nervous system contains much thinner axons e.g. C fibres or cortical axon collaterals ($d=0.1-0.3 \mu\text{m}$). Since Na⁺K pumps operate 50 times slower than the duration of an AP, firing rates of thin axons may be limited by rapid depletion of energy. We investigate how homeostatic and metabolic constraints limit neuronal activity using a Hodgkin-Huxley type model which tracks changes of ionic concentrations. We establish a minimum Na⁺K pump density to support a given firing rate and find that in thin axons, a few APs (30 for $d=1 \mu\text{m}$) are sufficient to deplete concentration gradients to the point where the resting potential lies above the AP threshold. This results in a persistent burst of APs likely to provoke cell death (Figure 1). This effect could be prevented by sufficient pumping, but the small diameter of axons limits the diffusion rate of ATP from the mitochondria. Mitochondria are only found in the soma or in varicosities along the axon. We can thus derive the maximum amount of sustainable neuronal activity as a function of the distance to the nearest mitochondria, and axonal diameter. Neurons may need to manage their energy budget [3, 4, 5] more carefully in compact circuits of the cortex. As with noise [6], energetic considerations pose constraints on the anatomy of axons and limit the miniaturisation of neural circuits, the effects of which could be directly observed.

I-98. Adaptive gaussian poisson process: a model for in vivo neuronal dynamics

Simone Carlo Surace
Jean-Pascal Pfister
University of Bern

SURACE@PYL.UNIBE.CH
PFISTER@PYL.UNIBE.CH

Recent advances in experimental techniques have led to an abundance of intracellular single-neuron recordings in vivo. However, it is unclear how to accurately characterize the statistical properties of both the sub-threshold membrane potential and the spike timings. In particular, an important feature of neuronal activity is the presence of refractoriness and adaptation. Here, we propose a modified Cox process model with an additional adaptive mechanism. The membrane potential is modeled as a Gaussian process, and through a nonlinearity, a time-varying rate for an inhomogeneous Poisson spiking process is obtained. We include a simple adaptation mechanism into this framework and apply it to in vivo intracellular recordings from HVC region of the Zebra Finch. We propose a method for estimating the sub-threshold membrane potential at the time of the spike and learning the parameters of the Gaussian process, nonlinearity and adaptation kernel. The resulting model parameters are used to generate artificial neuronal data which are very similar to the experimental one.

I-99. Rapid update in state estimation accounts for fast feedback control

Frederic Crevecoeur¹
Stephen H Scott²

FC12@QUEENSU.CA
STEVE.SCOTT@QUEENSU.CA

¹Centre for Neurosciences Studies

²Centre for Neuroscience Studies

It is well known that the brain relies on predictive processes to compensate for neural transmission delays during movement execution such as altering grasp force when lifting an object. However, internal prediction cannot anticipate unexpected disturbances and the extent to which time delays alter the performance of motor responses to perturbations remains unknown. To address this issue, we tested whether delayed feedback control could account for human motor performances. Participants (N = 6) maintaining their hand at a spatial target had to respond to perturbation pulses (5Nm for 50ms) applied to the elbow joint and return to the goal target within 600ms or 300ms. We compared participants' performances with feedback controllers with or without the presence of a state estimator (Kalman filter). For the delay-uncompensated controllers, we used a continuous pole placement algorithm for state feedback control¹, and an algorithm that determines the region of position, velocity and acceleration gains achieving stable control in the presence of time delays². The second class of controllers was based on optimal feedback control coupled with an optimal state estimator. All parameters were measured or taken from physiological models. The first class of controllers (delays uncompensated) could stabilize the joint in less than 600ms. However, returning to the target within 300ms generated increasing oscillations even if delays were as short as 30ms. In contrast, human participants were able to smoothly modulate their feedback responses for both target times. Similar stable responses were also possible for the controller that relied on optimal state estimation. These results suggest that motor responses are not solely mediated by sensory feedback and that a rapid update in the estimation of the state of the limb drives feedback responses to mechanical perturbations.

I-100. Dynamical entropy production in cortical circuits with different network topologies

Rainer Engelken^{1,2}
 Michael Monteforte¹
 Fred Wolf¹

RAINER@NLD.DS.MPG.DE
 MONTE@NLD.DS.MPG.DE
 FRED@NLD.DS.MPG.DE

¹MPI for Dynamics and Self-Organisation
²Bernstein Center for Comp. Neuroscience

The prevailing explanation for the irregularity of spike sequences in the cerebral cortex is a dynamic balance of excitatory and inhibitory synaptic inputs — the so-called balanced state [1]. Nevertheless its statistical properties are well described by a mean field theory that is independent of the single neuron dynamics, its dynamics is far from being understood. Recently it was found that the stability of the balanced state dynamics depends strongly on the detailed underlying dynamics of individual neurons. Inhibitory networks of leaky integrate-and-fire neurons show stable chaos [2,3], while a balanced network of neurons with an active spike generation mechanism exhibits deterministic extensive chaos [4]. Previous studies of the dynamics of the balanced state used random (Erdős-Rényi) networks. We extended this analysis to arbitrary network topologies and analyzed the entropy production in small world topologies [5], ring networks [6], clustered networks [7], multi-layered networks [8] and topologies with different frequencies of certain network motifs [9]. We derived an analytical expression for the single spike Jacobian containing elements of the coupling matrix, which enabled us to calculate the full Lyapunov spectrum for any desired topology. Using a single neuron model in which action potential onset rapidness [10] and synaptic time constant are adjustable, we simulated the dynamics in numerically exact event-based simulations and calculated Lyapunov spectra, entropy production rate and attractor dimension for a variety of connectivities. Stable chaos - characterized by irregular activity despite of a negative largest Lyapunov exponent — was previously only shown in networks of inhibitory neurons [3]. Surprisingly, we found stable chaos in the multi-layered network with mixed excitatory and inhibitory neurons.

II-1. Long-term memory . . . now longer than ever

Marcus Benna
 Stefano Fusi
 Columbia University

MKB2162@COLUMBIA.EDU
 SF2237@COLUMBIA.EDU

Every time we store a new memory, we modify a population of synapses, which inevitably disrupts previously stored memories. This leads to a tradeoff between memory lifetime and the initial amount of information stored per memory (its strength). This rigidity-plasticity dilemma imposes severe limitations on memory performance in the biologically relevant context of online learning with bounded synapses. Previous attempts to overcome these limitations revealed the importance of the complexity of biochemical processes that result in long-term synaptic modifications (Fusi et al. 2005). Models that incorporated cascades of processes operating on multiple timescales could achieve memory lifetimes and initial strengths that scale with the square root of N , where N is the number of synapses. Here we introduce a new class of complex synapses that outperform all previous models, allowing for a memory lifetime that scales almost linearly with N and an initial memory strength that still increases approximately as the square root of N . The complexity of each individual synapse grows only logarithmically with the memory lifetime. These models have been constructed guided by the idea that memories should fade as gracefully as possible. The best performance is achieved when the memory trace decays just slightly faster than the inverse square root of the time since storage. This decay can naturally and parsimoniously be accomplished by intrasynaptic dynamics resembling a diffusion process. It couples variables that represent biochemical processes operating on a wide range of different timescales. While motivated by continuous diffusive systems, these models can be implemented as fully discretized stochastic processes without a significant drop in performance. Our study shows that biological complexity can be harnessed to achieve the optimal scaling of memory strength and lifetime with the number of synapses. We make specific predictions about the bidirectional interactions between

biochemical processes involved in memory consolidation.

II-2. Here's Waldo! A mechanistic model of visual search predicts human behavior in an object search task

Thomas Miconi^{1,2}

THOMAS.MICONI@GMAIL.COM

¹Children's Hospital Boston

²Harvard Medical School

When searching for a given object in a visual scene, how does the brain decide where to look next? Current psychological theories of visual search suggest the existence of a global “attentional map”, computed by integrating incoming visual information (including saliency) with top-down, target-specific signals. Where and how this integration is performed remains unclear. Recent experimental evidence suggests that attentional detection occurs in attentional areas (such as FEF) before visual areas (IT and V4), in contrast with many existing models of visual search. Furthermore, most existing models are only coarsely (if at all) validated against human behavior. Here we describe a simple mechanistic model of visual search which both fits the existing evidence and predicts single-trial human behavior in a visual search task among complex objects. Our model posits that a target-specific modulation is applied at every point of a retinotopic area, selective for visual features of intermediate complexity (identified with LIP) with local normalization through divisive inhibition. The resulting field is then summed locally to constitute a single attentional map, the maxima of which are the foci of covert or overt attention (presumably acting through FEF). We first show that this model can localize natural objects in both composite and natural scenes. Then, to validate this model against human behavior, we collected data from human subjects during an object search task, and ran the model on the same task. The model is able to predict human fixations on single trials well above chance, including error and target-absent trials. Importantly, this ability is preserved after controlling for low-level similarity (pixelwise correlation or Euclidean distance) between objects and target. We conclude that our model captures non-trivial properties of the system that guides visual search in humans.

II-3. NMDA-mediated feedback accounts for effects of visual spatial attention in Neurogrid simulations

Nicholas Steinmetz

Ben Benjamin

Kwabena Boahen

Stanford University

NICK.STEINMETZ@GMAIL.COM

BENVB@STANFORD.EDU

BOAHEN@STANFORD.EDU

Gain modulation by corticocortical feedback has been observed in many cortical regions and behavioral contexts (Salinas & Thier, 2000; Silver, 2010) including during selective spatial attention, when the spiking responses of visual cortical neurons are multiplicatively modulated (McAdams & Maunsell, 1999). The mechanisms by which feedback modulates gain in selective attention remain unknown, though in vivo (Fox, Sato, & Daw, 1990) and in vitro experiments (Polisky, Mel, & Schiller, 2004) have suggested that NMDA-dependent mechanisms may be involved in nonlinear neocortical modulations. We hypothesized that the voltage-dependence of NMDA receptors could lead to an increasing contribution of the modulatory feedback with greater visual drive, resulting in a multiplicative relationship between output firing rates with and without modulation. To examine the hypothesis that NMDA receptors could mediate gain modulation during spatial attention, we designed a large-scale model of a visual cortical area (V4) and a frontal cortical area (the Frontal Eye Field, FEF). Each area was modeled with an excitatory and an inhibitory population of spiking neurons. The FEF contained recurrent excitatory and inhibitory connections to allow for self-sustaining activation of a restricted population of neurons representing the spatial location of attentional focus. The FEF fed back topographically to NMDA synapses on the dendritic compartment of excitatory V4 neurons. Visual drive was modeled as spike trains synapsing on dendritic AMPA receptors. We sim-

ulated the model on Neurogrid (Silver, Boahen, Grillner, Kopell, & Olsen, 2007) and found that top-down feedback onto NMDA synapses could account for the particular multiplicative gain modulations observed in recordings from visual cortical neurons. Suppressive interactions between neurons at the focus of attention and those elsewhere were not required to reproduce the electrophysiological results. Based on these results we predict that blocking NMDA receptors in extrastriate visual cortex would eliminate attention-driven firing rate modulations.

II-4. The laminar origin of sensitivity to high-order image statistics in macaque visual cortex

Yunguo Yu
Anita Schmid
Jonathan Victor
Weill Cornell Medical College

YUY2002@MED.CORNELL.EDU
AMS2031@MED.CORNELL.EDU
JDVICTO@MED.CORNELL.EDU

Extraction of local features (e.g., edges and lines) and surface properties are important first steps in extracting meaning from natural images. Analysis of higher-order statistics (HOS's) is critical for these processes, but it is as yet unclear where and how the underlying neural computations occur. We therefore made multi-tetrode recordings in anesthetized macaque V1 and V2 while presenting 320-ms snapshots representing seven kinds of higher-order statistical structure: standard random binary checkerboards, two kinds of third-order structure, and four kinds of fourth-order structure. We determined whether responses of individual neurons distinguished between different kinds of HOS's, and if so, at what latency. Results showed a systematic dependence of dynamics on cortical area and layer. In V1 granular (input) layers, the smallest fraction of cells showed HOS sensitivity: on average, 25% of cells could distinguish any given HOS type from one of the others. The latency distribution was bimodal: one peak at 80 to 130 ms, and a later peak at 150 to 250 ms. Outside of the granular layers, a larger fraction of cells showed HOS sensitivity: about 45% of supragranular cells, and 55% of infragranular cells. In the supragranular layers, latencies were similar to the early granular mode (80 to 110 ms), but had an earlier peak. In the infragranular layers, latencies corresponded to the later granular mode. In V2, 70% of cells in the granular layer were sensitive to HOS's, and this fraction increased further in the extragranular layers (>80%). V2 latencies were generally longer than V1 latencies, and more dispersed. In sum, neural selectivity for HOS's appears first in V1 supragranular layers and later in its input layers, and then becomes more prominent in V2. The pattern of laminar dynamics suggests that the overall feedback architecture of the cortical microcircuit plays a critical role in these computations.

II-5. A biophysical model of Bayesian inference and MCMC sampling in neural circuits

Alessandro Ticchi
Aldo Faisal
Imperial College London

A.TICCHI12@IMPERIAL.AC.UK
ALDO.FAISAL@IMPERIAL.AC.UK

Experimental evidence at the behavioural-level shows that brain make Bayes optimal decisions, yet at the circuit level little is known experimentally about how brains may implement simultaneously Bayesian learning and inference. Here we show how wiring anatomy and local synaptic learning rules can work together with molecular sources of noise enabling populations of neurons to unify three features of Bayesian computations: a) simple local learning rules enable our model to learn optimal statistical representations of sensory information. b) In absence of novel sensory information, the ubiquitous ion channel noise in neurons drives our model to autonomously produce samples from learned sensory input distributions i.e. representing the prior and c) local diffusive signals (e.g. nitric oxide) or recurrent wiring patterns suffice to enable our model to integrate any new information with the internally represented prior, thereby implementing a Markov Chain Monte Carlo sampling process which re-

flects inferred posterior distributions. Our model simulations shows a population of 20 stochastic neurons and demonstrate the 3 above features, by a) learning sensory tuning curves for the population in good agreement to theoretically derived optimal ones (R2 value >0.9 for density, weight of tuning curves), b) generating samples from learned prior distributions without sensory information and c) correctly computing posterior distributions with incoming sensory information (KL-divergence between model and analytically distributions <0.001). Specifically we tried a broad range of sensory input distributions from Gaussian, uniform to complex bi-modal distributions, achieving consistent results. In achieving these global behaviours microscopic noise, that represents a fundamental problem for information processing in brain, plays an unexpected constructing role, as it allows to get in correspondence with the statistical properties of the environment optimally.

II-6. Adaptive estimation of firing rate maps under super-Poisson variability

Mijung Park¹
 J. Patrick Weller
 Gregory Horwitz²
 Jonathan W Pillow¹

MJPARK@MAIL.UTEXAS.EDU
 JPWELLER@UW.EDU
 GHORWITZ@U.WASHINGTON.EDU
 PILLOW@MAIL.UTEXAS.EDU

¹University of Texas at Austin

²University of Washington

Adaptive stimulus selection in closed-loop neurophysiology experiments can significantly speed up the estimation of a neuron's stimulus tuning. Recent work on this problem has focused on the development of optimal stimulus selection methods based on a Poisson response model and a Gaussian Process (GP) prior over the firing rate map. Although this approach offers substantial improvements relative to randomized designs, it may perform poorly if data do not conform to model assumptions. In visual cortex, spike counts are often over-dispersed relative to the Poisson distribution, meaning that the ratio between the spike count mean and variance is greater than 1. In simulation, we find that firing rate maps estimated from over-dispersed data under a GP-Poisson model exhibit substantial undersmoothing, meaning that the inferred map is much rougher than the true map. Here we introduce a method for adaptive stimulus selection using a model that incorporates (and infers) the degree of overdispersion in neural spike responses. We model overdispersion by extending the GP-Poisson model to include an additional latent Gaussian noise source: responses are modeled as Poisson conditioned on the map plus Gaussian noise. We have applied our method to estimate the color tuning of V1 neurons recorded from awake, fixating monkeys. Stimuli were drifting bars that varied across trials over a 2-D color space that is specified by contrast to long- and medium- wavelength sensitive cones. In interleaved trials, individual V1 neurons were tested with stimuli that were chosen either adaptively or non-adaptively. In non-adaptive trials, stimulus selection was independent of responses. In adaptive trials, stimulus selection was based on the posterior variance of firing rate map under the overdispersed GP-Poisson model. Results showed that the adaptive method yielded faster convergence and more accurate firing map estimates than the non-adaptive method.

II-7. Introducing MEDAL: a Matlab Environment for Deep Architecture Learning

Dustin Stansbury
 Jack Gallant
 University of California, Berkeley

STAN_S_BURY@BERKELEY.EDU
 GALLANT@BERKELEY.EDU

Recent advances in machine learning have sparked enthusiasm for algorithms and model architectures that learn to transform data through multiple layers of nonlinearities. These techniques, known as deep architecture learning (DAL), have provided state-of-the-art performance on a array of complex tasks in computer vision, audio processing, speech recognition, natural language processing, robotics, and information retrieval. The hierarchical and

nonlinear properties of DAL frameworks also offer the neuroscience community with promising models of neural processing within the sensory and perceptual hierarchies of the brain. However, implementation and application of DAL models for general problems have been difficult for neuroscientists that lack extensive experience in machine learning. To facilitate the use of DAL methods we developed the Matlab? Environment for Deep Architecture Learning (MEDAL). This software suite we are introducing provides simple object-oriented implementations and demonstrations of many recently-developed DAL models. Included in the suite are implementations of Bernoulli- and Gaussian-input Restricted Boltzmann Machines (RBMs), Factored 3-Way RBMs, Mean-Covariance RBMs, Conditional RBMs, Convolutional RBMs, Deep Belief Networks (DBNs), Convolutional DBNs, stacked, denoising, and sparse Autoencoders, Multilayer Neural Networks (MLNNs), and Convolutional MLNNs. We demonstrate the functionality of the software suite by replicating and extending a number of previous results based on each model implementation. We further demonstrate the utility of DAL for deriving features used in receptive field models of the primate visual system.

II-8. Encoding the direction of ego-motion in the fly brain

James Trousdale

Sam Carroll

Kresimir Josic

University of Houston

JRTRUSD@MATH.UH.EDU

SRCARROLL314@GMAIL.COM

JOSIC@MATH.UH.EDU

The lobula plate of the blow fly *Calliphora vicina* is a system of approximately 60 motion-sensitive neurons. Both sides of the blow fly brain contain mirror copies of this system. The Vertical System (VS) cells form a subset of the lobula plate consisting of 10 non-spiking neurons per side. These cells are electrically coupled, and code for the azimuth of the axis of rotation of the fly during flight. Despite a plethora of experimental information, theoretical studies of the coding properties of this network have been somewhat limited. Mathematical and computational studies have been restricted primarily to understanding how the angle of rotation is encoded in the network response in steady-state. We extend previous studies, examining the fundamental role that coupling plays in shaping the transient responses to ego-motion of the fly, focusing on how coupling between VS assists in the encoding of the rotational azimuth. We employ a carefully calibrated model of part of the lobula plate system in studying the response of the VS population to different optic flow patterns. In particular, we manipulate the statistics of images of randomly distributed bars projected to and rotated on the surface of a sphere. We find that at steady-state, under the Optimal Linear Estimator, the mean-square error (MSE) of the angle of rotation estimated from time-integrals of the potential is approximately independent of coupling. We then consider a more general encoding scheme to assess the information the population response of the VS cells contains about the stimulus while the system is in the more biophysically relevant transient state. Our results indicate that correlations are integral to the ability of the population to accurately encode the stimulus when images are sufficiently sparse. We explore the role of coupling and correlations in the VS population in relation to the input statistics.

II-9. Representing correlated retinal population activity with Restricted Boltzmann Machines

David Schwab¹

Kristina Simmons²

Jason Prentice¹

Vijay Balasubramanian²

¹Princeton University

²University of Pennsylvania

DSCHWAB@PRINCETON.EDU

SIMMK@MAIL.MED.UPENN.EDU

JASONSP@PRINCETON.EDU

VIJAY@PHYSICS.UPENN.EDU

It is challenging to discover and model patterns of correlation in large-scale neural recordings because the number

of possible correlated subsets grows exponentially with the number of neurons. Previous approaches (e.g. Ising models and Generalized Linear Models) attempt to capture correlated activity by positing that pairwise correlations are dominant. We measured pairwise correlations between populations of ~40 retinal ganglion cells (GCs) responding to stimulus ensembles with widely differing spatio-temporal noise structures (white, exponentially-correlated, scale-free, full-field flicker) and found surprisingly invariant pairwise correlations. This suggested that higher-order correlations may be an important channel for useful population information. To capture such correlations we modeled the population response using Restricted Boltzmann Machines (RBMs), which use hidden units to parsimoniously reproduce arbitrary-order correlations. We trained an RBM with a small number of hidden units (~15% of the number of GCs) on each stimulus ensemble. The model accurately reproduced second- and third-order correlations, as well as pattern probabilities, requiring a fraction (<15%) of the number of parameters in previous population models. By marginalizing over the hidden units, we arrived at a population representation containing many different orders of interaction, finding that the RBM relies heavily on higher-order terms. The small number of hidden units also provides a compressed representation of neuronal activity – for each GC activity pattern, we computed the most likely arrangement of hidden activations. This low-dimensional representation showed subtle dynamical adaptation in the neural population over ~10 minutes following a change in input statistics. Finally, we trained an additional decision layer to discriminate between neuronal responses to different stimulus ensembles, demonstrating that these techniques indeed extract information relevant for stimulus decoding. Our approach is broadly useful for learning population codes when the relevant interaction orders are unknown, especially when the correlation structure is largely due to common input.

II-10. Dendritic nonlinearities shape beyond-pairwise correlations and improve (spiking) population coding

Joel Zylberberg
N. Alex Cayco-Gajic
Eric Shea-Brown

University of Washington

JOELZY@UW.EDU
CAYCOGAJIC@GMAIL.COM
ETSB@WASHINGTON.EDU

Neural populations often show correlated activity, with co-active neural pairs, triplets, etc. occurring more frequently than would be expected from the lower-order statistics alone – means, covariances, etc., respectively. Much research has investigated the impact of pairwise correlations on coding performance in neural populations. Recent experiments with relatively large neural populations, however, show significant higher-order correlations (HOC: beyond pair-wise): the data are poorly fit by pair-wise maximum entropy models, but well-fit by higher-order models. We seek to understand how HOC are shaped by the properties of neural networks and of the neurons therein, and how these HOC affect population coding. In our presentation, we will first demonstrate that dendritic non-linearities similar to those observed by Polsky et al (2004) are equivalent to beyond-pairwise interactions in a spin-glass-type (maximum entropy) statistical model: they can either increase, or decrease, the magnitude of the HOC relative to the pair-wise correlations. We will then discuss our studies of a population coding model (which generalizes that of Tkacik et al. (2010)) with parameterized pairwise- and higher-order interactions. These studies reveal the conditions under which the beyond-pairwise interactions (dendritic nonlinearities) can increase the mutual information between a given set of stimuli, and the (noisy) population spiking responses. When the stimuli are jointly Gaussian, coding performance can be improved by modifying the output HOC via dendritic nonlinearities, if the neurons have low firing rates. Nonlinearities improve coding over a broader range of firing rates for skewed stimulus distributions, like the distribution of luminance values in natural images, so long as the neurons receive correlated inputs (for example, due to overlapping receptive fields). Normative theories might therefore predict differences in the dendritic summation properties in neural populations with different mean firing rates, or ones that encode stimuli drawn from different distributions.

II-11. The perils of inferring information from correlations

Jeffrey Beck¹
 Ingmar Kanitscheider²
 Xaq S Pitkow¹
 Peter Latham³
 Alexandre Pouget²

JBECK@BCS.ROCHESTER.EDU
 INGMAR.KANITSCHIEDER@UNIGE.CH
 XPITKOW@BCS.ROCHESTER.EDU
 PEL@GATSBY.UCL.AC.UK
 ALEX@CVS.ROCHESTER.EDU

¹University of Rochester

²University of Geneva

³University College, London

Estimating the information in neural activity is a key step toward understanding how neural activity controls behavior. Fisher information in a population code depends critically on noise correlations, in particular on the noise correlations that cannot be distinguished from the signal. For population codes in which the activity of the i th neuron in response to a stimulus, s , is $f_i(s)$, the signal is the vector $f'(s)$. For such codes, information is limited only by noise aligned with the signal; that is, noise in the f' direction; what we call $f'f'^T$ correlations. Experimentally, it seems that correlations of this $f'f'^T$ form are not present in neural activity. Instead correlations seem to depend mostly on the difference in preferred stimuli. This implies that large populations have much more information than is observed in animal behavior. This seems rather unlikely; it is much more likely that information is limited in the sense that it saturates as the number of neurons increases. When this is the case, we show that correlations must be present, but may be quite small. As a result they can be masked by other correlations which do not limit information. Consequently, the standard approach — measuring a few pairwise correlations, using regularization techniques to fill in the missing entries, and then estimate information — is likely to give very wrong answers. This is because, in high dimensions, even very small errors in estimating either f' or the correlations can lead to bad estimates of information. Fortunately, there is a better approach: forget about the tuning curves and the correlations, record simultaneously from a large population, and directly decode population activity. Decoding accuracy can then be easily translated into Fisher information.

II-12. Temporal decorrelation by power-law adaptation in pyramidal neurons

Christian Pozzorini¹
 Richard Naud²
 Skander Mensi^{3,4}
 Wulfram Gerstner¹

CHRISTIAN.POZZORINI@EPFL.CH
 RNAUD@UOTTAWA.CA
 SKANDER.MENSI@EPFL.CH
 WULFRAM.GERSTNER@EPFL.CH

¹EPFL, BMI

²University of Ottawa

³EPFL

⁴BMI

Firing rate adaptation is ubiquitous in the cortex, however its functional role remains unclear. In pyramidal neurons, adaptation occurs on multiple timescales ranging from milliseconds to seconds. Combining in-vitro recordings with single neuron modeling, we found that spike-frequency adaptation is due to both a spike-triggered adaptation current and a spike-triggered movement of the firing threshold. Our results show that both adaptation processes last for more than 20 seconds and decay over multiple time scales following a power-law. Using in-vivo intracellular recordings from awake mice, we then demonstrate that the power-law decay of the spike-triggered adaptation mirrors and cancels the temporal correlations of input current received at the soma of somatosensory pyramidal neurons. These findings indicate that spike-frequency adaptation causes temporal decorrelation of output spikes, an efficient coding procedure that, at high signal-to-noise ratio, maximizes the information transfer.

II-13. Theory of higher-order correlations in neural network with columnar structure

Yasuhiko Igarashi¹
Masato Okada²

IGAYASU1219@MNS.K.U-TOKYO.AC.JP
OKADA@K.U-TOKYO.AC.JP

¹University of Tokyo

²Graduate School of Frontier Sciences, The Uni

Dependencies among cells determine the detailed nature of the code of neural populations. To understand the complex relationship between the structure of a neural network and the code of neural populations, the use of multi-electrode neural population recordings have become common. Neurophysiological experiments have recently shown that not only pair-wise correlated neural activities but also higher-order correlated patterns of activity are often observed in the brain by sequences of action potentials of neural populations [Ohiorhenuan et al. 2010, 2011; Ganmor et al. 2011]. However, very little is theoretically known about the relationship in a network of structural connections linking sets of neurons. We constructed a theory on the origin of structured higher-order neural activities in a minimum network model that can elucidate experimental observations. We particularly focus on the comparison between our theoretical results and the electrophysiological experiment reported by Ohiorhenuan et al. involving the primary visual cortex (V1) [Ohiorhenuan et al. 2010, 2011]. Unlike a homogeneous network [Amari et al. 2003; Macke et al. 2011], a network with columnar structure can provide not only the tuning curve of firing rates but also the relationship between higher-order correlations. Although in the homogeneous feedforward network where higher-correlations among neurons are uniform regardless of external stimulus, it is believed that visual stimulation clearly reorganizes the activity of structured V1 circuits by preferentially activating V1 neurons. We therefore calculated a triplet correlation among V1 neurons derived by light bars. We found that the heterogeneous structure can dynamically control the structure of higher-order correlations and generate both sparse and synchronized neural activity, which has been observed in neurophysiological experiments [Yu et al. 2010]. We expect our study to promote theoretical studies on how structured interaction affects higher-order correlated neural activity and information processing in the brain.

II-14. Perceptual decisions are limited primarily by variability in early sensory cortex

Charles Michelson
Jonathan W Pillow
Eyal Seidemann

CHUCKMICHELSON@UTEXAS.EDU
PILLOW@MAIL.UTEXAS.EDU
EYAL@MAIL.CPS.UTEXAS.EDU

University of Texas at Austin

Understanding how perceptual decisions are formed in the brain is a long-standing goal of systems neuroscience. Previous research in monkeys has found neural activity that correlates with decisions as early as primary sensory cortex, but these single-neuron recordings show only weak correlations with choice. This weak relationship implies that single sensory cortical neurons capture only a small fraction of the choice-related variability: such variability may be distributed across a large population of cells in sensory cortex, and may also occur in other areas downstream from early sensory cortex. Here we use a combination of direct measurement of neural population responses in primary visual cortex (V1) by voltage-sensitive dye imaging (VSDI) and a simple computational model to show that V1 encompasses most of the choice-related variability in highly trained monkeys engaged in a demanding visual pattern detection task. Our results are consistent with previously reported observations of weak single-neuron choice-related signals, and weak inter-neuronal correlations. Overall, our results suggest that simple perceptual decisions are formed by pooling information across a large population of weakly correlated sensory neurons, and that most choice-related variability is already present at the earliest stages of sensory cortical processing.

II-15. Different response properties of rat parietal and frontal cortices during evidence accumulation

Timothy Hanks^{1,2}
Chunyu Ann Duan^{2,1}
Jeffrey Erlich^{2,1}
Bingni Brunton^{2,1}
Carlos Brody^{2,1}

¹HHMI

²Princeton University

THANKS@PRINCETON.EDU
CDUAN@PRINCETON.EDU
JERLICH@PRINCETON.EDU
B.W.BRUNTON@GMAIL.COM
BRODY@PRINCETON.EDU

Gradual accumulation of evidence is thought to be a fundamental component of decision-making. Similar neural correlates of evidence accumulation have been found in both parietal and frontal cortices. As a result of this similarity, it has not yet been possible to establish different roles for parietal and frontal cortices in evidence accumulation. Here we report new electrophysiological recordings from posterior parietal cortex (PPC) and frontal orienting fields (FOF) of rats performing a two-alternative evidence accumulation task. Consistent with primate studies, we find that in both areas, neural firing rates gradually change over time in this type of task, with the rate of change depending on the strength of the evidence. However, despite the apparent parietal/frontal response similarity, a new finer-grained quantitative analysis method reveals that different types of information are encoded in the two areas: PPC represents the graded value of the accumulated evidence, while FOF has a binary representation, encoding at each timepoint which of the two decision alternatives is the best choice (given the evidence so far). Our results suggest that circuits involved in the accumulation of evidence are not composed of diffusely distributed brain areas performing the same function, but that instead, each brain area has a specific and identifiable role in the decision-making process.

II-16. Orbitofrontal cortex is sensitive to natural behavioral categories during social exploration

Geoffrey Adams
John Pearson
Michael Platt
Duke University

GEOFFREY.K.ADAMS@DUKE.EDU
PEARSON@NEURO.DUKE.EDU
PLATT@NEURO.DUKE.EDU

Orbitofrontal cortex (OFC) is a multimodal area implicated in value-based decision-making as well as normal social behavior. Its strong connections with inferior temporal cortex, the basolateral amygdaloid nuclei, the hippocampal formation, and other prefrontal cortical areas make it well situated to receive and interpret social information from the environment according to its biological function to guide immediate investigative behaviors as well as future social decision-making. To understand the role that OFC plays in adaptive exploration of the social environment, we recorded single units in rhesus monkey OFC while the monkeys performed a video selection and viewing task. Several hours of video of spontaneous behaviors produced by unfamiliar free-ranging rhesus monkeys comprised the video database, from which the subject monkeys viewed a series of five-second clips. OFC units' firing rates were strongly modulated by the presentation of movies of unfamiliar conspecifics. Individual units also exhibited specific patterns of modulation to various natural behavioral categories presented in the videos.

II-17. Humans exploit the uncertainty in priors to improve direction perception.

Steeve Laquittaine¹
Justin Gardner²

STEEVE_LAQUITAIN@MSN.COM
JUSTIN@BRAIN.RIKEN.JP

¹Riken BSI

²Gardner Research Unit

How well can humans exploit prior information when estimating perceptual quantities like the direction of motion? If one knows in advance the most likely direction, then biasing perception toward that prior when sensory evidence is weak is a good strategy. While humans have been shown to use priors in a variety of different tasks, less is known about how accurately priors are represented. For instance, subjects may bias sensory estimates towards the best prior estimate (mean of the prior), but ignore uncertainty in the prior (variance of the prior). Alternatively, subjects may use both the best prior estimate and the level of uncertainty in that prior estimate to bias sensory decisions — a strategy consistent with optimal Bayesian inference. We used a motion estimation task in which we could independently vary the level of sensory evidence by changing the coherence of the motion and the uncertainty in the prior by manipulating the set of displayed directions. By examining subject's estimated directions relative to the displayed directions we could distinguish three different hypotheses: 1) subjects ignored the prior 2) subjects used only the mean of the prior 3) subjects used both the mean and the uncertainty of the prior. Consistent with the last hypothesis, we found that the directions reported were biased toward the prior for weak motion, and that this bias was stronger when the prior was more certain. We therefore conclude that subjects exploit both the mean and the uncertainty in their prior to improve their perception of motion directions. Our results support the view that the representation of priors by humans is consistent with optimal Bayesian inference.

II-18. Hebbian mechanisms underlying the learning of Markovian sequence probabilities

Kristofer Bouchard^{1,2}
Surya Ganguli³
Michael Brainard¹

KRIS@PHY.UCSF.EDU
SGANGULI@STANFORD.EDU
MSB@PHY.UCSF.EDU

¹University of California, San Francisco

²UCB

³Stanford University

Most natural sensorimotor events do not occur in isolation, but are embedded in sequences with rich probabilistic structure. Despite the ubiquity of such learned probabilistic sequences, little is known about neural mechanisms allowing the statistical structure of sequential experience to be embedded within patterns of synaptic weights. Here, we investigate conditions under which Hebbian synaptic plasticity sculpts unstructured networks to quantitatively reflect the conditional probability of sequential events in their synapses. We show through analytics and simulations that Hebbian plasticity with presynaptic competition develops synaptic weights proportional to conditional forward transition probabilities present in the input sequence, and are thus appropriate for sequence generation. In contrast, postsynaptic competition develops weights proportional to the conditional backward probabilities $P(s(t-1)|s(t))$, which interestingly are reflected in auditory responses of Bengalese finch (Bf) song circuitry. We demonstrate that to stably but flexibly reflect the conditional probability of a neuron's inputs and outputs, local Hebbian plasticity should approximately balance the magnitude of synaptic depression relative to potentiation (a competitive force that triggers weight differentiation) with the weight dependence of synaptic change (a homogenizing force that stabilizes weights). These forces control the rate at which structure is learned and the entropy of the final distribution of synaptic weights. Thus their relative balance induces a prior over learnable transition distributions. For a range of balances, we find robust sequence learning, including the learning of probabilistic syllable sequences generated by Bfs. Together, these results demonstrate remarkably simple correspondences between biophysics and probabilistic sequence learning: the site of synaptic competition dictates the temporal

flow of learned probabilistic structures and the balance between competitive and homogenizing forces dictates a prior expectation of the randomness of the sequence to be learned. This yields a novel mechanism for priors over sequence distributions to be embedded within synaptic biophysics.

II-19. Controllability and resource-rational planning

Falk Lieder

FALK.LIEDER@GMAIL.COM

University of Zurich

Learned helplessness experiments involving controllable vs. uncontrollable stressors have shown that the perceived ability to control events has profound consequences for decision making. Normative models of decision making, however, do not naturally incorporate knowledge about controllability, and previous approaches to incorporating it have led to solutions with biologically implausible computational demands [1,2]. Intuitively, controllability bounds the differential rewards for choosing one strategy over another, and therefore believing that the environment is uncontrollable should reduce one's willingness to invest time and effort into choosing between options. Here, we offer a normative, resource-rational account of the role of controllability in trading mental effort for expected gain. In this view, the brain not only faces the task of solving Markov decision problems (MDPs), but it also has to optimally allocate its finite computational resources to solve them efficiently. This joint problem can itself be cast as a MDP [3], and its optimal solution respects computational constraints by design. We start with an analytic characterisation of the influence of controllability on the use of computational resources. We then replicate previous results on the effects of controllability on the differential value of exploration vs. exploitation, showing that these are also seen in a cognitively plausible regime of computational complexity. Third, we find that controllability makes computation valuable, so that it is worth investing more mental effort the higher the subjective controllability. Fourth, we show that in this model the perceived lack of control (helplessness) replicates empirical findings [4] whereby patients with major depressive disorder are less likely to repeat a choice that led to a reward, or to avoid a choice that led to a loss. Finally, the model makes empirically testable predictions about the relationship between reaction time and helplessness.

II-20. Neural responses in the rat parietal cortex during decision formation and movement

David Raposo^{1,2}

Matthew Kaufman¹

Anne Churchland¹

DRAPOSO@CSHL.EDU

MKAUFMAN@CSHL.EDU

CHURCHLAND@CSHL.EDU

¹Cold Spring Harbor Laboratory

²Champlain Neuroscience Programme

During decisions about random dot motion in monkeys, many neurons in the posterior parietal cortex (PPC) have 'congruent' responses during decision formation and movement: they tend to have similar direction selectivity throughout a trial. It is not known whether this congruence is a standard feature of decision-making circuitry, or whether it is task dependent. We developed a two-alternative forced-choice task in which subjects are presented with a stream of visual and/or auditory events and are asked to determine if the event rate is high or low relative to an experimenter-defined boundary (Raposo et al. 2012). We recorded 186 single neurons from the PPC of 3 rats performing the task. Using ROC analysis we examined the direction selectivity of these neurons in two different epochs in the trial — decision formation and movement. During decision formation, 28% of neurons had significantly different responses for trials ending in high- vs. low-rate choices. Data collected in correct and error trials suggest that these neurons encode the future choice direction. As in monkeys, the ability of these neurons to distinguish high- from low-rate choices typically built up for hundreds of milliseconds, suggesting accumulation of evidence toward a decision. During movement, 66% of neurons were strongly direction selective. To evaluate the

degree of congruence over the course of a trial, we compared the direction selectivity during decision formation and movement. Critically, we observed that many neurons were incongruent: their direction selectivity changed around movement time. This result differs from what is frequently reported in decision-making studies. Our hypothesis is that the difference stems from our task design: unlike most perceptual decision-making studies, the nature of the stimulus (high or low rate) is unrelated to the eventual response (left or right movement).

II-21. Do humans account for stimulus correlations in visual perception?

Manisha Bhardwaj¹
 Ronald Van den Berg²
 Wei Ji Ma²
 Kresimir Josic¹

MANISHA@MATH.UH.EDU
 RB2@BCM.EDU
 WJMA@BCM.EDU
 JOSIC@MATH.UH.EDU

¹University of Houston

²Baylor College of Medicine

Bayesian optimality is a general principle that has been successfully used to build quantitative models of perception. However, it is unclear whether Bayesian models can adequately describe perception in natural scenes containing multiple objects. A number of recent studies have taken a first step towards addressing this question in simplified scenes containing multiple objects. Human behavior was found to be close to Bayes-optimal, for example in visual search [1], sameness judgment [2], and change detection [3]. However, these studies did not introduce any higher-order structure among the stimuli (but see [4]). For instance, distractor stimuli in the visual search task were chosen independently of one another. By contrast, visual stimuli in natural scenes possess a complex structure. It is therefore important to examine how visual perception is affected by structured input. We examined this question in a target detection task. Subjects reported whether a vertical target stimulus was present in a set of four stimuli. Distractor orientations were chosen from a multivariate normal distribution. We manipulated the amount of structure in the scene by varying the strength of the correlation between distractor orientations across experimental sessions. We fitted two Bayesian models, which differed only in how observers took into account distractor correlations: in Model 1 ('optimal'), the observer used the correct correlation strengths, whereas in Model 2 ('suboptimal'), the observer assumed potentially different values. We found that subjects were suboptimal in the sense that they did not take into account the correct correlations. They overestimated low correlations, but performed nearly optimally when distractors were perfectly correlated. This suggests that subjects perceive structure in visual scenes even if there is none.

II-22. Neural correlates of arbitration between model-based and model-free reinforcement learning systems

Sang Wan Lee
 Shinsuke Shimajo
 John O'Doherty

SWLEE@CALTECH.EDU
 SSHIMOJO@CALTECH.EDU
 JDOHERTY@CALTECH.EDU

California Institute of Technology

Human and animal action-selection is known to be mediated by at least two distinct mechanisms, a goal-directed and a habitual mechanism, which on the computational level can be implemented via divergent forms of reinforcement-learning: a model-based and a model-free system. Although there is considerable evidence for the existence of those computations, a fundamental question that has remained as yet unaddressed is how control is passed between model-based and model-free learning systems, and where in the brain adjudication over which system directs behavior is mediated. Here we formulate a model in which a computation about the degree of reliability of the model-based system is used to arbitrate between the two learning strategies. For this, a Bayesian computation is made over the probability of the state/reward prediction error being set to zero, and this signal is fed

into a biophysical two-state transition model that allocates control to the model-based and model-free systems in a manner proportional to the degree of reliability. We used a sequential decision paradigm in which the structure of the decision task was changed at different points in the experiment in order to optimally favor model-based or model-free control. Behavioral analysis indicates that our computational model successfully predicted when and how subjects' behavior is shaped by the two control processes in terms of the degree of behavioral sensitivity exhibited to task structure changes. Subsequent functional neuroimaging analyses indicated key roles for a region of dorsomedial frontal cortex and inferior lateral prefrontal cortex in encoding beliefs and uncertainty about the reliability of the model-based system. Our study therefore reveals evidence for the first time of a neural arbitration signal between model-based and model-free systems in some cortical brain areas, suggesting these regions may play a fundamental role in switching control between model-based and model-free reinforcement-learning.

II-23. The neural correlates of counterfactual-Q-learning in a strategic sequential investment task

Rong Guo^{1,2}
Michael Tobia³
Wendelin Böhmer¹
Tobias Sommer³
Christian Büchel³
Klaus Obermayer¹

RONG@NI.TU-BERLIN.DE
M.TOBIA@UKE.DE
WENDELIN@NI.TU-BERLIN.DE
TSOMMER@UKE.UNI-HAMBURG.DE
BUECHEL@UKE.UNI-HAMBURG.DE
KLAUS.OBERMAYER@MAILBOX.TU-BERLIN.DE

¹Technische Universität Berlin

²BCCN Berlin

³Universitätsklinikum Hamburg-Eppendorf

Comparison of factual and counterfactual outcomes associated with selected and unselected actions results in a fictive prediction error (FPE). Reinforcement learning models do not incorporate this fictive error signal although there is neural evidence that people do indeed take counterfactual outcomes into account when making subsequent decisions. The purpose of this experiment was to test a reinforcement learning model with and without fictive error signals. Q-learning models update the value of a state-action pair according to a temporal difference (TD) computed from observed outcomes and anticipated future rewards. These models have been shown to fit both human and animal choice behavior. We extended the standard Q-learning model by incorporating both counterfactual gains (rejoice) and losses (regret) as potential error signals. This FPE-enhanced-Q model updates the values of both the selected and unselected actions using counterfactual outcome information that is not used in the standard Q-learning model. fMRI data were collected while 30 healthy participants (ages 18-30) completed 80 rounds of a strategic sequential investment task. A round of trials consisted of three investment decisions, each associated with a state defined by a unique expected reward magnitude and win probability. The action of each investment, rather than the outcome of the trial, determined the transition to each subsequent state. Models were fitted to individual subject data by maximum likelihood, and then evaluated by contrasting BIC scores. The FPE-enhanced-Q model fits the data significantly better than standard Q-learning model. In addition, the model-based Q, TD and FPE time series were significantly correlated with fMRI activity in the striatum and orbitofrontal cortex. This is the first demonstration of a significant relation between model derived Q-values and neural activity. In conclusion, incorporating fictive prediction error signals in Q-learning improves the learning fit and can be mapped to the brain.

II-24. Dynamics of decision and action in rat posterior parietal cortex

Matthew Kaufman¹
David Raposo^{1,2}
Anne Churchland¹

MKAUFMAN@CSHL.EDU
DRAPOSO@CSHL.EDU
CHURCHLAND@CSHL.EDU

¹Cold Spring Harbor Laboratory

²Champlain Neuroscience Programme

Activity in posterior parietal cortex (PPC) has been extensively studied in the context of decision making. PPC activity reflects accumulated evidence needed to make a choice, and this activity has generally been assumed to be a one-dimensional representation of the decision variable. However, the dimensionality of such responses is rarely computed, and thus the possibility remains that responses are higher-dimensional than typically thought. Here, we take a dynamical systems approach to understanding this activity. We recorded in PPC of three rats (48-69 isolated single units/animal using 8 tetrodes) during a multisensory rate-judgment task with an arbitrary association of stimulus rate and response direction. We found that activity patterns were substantially more complex than would be expected for a simple integrator: the dimensionality of PPC activity (assessed using PCA) ranged from ~3-7 during the stimulus epoch, and ~8-12 when movement was considered as well. Moreover, using a novel mathematical technique, we found that neurons achieved different patterns of activity during stimulus presentation versus during movement. Specifically, we found that neural dimensions containing high variance during the stimulus epoch contained relatively little variance during the movement, and vice versa. However, this does not mean that the relationship between decision-related and movement-related activity is arbitrary. We applied jPCA, a recent technique for finding rotational patterns in high-dimensional data. We found that rotations were strongly present during movement, in somewhat a similar manner as in monkey motor cortex. These findings together imply that rat PPC is likely acting as more than a simple one-dimensional integrator during the stimulus epoch, and explores more patterns still when movement must be produced. Thus, PPC appears to also be involved in translating the decision into action.

II-25. Normalization predicted by optimal inference in sensory systems

Matthew Chalk¹
Paul Masset^{2,3}
Boris Gutkin
Sophie Deneve

MATTHEWJCHALK@GMAIL.COM
PAUL.MASSET@GMAIL.COM
BORIS.GUTKIN@GMAIL.COM
SOPHIE.DENEVE@ENS.FR

¹École Normale Supérieure

²Watson School of Biological Sciences

³Cold Spring Harbor Laboratory

The responses of sensory neurons are non-linear. For example, they typically saturate, and are modulated by the spatiotemporal context of presented stimuli. 'Normalization', in which neural responses encode the ratio between their feed-forward input and the summed activity of nearby neurons, provides a common framework to explain these non-linear response properties across many neural systems and organisms. Here, we show that normalization is a fundamental consequence of a system that performs optimal estimation of presented stimulus features given sensory noise. The trial-by-trial variance in neural firing rates typically scales proportionally with the mean. Given this signal-dependent noise, we show that optimal estimation of presented stimulus features results in normalization, with each neuron encoding the ratio between its received and predicted input. In contrast, previous 'predictive coding' models, in which each neuron encodes the difference between its received and predicted input, assume constant sensory noise. We show that optimal estimation with signal-dependent noise can account for several aspects of sensory neural responses not explained by these previous models, including response saturation, contextual shifts in contrast response curves, and the stimulus-dependent time-course of neural responses. We extend our model to perform inference on a hierarchy of features. Neurons at each level of the network estimate increasingly complex stimulus features, based on inputs from the previous level. Thus, nor-

malization of responses at one level of the network reshapes the stimulus selectivity of higher-level neurons. As a result, neural receptive fields (RFs) are dynamically reshaped by contextual stimuli that do not elicit a response when presented alone. While RFs vary depending on the stimuli used to measure them, an iterative method can efficiently recover the neuron's 'true' selectivity: the invariant stimulus feature that it encodes.

II-26. The synaptic sampling hypothesis

Laurence Aitchison
Peter Latham
University College, London

LAURENCE@GATSBY.UCL.AC.UK
PEL@GATSBY.UCL.AC.UK

There is strong evidence that the brain infers a probability distribution over the external world's state given sensory input; the sampling hypothesis states that neural activity samples from this distribution. The sampling process provides a possible explanation for the highly variable activity patterns observed in the brain. Here we apply the same principle to synaptic weights. We hypothesise that neurons infer a probability distribution over the ideal setting of their weights, and that the neuron's actual, physical synaptic weights are sampled from this distribution. Importantly, the weight is sampled each time there is an incoming spike, providing a possible explanation for the noisy and unreliable synaptic transmission observed in the brain. To develop quantitative and testable predictions, we build a model based loosely on the cerebellum, in which a Purkinje cell adjusts its weights given parallel fiber input and a climbing fiber error signal. Because the world is nonstationary, we choose the prior to allow the true weights to change over time; we use an exponentiated Ornstein Uhlenbeck process, as it gives rise to log-normal weight frequency distributions, and is consistent with spine size changes (Y. Loewenstein, A. Kuras and S. Rumpel. 2011. J. Neuroscience). We then derive a neurally plausible approximation to exact inference of the distribution over weights. This allows us to make predictions about how the mean weight and the noise level should behave at both steady state and during LTP/LTD. The set of predictions is novel, falsifiable, and specific enough that experimental corroboration would go a long way toward supporting our hypothesis.

II-27. Information-theoretic limits on encoding over diverse populations

O. Ozan Koyluoglu
Ila Fiete
University of Texas at Austin

OZAN@MAIL.UTEXAS.EDU
ILAFIETE@MAIL.CLM.UTEXAS.EDU

Population coding refers to a setting where a given stimulus is represented by the activities of a population of neurons. For instance, in orientation-tuned V1 neurons, each neuron fires near its preferred stimulus, with an activity profile given by the tuning curve. When combined with an estimator, these activities constitute a fully identified coding system in which the efficiency of the system is quantified by a measure of distortion (error) between the estimated stimulus and its actual value. Here, we use an information-theoretic approach to bound distortion (in a mean-square sense) for populations of neurons: a stimulus s is first sent to an encoder, that computes a vector-valued function of the stimulus, and each entry of the vector is then represented in a separate population code. We assume the total number of neurons is fixed at Q , and that the Fisher information in each neural population scales linearly with the number of neurons in that population (as seen for unimodal tuning curves with Poisson spike variability, among various examples). We consider two scenarios: The encoder simply passes out s , to one population of Q total neurons; or, it passes out elements of the N -d vector $x(s)$, to N populations of $M=Q/N$ neurons each. For these scenarios, we use joint source-channel coding theory to bound how the information-theoretically minimal distortion will scale with M, N . We show that breaking the neurons into N populations can, with appropriate encoding, result in distortions that scale as $M^{-(N)}$, whereas directly representing the stimulus in a single population, by necessity, produces distortions that scale as $1/(MN)$. Our results show that diverse population encoding can result in potentially much lower distortion, and quantify how distortion scales with number of populations.

II-28. Matching encoding and decoding with spiking neurons

Fleur Zeldenrust
 Sophie Deneve
 Boris Gutkin
 École Normale Supérieure

FLEURZELDENRUST@GMAIL.COM
 SOPHIE.DENEVE@ENS.FR
 BORIS.GUTKIN@GMAIL.COM

The Generalized Linear Model (GLM) is a powerful tool in assessing neural spike responses ([Pillowetal2008]; for an overview, see [Pillow2007]). The model assumes that the output of a neuron is an inhomogeneous Poisson process, of which the instantaneous rate is given by a thresholded sum of the linearly filtered input and output. It can incorporate effectively both the neuron's receptive field and history-dependent effects such as the refractory period and spike-frequency adaptation [Paninski2004], [Truccoloetal2005]. While the GLM is a descriptive model of how neurons respond to their input, we show how it can be used to unify encoding (how a neuron represents its input in its output spike train) and decoding (how the input can be reconstructed from the output spike train) properties. We analytically derive a GLM that can be interpreted as a recurrent network of neurons that optimally tracks a continuously varying input. In this GLM, every neuron only fires a spike if this reduces the mean-squared error between the received input and a prediction of the input based on the output spike trains of the network, implementing a form of Lewicki's 'matching pursuit' [5]. Contrary to the standard GLM, where input and output filters are independently fitted to a neuron's response, here the filters have a direct interpretation. This theory predicts that the feature the neuron represents directly determines its input and the output filters. Moreover, the representing feature determines the neuron's spike-generating dynamics and its connectivity to other neurons. Therefore, we predict that the encoding and decoding properties of sensory neurons are two sides of the same coin. We use this approach to investigate the coding properties of several types of neurons recorded in in vitro patch clamp experiments.

II-29. Thalamic synchrony drives cortical feature selectivity in standard and novel visual stimuli

Sean Kelly¹
 Jens Kremkow²
 Jianzhong Jin²
 Yushi Wang²
 Qi Wang¹
 Jose-Manuel Alonso²
 Garrett Stanley¹

SKELLY32@GATECH.EDU
 JENS@KREMKOW.INFO
 JJIN@SUNYOPT.EDU
 YUSHIWANG@SUNYOPT.EDU
 QI.WANG@BME.GATECH.EDU
 JALONSO@SUNYOPT.EDU
 GARRETT.STANLEY@BME.GATECH.EDU

¹Georgia Institute of Technology
²SUNY-Optometry

Although a growing body of research has implicated neuronal synchrony as a potentially important part of the neural code, the role of synchrony in the emergence of cortical feature selectivity remains unclear. It has been shown that cortical layer 4 neurons are sensitive to the timing of the projecting thalamic input, but this has only been addressed using the recorded activity of pairs of thalamic neurons due to experimental limitations. Here, we have created a model of cortical feature tuning that relies only on the synchrony of spiking activity among populations of neurons in the cat Lateral Geniculate Nucleus (LGN) in order to quantify the role of input synchrony in the emergence of cortical orientation tuning. Given the limitations of experimental neural recording, we have created a framework that uses our large database of simultaneously recorded neurons to 'fill-in' the unobserved LGN population. This framework uses stimulus geometry to manipulate the recorded response and shift spatial locations, with the resulting populations revealing that the synchrony across the thalamic population was a function of stimulus orientation. We use these populations as input to a large-scale biophysical integrate-and-fire model of the cortical response. Using Fisher Information to determine the ability of the simulated cortical signal to transmit information about stimulus orientation, we show that information efficiency saturates for levels of synchrony on

the order of 20 ms. We extended this finding to a stimulus which combines perspective and observer motion effects with sinusoidal gratings. Preliminary results using recorded LGN inputs to the cortical model indicate that dynamic motion in the stimulus appears to be reflected in the synchrony across the LGN sub-population to which the cortical response is particularly sensitive. Taken together, these results further suggest the role of thalamic synchrony in the emergence of cortical feature selectivity in the visual pathway.

II-30. Encoding and decoding stimuli that generate persistent activity

Kyriaki Sidiropoulou^{1,2}
Panayiota Poirazi²

SIDIROPOULOUKIKI@GMAIL.COM
POIRAZI@IMBB.FORTH.GR

¹University of Crete
²IMBB-FORTH

Proper functioning of working memory involves the expression of stimulus-selective persistent activity in pyramidal neurons of the prefrontal cortex (PFC), which refers to neural activity that persists for seconds beyond the end of the stimulus. The mechanisms which PFC pyramidal neurons use to discriminate between preferred vs. neutral inputs at the cellular level are largely unknown. Here, we use a compartmental modeling approach to search for discriminatory features in the properties of incoming stimuli to a PFC pyramidal neuron and/or its response that signal which of these stimuli will result in persistent activity emergence. We identify synaptic location within the basal dendrites as a feature of stimulus selectivity. Specifically, persistent activity-inducing stimuli consist of activated synapses that are located more distally from the soma compared to non-inducing stimuli, in both model cells. In addition, the action potential (AP) latency and the first few inter-spike-intervals of the neuronal response can be used to reliably detect inducing vs. non-inducing inputs, suggesting a potential mechanism by which downstream neurons can rapidly decode the upcoming emergence of persistent activity. Furthermore, we tested whether these features are subject to modulation by modulation of ionic mechanisms of the PFC model neuron and found that specific mechanisms, such as the N-type calcium current and the slow calcium-activated potassium current modulate the ability of the neuron to decode persistent activity induction. Collectively, our results pinpoint to specific features of the incoming stimulus and neuronal output that encode and decode, respectively, the induction of persistent activity.

II-31. Evaluation of single unit error contribution to neural state space dynamics in linear BMI decoders

Islam Badreldin¹
Karthikeyan Balasubramanian²
Mukta Vaidya²
Joshua Southerland³
Andrew Fagg³
Nicholas Hatsopoulos²
Karim Oweiss¹

BADRELDI@EGR.MSU.EDU
KARTHIKEYANB@UCHICAGO.EDU
MVAIDYA@UCHICAGO.EDU
JS@OU.EDU
FAGG@CS.OU.EDU
NICH@UCHICAGO.EDU
KOWEISS@MSU.EDU

¹Michigan State University
²University of Chicago
³University of Oklahoma

Neural decoders in Brain-Machine Interfaces (BMIs) translate neural activity of ensembles of neurons into control signals to actuate artificial devices. The class of linear BMI decoders, particularly the Wiener and Kalman filters, have been shown to perform well in goal-directed reach and grasp tasks. They require, however, frequent calibration to adjust their coefficients to cope with the non-stationary nature of neural spiking patterns, particularly during long term experiments. In this work, we characterize the long term performance of Wiener decoders during

one dimensional control of a robotic arm performed by two adult male rhesus macaques (*Macaca mulatta*). We introduce a novel error metric computed in the neural state space to assess how well the monkeys learned to use the decoder to produce biomimetic reach kinematics similar to those recorded in humans.

II-32. The geometry of excitation-inhibition balance in human neocortex

Nima Dehghani¹
Adrien Peyrache²
Eric Halgren³
Sydney Cash⁴
Alain Destexhe⁵

NIMA.DEHGHANI@WYSS.HARVARD.EDU
ADRIEN.PEYRACHE@GMAIL.COM
EHALGREN@UCSD.EDU
SCASH@PARTNERS.ORG
DESTEXHE@UNIC.CNRS-GIF.FR

¹Harvard University

²New York University

³University of California, San Diego

⁴Harvard Medical School; MGH

⁵Centre National de la Recherche Scientifique

High-density intracranial recordings were obtained using a 96-electrode 'NeuroPort' silicon array with 400-um spacing, covering an area of 4 x 4 mm, 1mm in length, placed in layers II/III of the middle temporal gyrus from the neocortex of humans during epilepsy monitoring. Up to 90 simultaneously recorded units could be discriminated and were separated between regular-spiking (RS) and fast-spiking (FS) cells based on spike shape. Many pairs of cells showed functional interaction, which allowed to directly prove the inhibitory or excitatory nature of FS and RS cells, respectively (Peyrache et al., PNAS 2012). Periods devoid of epileptic inter-ictal activity were used to assess the differential firing of RS and FS cells during different brain states, such as wakefulness and different phases of sleep (light SWS, deep SWS, REM sleep). We compared the multi scale balance of the two interacting group of cells.....

II-33. State-dependent impact of distinct interneuron types on visual contrast gain

Ulf Knoblich
Jessica Cardin
Yale University

ULF.KNOBLICH@YALE.EDU
JESS.CARDIN@YALE.EDU

Neuronal gain is a measure of neuronal sensitivity to the stimulus and a defining element of the contribution of single neurons to network operations. Previous studies have suggested that neuronal gain contributes to perception of stimulus parameters, such as visual contrast or orientation, by enhancing neuronal response discriminability. Recent work has shown that modulating inhibitory interneuron activity in visual cortex influences visual perception in awake mice (Lee et al. 2012). In addition, distinct sources of inhibition exert different impacts on pyramidal cell gain control under anesthesia (Atallah et al. 2012; Lee et al. 2012, Wilson et al. 2012). However, changes in behavioral state cause significant shifts in network activity and neuromodulatory environment, and the state-dependence of inhibitory recruitment and impact remains unknown. To explore the state-dependent impact of different types of inhibitory interneurons on neuronal gain, we recorded the responses of primary visual cortex neurons to drifting grating stimuli of varying contrast. Using optogenetic manipulation of genetically targeted cell classes, we enhanced or suppressed the activity of excitatory and specific inhibitory neurons (parvalbumin-, somatostatin- and vasoactive intestinal peptide-expressing cells) during extracellular multi-tetrode recordings in V1 of lightly anesthetized and awake (quiescent or running) animals. In addition to specifically manipulating these cells, we also identified their spontaneous and visually evoked activity by optogenetic 'tagging'. Interneuron contrast response curves were more linear in the awake state, mostly due to elevation of spontaneous firing rates.

Somatostatin-expressing cells showed the strongest increase in firing, significantly changing the predominant sources of inhibition across states. Changes in pyramidal cell contrast response curves induced by modulation of interneuron activity showed both multiplicative and additive effects in anesthetized animals, but became less multiplicative in awake animals. These findings reveal significant state-dependent changes in the interaction between excitation and inhibition, suggesting more complex network interactions in the awake cortex.

II-34. Receptive field formation by interacting excitatory and inhibitory plasticity

Claudia Clopath¹
Henning Sprekeler²
Tim P Vogels³

CC3450@COLUMBIA.EDU
H.SPREKELER@BIOLOGIE.HU-BERLIN.DE
TIMVOGELS@GMAIL.COM

¹Columbia University

²Humboldt University of Berlin

³Ecole Polytechnique Federale de Lausanne

Cortical neurons receive a balance of excitatory and inhibitory currents. This E/I balance is thought to be essential for the proper functioning of cortical networks, because it ensures their stability and provides an explanation for the irregular spiking activity observed in vivo. We recently suggested that activity-dependent Hebbian plasticity of inhibitory synapses could be a self-organization mechanism by which inhibitory currents can be adjusted to balance their excitatory counterpart (Vogels et al. Science 2011). The E/I balance not only generates irregular activity, it also changes neural response properties to sensory stimulation. Activity-dependent excitatory synaptic plasticity should therefore be sensitive to the E/I balance and should in turn be indirectly controlled by inhibitory plasticity. The question under which conditions excitatory Hebbian learning rules can establish receptive fields needs therefore to be reevaluated in the presence of inhibitory plasticity. In particular, it is of interest under which conditions neurons can simultaneously develop stimulus selectivity and a co-tuning of excitatory and inhibitory inputs. To address these questions, we analyze the dynamical interaction of excitatory and inhibitory Hebbian plasticity. We show analytically that the relative degree of plasticity of the excitatory and inhibitory synapses is an important factor for the learning dynamics, where faster inhibitory than excitatory learning is required for stable weights. We also find that the stimulus tuning of the inhibitory input neurons has a strong impact on receptive field formation. Our analysis suggests that the sliding threshold of BCM rules may not be implemented on a cellular level but rather by plastic inhibition arising from interneurons without stimulus tuning. If the stimulus tuning of the inhibitory input neurons is broader than that of the excitatory inputs, we observe a 'local' BCM behavior that leads to stimulus selectivity on the spatial scale of the inhibitory tuning width.

II-35. In vivo dissection of layer 1 inputs in the barrel cortex

Wanying Zhang
Randy Bruno
Columbia University

WANYING.ZHANG@GMAIL.COM
RANDYBRUNO@COLUMBIA.EDU

Layer 1 of the cerebral cortex is a largely acellular layer that consists mainly of long-range projection axons and apical dendrites of deeper pyramidal neurons. In the rodent barrel cortex (BC), layer 1 contains axons from both higher motor and sensory areas of the brain. Despite the abundance of synapses in L1 their actual contribution to sensory processing remains unknown. We investigated the impact of activating L1 long-range axons on BC layer 2/3 (L2/3) pyramidal neurons. We focused our study on three main sources of L1 synapses: the posterior medial nucleus of the thalamus (POm, the secondary somatosensory nucleus), the primary motor cortex (M1), and the secondary somatosensory cortex (S2). In each animal, we delivered the gene for channelrhodopsin (ChR2) to one of these three regions, and then photostimulated the ChR2-positive axons in BC L1 while recording whole-cell

recording from L2/3 cells in vivo. We found that while activation of POM axons elicits strong EPSPs in all recorded L2/3 cells, activation of M1 axons elicit long-lasting IPSPs in most of the cells. On the other hand, S2 axons evoke small or no detectable responses. These results indicate that even though the projection axons from different regions all project to BC L1, they form different sub-networks locally and can have dramatically different effects on signal processing in barrel cortex. We are also currently investigating how activations of these L1 axons interact with direct sensory inputs (whisker stimulation) in L2/3 pyramidal neurons.

II-36. Out of the zoo? Interneuron subtypes encode specific behavioral variables in the cingulate cortex

Sachin Ranade
Duda Kvitsiani
Balazs Hangya
Z. Josh Huang
Adam Kepecs

SACHIN.CSHL@GMAIL.COM
DKVITSIANI@GMAIL.COM
BALAZS.CSHL@GMAIL.COM
HUANGJ@CSHL.EDU
KEPECS@CSHL.EDU

Cold Spring Harbor Laboratory

Neurons in prefrontal cortex exhibit large response heterogeneity during behavior, encoding distinct combinations of sensory, motor and other features with diverse temporal dynamics. Many models (e.g. random-network models and liquid state-machines) assume that such diversity is intrinsic to cortical circuits and provide a computational role for heterogeneity. Moreover this perplexing 'cortical response zoo' has often been ascribed to cell-type diversity within cortical circuits. However, little is known about the behavioral correlates of identified cortical cell classes during behavior. To link identified neuron types with their network and behavioral functions, we recorded from the two largest genetically-defined inhibitory interneuron classes, the perisomatic-targeting parvalbumin (Pv) and the dendritic-targeting somatostatin (Som) neurons in anterior cingulate cortex using channelrhodopsin assisted activation as a physiological tag. We found that Pv and a subtype of Som neurons form functionally homogeneous populations showing a double dissociation between both their inhibitory impact and behavioral correlates. While a subtype of Som neurons selectively responded at reward approach, Pv neurons responded while leaving the reward site and encoded the preceding stay duration. These behavioral correlates of Pv and Som neurons defined a behavioral epoch and a decision variable important for foraging (whether to stay or to leave), a crucial function attributed to ACC. Furthermore, in contrast to Som neurons, Pv neurons fired in millisecond synchrony and exerted fast and powerful inhibition of principal cell firing, consistent with the idea that they respectively control the outputs of and inputs to principal neurons. These results suggest a connection between the circuit-level function of different interneuron-types in regulating the flow of information, and the behavioral functions served by the cortical circuits. Moreover these observations bolster the hope that functional response diversity during behavior can in part be explained by cell-type diversity and will further inform future theoretical studies.

II-37. A cellular mechanism for system memory consolidation

Urs Bergmann^{1,2}
Michiel Remme³
Susanne Schreiber³
Henning Sprekeler³
Richard Kempter⁴

URS.BERGMANN@BIOLOGIE.HU-BERLIN.DE
MICHIEL.REMME@HU-BERLIN.DE
S.SCHREIBER@HU-BERLIN.DE
H.SPREKELER@BIOLOGIE.HU-BERLIN.DE
R.KEMPTER@BIOLOGIE.HU-BERLIN.DE

¹Humboldt University Berlin

²BCCN Berlin

³Humboldt University of Berlin

⁴Humboldt-Universität zu Berlin

Declarative memories initially depend on the hippocampus. Over a period of weeks to years, however, these memories become hippocampus-independent through a process called system memory consolidation. The underlying cellular mechanisms are unclear. Here, we suggest a consolidation mechanism based on an anatomical network motif and spike timing-dependent plasticity. As a first step in the memory consolidation process, we focus on pyramidal neurons in the hippocampal CA1 area. Information from entorhinal cortex reaches these cells through two pathways: an indirect pathway via CA3 and the Schaffer collaterals (SC), and the direct perforant path (PP). Memory patterns are assumed to be initially stored in the recurrent CA3 network and SC synapses during the awake exploratory state. During a subsequent consolidation phase, CA3/SC-dependent memory patterns are then partly transferred to the PP synapses. Using numerical simulations and mathematical analysis, we show that this consolidation process occurs as a natural result from the combination of (1) spike timing-dependent plasticity at PP synapses and (2) the temporal correlations between SC and PP activity, since the SC input is delayed compared to the PP input (5-15 ms). Investigation of alternating wake and sleep phases reveals a decay of the memories in SC, but a slower decay in the PP connection to CA1. Less memory consolidation to the PP allows longer memory retention times, but comes at the price of worse initial memory storage. In a hierarchical network model that repeats the network motif across many levels, each direct connection at one level is part of the indirect pathway of the next level. Learning rates decrease with increasing level. Theoretical analysis and simulations of the hierarchical system show power-law forgetting, as seen with psychophysical forgetting functions (Wixted and Ebbesen 1991). Furthermore, consolidated memories yield faster responses because they are stored in increasingly shorter synaptic pathways.

II-38. Temporally evolving surround suppression helps decoding in a spiking model of motion processing

Philip Meier
Micah Richert
Jayram Moorkanikara Nageswaran
Eugene Izhikevich
Brain Corporation

MEIER@BRAINCORP.COM
RICHERT@BRAIN CORPORATION.COM
NAGESWARAN@BRAIN CORPORATION.COM
EUGENE.IZHKEVICH@BRAIN CORPORATION.COM

We find that a neural network with temporally evolving surround suppression improves the linear decodability of the network's population response. We present a novel model of motion processing that is fully implemented in a spiking neural network. We examine the role of lateral inhibition in V1 and MT. We model the response of the retina, V1 and MT, where each neuron is a single compartment conductance-based dynamical system. We apply a linear decoder to estimate the speed and direction of optic flow from a population response of a small spatial region in MT. Before training the decoder on population vector responses from MT with labeled speeds, we allow the spiking neural network to adapt the weights of the recurrent inhibitory neurons with spike-timing dependent plasticity (STDP). This allows the individual cells to adapt their dynamic range to the statistics reflected in the activity of the excitatory feed-forward network. Also, we impose a random onset latency of 1-10 ms for each feed-forward neuron. The combination of the onset latency and the inhibitory STDP results in a surround suppression with a magnitude that modulates throughout the course of the response, balancing the incoming excitatory drive. The temporally evolving surround suppression affects the activity of excitatory and inhibitory units in V1 and MT. The result is a population response of MT excitatory units that is more informative for decoding. The early response is less direction selective but drives the inhibition that sculpts the later responses. One source of improvement is that inhibition removes the non-selective response, but still preserves a robust selective response. Also the inhibition acts as gain control which limits how much saturation corrupts the linear code. We measure decoding performance by calculating the sum squared error of an estimate of the direction and speed of optic flow.

II-39. Functional bases for multidimensional neural computations

Joel Kaardal^{1,2}
 Jeff Fitzgerald¹
 Michael J. Berry II³
 Tatyana O. Sharpee^{1,2}

JKAARDAL@UCSD.EDU
 JEFFFITZGERALD@SALK.EDU
 BERRY@PRINCETON.EDU
 SHARPEE@SALK.EDU

¹Salk Institute for Biological Studies

²University of California, San Diego

³Princeton University

Despite the recent advances in dimensionality reduction methods for characterizing neural feature selectivity, it remains an open problem of how to relate the obtained relevant stimulus dimensions to the underlying neural circuitry. Part of the problem is that many of the dimensionality reduction methods specify features not uniquely, but up to a linear combination of them. Often, the results are presented in terms of orthogonal bases. Such orthogonal representations make it difficult to infer the corresponding neural computation that often involve a set of overlapping stimulus features, as in many types of motion computations (e.g. Gollisch & Meister, 2010). Here we introduce the idea of searching for a functional basis – a set of dimensions whose output can be combined according to a predefined function, such as the logical AND and OR operations. These bases may be non-orthogonal, over-complete, and, when applied to experimental data, often yield features with simple and easily interpretable form. The coefficients that describe the functional features within the relevant subspace are found by maximum likelihood optimization of a model with the chosen nonlinearity. We illustrate the approach on a variety of model neurons created to mimic properties of neurons of different stages of visual processing from the retina to extrastriate areas, as well as by applying to recordings of salamander retinal ganglion cell (RGC) responses. In the case of RGC responses, we find that the functional basis transformation yield spatially overlapping features that may be interpreted as arising either from individual bipolar cells or from irregularities in the dendritic computations of retinal ganglion cells. The functional basis representation yielded improved predicting power compared to models based on identical position-shifted features as well as models where no a priori assumptions about the form of the nonlinear function are made, as in the spike-triggered covariance (STC) dimensions).

II-40. Synaptic plasticity shapes the excitation/inhibition balance during on-going cortical activity

Aryeh Taub
 Yonatan Katz
 Ilan Lampl
 Weizmann Institute

ARYEH.TAUB@WEIZMANN.AC.IL
 YONATAN.KATZ@WEIZMANN.AC.IL
 ILAN.LAMPL@WEIZMANN.AC.IL

The prominent feedback connections between excitatory and inhibitory neurons in the cortex suggest that the cortex operates at a balanced state where inhibition modulates excitation. To examine this, we recorded the spontaneous excitatory and inhibitory inputs onto cortical neurons, while inducing shifts in brain-state by altering the depth of anesthesia. Although the rate of both excitatory and inhibitory events decreased under deeper anesthesia, the magnitude of inhibition increased, while excitation was unaffected. Importantly, that excitation was indifferent to the change in inhibition implies that spontaneous cortical activity is not at a balanced-state. To examine the relationship between the magnitude of inhibition and cortical-states, we replayed the temporal patterns of spontaneous inhibitory activity using cortical electrical-stimulation while blocking local excitation. The magnitude of inhibition increased as the rate of stimulation decreased, similar to the observation under deep anesthesia. Surprisingly, this occurred irrespectively of the depth of anesthesia, suggesting that the excitation-inhibition balance during spontaneous cortical activity is determined mainly by the short-term synaptic properties of feedforward inhibitory inputs.

II-41. Attention improves information processing by tuning cortical networks towards critical states

Nergis Toemen
Udo Ernst
University of Bremen

NERGIS@NEURO.UNI-BREMEN.DE
UDO@NEURO.UNI-BREMEN.DE

Cortical information processing is highly flexible and adapts rapidly to the current behavioral task. For example, attention has been shown to boost performance of subjects in psychophysical tasks, and to improve cortical representations of visual stimuli in electrophysiological recordings. There is evidence that these phenomena rely on collective dynamics in visual cortical networks which is more complex than simple firing rate modulations or improved signal-to-noise ratios. However, it is currently unclear which particular neural mechanisms underlie improved stimulus processing under attention. The dynamics of cortical networks exhibit certain properties that have been suggested to be linked to optimization of information processing. Experimental evidence suggests that cortical networks operate near a 'critical' state in which scale-free avalanches of spike events occur, generating neural patterns which are 'rich' in structure. In addition, excitatory and inhibitory synaptic currents to cortical neurons can be balanced, which explains the high variability of neural activity and would allow cells to rapidly react to changes in the stimulus. Here we explore the relationship between criticality, balance, and enhancing stimulus representations in a simple network of integrate-and-fire neurons driven by an external stimulus. By increasing the efficacy of recurrent couplings, attention enhances spontaneous synchronization and renders the activation patterns for different stimuli more distinct. Their discriminability is maximized near the critical state of the network. Moreover, we link these results to experiments in awake behaving monkeys, by reproducing the power spectra of local field potentials and the observed increase in their difference under attention. Taken together, our study implies that attention drives the cortex towards a critical state, hereby maximizing the discriminability of different stimuli represented in the corresponding network activity. Furthermore, our framework suggests a novel role for synchronization in cortical information processing.

II-42. Beyond acute experiments: automated long-term tracking of socially-housed mice

Shay Ohayon¹
Ofar Avni²
Adam Taylor²
Roian Egnor²
Pietro Perona¹

SHAY.OHAYON@GMAIL.COM
AVNIO@GMAIL.COM
ADAMLYLETAYLOR@GMAIL.COM
EGNORR@JANELIA.HHMI.ORG
PERONA@CALTECH.EDU

¹California Institute of Technology

²Janelia Farm Research Campus

Social relationships in mice develop and evolve over the course of many days. The ability to carry out thorough, quantitative, long-term observations of mice in a natural environment would likely have transformative effects in understanding how genetic, pharmacological and environmental manipulations affect long term animal social behavior. However, current state of the art systems are limited to the observation of two mice sharing an unfamiliar enclosure for a period of 10-20 minutes, often in partition cages, which limit social interaction, and requires human verification and correction of tracking results due to identity swaps. We present a robust method for tracking multiple mice in a large enclosure, as they interact over multiple days through dark and light cycles. By integrating automatically trained computer-vision classifiers with a Hidden Markov Model identities are can be preserved across long periods even when mice are occluded or burrow, disappearing from view. From single-camera overhead video of the mouse enclosure our system computes the trajectory of each individual, identified by uniquely discriminable fur bleach marks. Trajectories can be analyzed to measure individual and social behavior, such as courtship, dominance and aggression as they develop over the course of days, beyond the range of acute experiments. We show the applicability of our system by tracking groups of four mice (two males and two females)

continuously for five days and analyzing their behavior. We find that place preference, social association and male dominance relationships gradually evolve over days. The ability to track individual mice during social interactions unfolding over many days without the need for tedious human verification makes our system an unprecedented tool for observing the effects of genetic, pharmacological and environmental manipulations on long-term complex social behavior.

II-43. Modeling adaptive changes in the motor program underlying birdsong

Baktash Babadi
Bence Olveczky
Harvard University

BBABADI@FAS.HARVARD.EDU
OLVECZKY@FAS.HARVARD.EDU

The zebra finch, a songbird, learns its courtship vocalization in much the same way that we learn many of our motor skills, thus offering a tractable experimental system for studying the neural mechanisms underlying complex motor learning. The circuits controlling vocal musculature in songbirds comprise premotor nucleus HVC and motor cortex analogue brain area RA. HVC neurons are thought to encode the temporal structure of song, whereas their connections to RA neurons translate this timing code into muscle activity patterns driving learned vocalizations. Hence, the HVC-RA connections are assumed to be an important locus for learning. Recent voltage-clamp recordings from RA projection neurons have revealed a developmental strengthening of HVC-RA synaptic inputs that are accompanied by a decrease in the total number of inputs. Additionally, significant paired-pulse depression (PPD) was observed at the HVC-RA synapse in juvenile, but not adult, birds, suggesting a possible role for PPD in shaping the motor circuits during learning. We implemented a biologically plausible computational model of the song circuit, in which RA projection neurons were modeled as integrate-and-fire neurons receiving precisely patterned temporal input from HVC and Poisson spike trains from LMAN, the outflow nucleus of a basal ganglia circuit known to be necessary for song learning. Our model incorporated the observed age-related trends in RA synaptic connectivity, and show that they account for the developmental maturation of song-related RA firing patterns. We further show that implementing Spike Timing-Dependent Plasticity (STDP) is an effective and plausible driver of the synaptic changes observed during development. Moreover, combining STDP and PPD, allowed RA neurons to incorporate time-varying 'instructive' inputs from LMAN. Our results suggest that STDP coupled with a depressing HVC-RA synapse can account for the adaptive strengthening and pruning of these synapses during development, and thus underlie song learning.

II-44. Do orientation preference maps arise from hexagonal retinal ganglion cell mosaics?

Manuel Schottdorf¹
Wolfgang Keil^{2,3}
Michael Schnabel⁴
David M. Coppola⁵
Siegfried Löwel⁶
Leonard E. White⁷
Matthias Kaschube⁸
Fred Wolf⁹

MANUEL@NLD.DS.MPG.DE
WOLFGANG@NLD.DS.MPG.DE
M-SCHNABEL@NORTHWESTERN.EDU
DCOPPOLA10@VERIZON.NET
SIEGRID.LOEWEL@BIOLOGIE.UNI-GOETTINGEN.DE
LEN.WHITE@DUKE.EDU
KASCHUBE@FIAS.UNI-FRANKFURT.DE
FRED@NLD.DS.MPG.DE

¹Max Planck Inst. Dynamics & Self-organization

²MPI for Dynamics & Self-organization

³BCCN Göttingen

⁴Northwestern University

⁵Randolph-Macon College

⁶Northwestern

⁷Duke University

⁸Frankfurt Institute for Advanced Studies

⁹MPI for Dynamics and Self-Organisation

We quantitatively evaluate a recently proposed model for Orientation Preference Maps (OPMs), in which the spatial distribution of ON- and OFF-center receptive fields of retinal ganglion cells (RGCs) seeds the structure of receptive fields and OPMs in the visual cortex and activity-dependent processes act as a mere refinement during postnatal development. In this model, OPMs with a quasi-periodic roughly hexagonal structure arise as the result of Moiré-Interference of feedforward inputs from hexagonal RGC mosaics. While the model has been shown to qualitatively account for several statistical properties of OPMs, both, a precise mathematical assessment of its predictions and a quantitative comparison to a robust experimental data set of OPMs is still lacking. Such quantitative investigation appears critical to elucidate the relative contributions of feedforward seeding and activity-dependent refinement in shaping the spatial layout of OPMs. First, we analytically determine visual cortical receptive fields and OPMs as predicted by the Moiré-Interference model and derive expressions for several spatial statistics, including the density of topological defects, called pinwheels. We then perform a statistical analysis of local hexagonal order in a large set of experimentally obtained OPMs from tree shrew, galago and ferret (>90 maps). We compare measured maps to model maps and the most general control ensemble consistent with the null-hypothesis that OPMs are statistically isotropic and lack excess hexagonal order. We find that different indicators of hexagonal order in experimentally obtained OPMs are statistically indistinguishable from these isotropic control families. Moreover, we show that several spatial statistics of OPMs predicted by the model, are qualitatively different from the ones observed experimentally. Our results challenge the Moiré-Interference model for OPMs and argue for a small contribution of subcortical constraints to the spatial layout of OPMs compared to activity-dependent processes during postnatal development.

II-45. Loss of theta modulation of the hippocampal cell firing is accompanied by deterioration of episode f

Yingxue Wang^{1,2}

Eva Pastalkova^{1,2}

WANGY12@JANELIA.HHMI.ORG

PASTAK@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus

²HHMI

The firing of cells in the hippocampus is modulated by theta oscillations (5-11 Hz), which appear to enable temporally precise cell interactions within the network. It is unclear, however, whether the theta modulation is necessary for the generation of neuronal activity patterns related to memory formation. To investigate this question, we eliminated theta oscillations in hippocampus by medial septum inactivation using GABA-agonist muscimol in rats during a delayed memory task: a two-arm-alternation task with 10 seconds of wheel running during the delay period. The task was chosen because CA1 neurons with place-cell-like activity appeared not only in the maze (place cells), but also during the wheel run (episode cells) when the perceived sensory cues were stationary and the firing was presumably generated within the network depending on the recent experience [1]. Following the loss of theta modulation of CA1 cell firing, the animal's performance in the memory task was impaired. At the same time, the transient firing pattern of episode cells changed completely: cells were either active throughout the entire wheel run or were mostly silent. In contrast, the external cue dependent place fields were retained to a large extent. This finding suggests that theta modulation of single cell firing and network interactions are necessary for the formation of the memory task dependent transient firing of episode cells and for the mnemonic function of the hippocampus. Reference: [1] E Pastalkova, V Itskov, A Amarasingham, G Buzsáki, 'Internally generated cell assembly sequences in the rat hippocampus', Science, Vol. 321, pp. 1322-1327, 2008

II-46. The sense of place: grid cells in the brain and the transcendental number e

Vijay Balasubramanian¹
 Jason Prentice²
 Xuexin Wei¹

VIJAY@PHYSICS.UPENN.EDU
 JASONSP@PRINCETON.EDU
 WEIXPKU@GMAIL.COM

¹University of Pennsylvania
²Princeton University

Grid cells in the brain represent place by responding when the animal occupies one of a periodic lattice of 'grid fields' during spatial navigation. The grid scale varies systematically along the dorso-ventral axis of the entorhinal cortex. Here we propose that the grid system minimizes the number of neurons required to encode location with a given spatial resolution. From this hypothesis, we derive a number of predictions about grid coding in two dimensions: (i) grid scales should follow a geometric progression, implementing a two-dimensional analog of a base- b number system, (ii) the mean ratio between adjacent grid scales should be $e^{1/2} \sim 1.6$ for idealized, noiseless neurons, and should robustly lie in the range ~ 1.4 - 1.7 for realistic neurons, independently of the decoding scheme used by the brain, (iii) the scale ratio should vary modestly within and between animals, (iv) the ratio between grid scale and individual grid field widths at that scale should also lie in this range, (v) grid fields should lie on a triangular lattice. All five predictions are supported quantitatively by recent experiments. Specifically, Barry et al. (2007) reported a mean ratio of adjacent grid scales of ~ 1.7 . Recent results based on larger data set reveal a geometric progression of grid cells with a geometric factor ~ 1.42 , accompanied with modest variability in the scaling of the grids (Stensola, et al., 2012). Meanwhile, Giacomo, et al. (2011) reported that the mean ratio between grid scale and individual grid field width (diameter) is ~ 1.65 . The theory makes new testable predictions for optimal grids supporting navigation in one and three dimensions. Our results suggest that a principle of economy organizes key characteristics of neural circuits supporting higher cognitive functions.

II-47. Characteristics prediction of STDP in hippocampal CA1 network by mutual information maximization

Ryota Miyata¹
 Keisuke Ota²
 Toru Aonishi¹

MIYATA@ACS.DIS.TITECH.AC.JP
 K-OTA@BRAIN.RIKEN.JP
 AONISHI@DIS.TITECH.AC.JP

¹Tokyo Institute of Technology
²RIKEN BSI

Spike-timing-dependent plasticity (STDP) is a biological process that adjusts synaptic efficacy depending on the relative timing of pre- and postsynaptic spikes. STDP has been found in a wide variety of nervous systems, and a rich diversity of temporal windows for STDP induction has been reported. There have been several approaches to interpreting the computational roles of STDPs in neural circuits. Lengyel et al. (2005) developed a normative theory of auto-associative memory on the basis of the Bayes' theorem, and derived pairs of STDPs and phase response curves (PRCs) optimally functioning as auto-associative memory, i.e., retrieval of temporal memory spike patterns from the noisy ones. In the hippocampal CA1 region, it has been reported that reverse replay and temporally structured replay of memory spike patterns occur during rest and REM sleep, respectively. Here, we explore optimal neural implementations recalling such reverse-ordered and spread-out spike patterns. First, we formulate a hetero-associative memory network recalling not only the normal spike patterns, but also the reverse-ordered and twice spread-out patterns as a phase oscillator model consisting of an STDP and a PRC. Next, we analytically derive the mutual information between a stored phase pattern and a network output for evaluating memory retrieval performance. By maximizing the mutual information, we search a set of optimal STDPs under the constraint of PRCs of hippocampal CA1 pyramidal neurons recorded in vitro. The typical STDPs observed in CA1 region are classified into two types: symmetric and asymmetric plasticity rules. We show both of these rules are included in the theoretically derived set of optimal STDPs. The theoretically derived STDPs qualitatively

coincide with the first two Fourier series approximations for those reported in CA1 neurons.

II-48. Time-scales of neural integration constrain a songbird operant behavior

Yoonseob Lim
Barbara Shinn-Cunningham
Timothy Gardner
Boston University

YSLIM@BU.EDU
SHINN@CNS.BU.EDU
TIMOTHYG@BU.EDU

In spite of long standing theoretical interest, the cortical machinery underlying temporal pattern recognition remains largely unknown, though the circuit principles, once known, could suggest new technologies for hard problems such as speech perception in noise. On a theoretical level, synfire chains [1], avalanches [2] or transient dynamics in recurrent networks [3] are proposed to underlie temporal processing. In other models, persistent currents in single cells bridge intervals of time [4]. Here, we examine how cortical auditory neurons acquire sensitivity to temporal patterns by training songbirds to recognize and respond to a sequence of clicks. The task is reminiscent of Morse code pattern detection. In the songbird task, the only difference between target and non-target stimuli is the sequential ordering of a fixed set of click intervals (click intervals range from 11ms to 40ms. Patterns are 250ms to 3sec long). By construction, no spectral cues exist, and no single interval of time between clicks provides a basis for discrimination. To solve the task, songbirds must form memories for ordered sets of multiple time-intervals. In awake birds, neurons in the first stage of auditory cortex respond synchronously with low latency and high temporal precision to each click. One synapse further along the auditory stream, neurons are selective for specific combinations of intervals, while still maintaining precise timing. When the duration of click patterns is stretched by a factor of two, songbird behavioral performance and neural sequence selectivity break down, revealing a fundamental time-scale of sensory integration in the songbird auditory cortex.

II-49. Unsupervised learning of binaural features from naturalistic stereo sounds

Wiktor Młynarski
Max Planck Institute for Math in the Sciences

MLYNAR@MIS.MPG.DE

Binaural hearing mechanisms utilize differences in the sound arriving at the left and right ear to extract information about spatial configuration of sound sources. According to the widely acknowledged Duplex Theory, sounds of low frequency are localized based on Interaural Time Differences (ITDs). Position of high frequency sounds is identified based on Interaural Level Differences (ILDs), since phase difference corresponding to ITDs becomes ambiguous with growing frequency. Natural sounds, however, possess a rich structure and contain multiple frequency components. This leads to the question - what are the contributions of different cues to sound position identification under natural stimulation? In this work I exploit a sparse coding model of stereo sounds. Such approach allows to find out, what kind of high-order structure is present in the binaural signal, and how informative it is about the position of the sound source. I simulate naturalistic stereo signal, by convolving speech sounds with experimentally measured, human Head Related Impulse Response Filters (HRIRs). In the next step, I learn sparse coding dictionaries of stereo sounds. This is done in different signal domains (raw waveform, log-spectrogram). Finally, I analyze properties of the basis functions and dependence of the encoding on the sound source position. Learned basis functions capture regularities present in the sound waveform as well as those resulting from differences between the ears. Features corresponding to ITDs/ ILDs can be identified in the sparse basis, which allows to study their informativity about sound position in a naturalistic setting. In addition, a subpopulation of spectro temporal basis functions displays complex interaural differences. This can be related to experimental findings showing, that binaural hearing relies not only on phase or level disparities, but also on comparison of sound spectra at both sides of the head.

II-50. Closing the loop; Inverse-model learning with a nonlinear avian syrinx and sparse auditory coding

Alexander Hanuschkin^{1,2}
Surya Ganguli³
Richard Hahnloser^{4,5}

ALEXANDER.HANUSCHKIN@INI.PHYS.ETHZ.CH
SGANGULI@STANFORD.EDU
RICH@INI.PHYS.ETHZ.CH

¹INI, University of Zurich and ETH Zurich

²Neuroscience Center Zurich

³Stanford University

⁴Institute of Neuroinformatics UZH / ETHZ

⁵Inst. Für Neuroinformatik UZH

Control-theoretic inverse models are very useful for learning and generating flexible sensory-goal directed motor behaviors. We have recently proposed a simple eligibility-weighted Hebbian learning rule capable of provably forming inverse models in high dimensional linear networks by associating random motor explorations with their future sensory consequences. In this theory the inverse model forms in the synaptic connections from sensory to motor neurons, allowing the conversion of a sensory memory (for example a tutor song template) into the necessary motor pattern required to reproduce the sensory memory. Here we study both a nonlinear extension of this model and analytically demonstrate the relationship between inverse models and mirror neurons. We test inverse model learning in a nonlinear mass-spring model of the avian syrinx and an efficient sparse sensory representation of sound. Our learning rule learns appropriate inverse models. The inverses we find are causal (map sensation to the same action) or predictive (map sensation to future action) depending on the stereotypy of the neural code for motor explorations. In a random code, the formed inverse is causal and maximally useful for feedforward motor control because it allows imitation of arbitrary sensory target sequences. We also show mirror neurons naturally arise during inverse model learning. Mirroring of motor and sensory evoked activity is either in precise temporal register, reflecting predictive inverses associated with stereotyped motor codes, or temporary delayed, reflecting causal inverses associated with variable motor codes. Overall, this work demonstrates that bird song can be learned in realistic models of sound production, sound perception, and synaptic learning rules, and creates new conceptual connections (consistent with differences between HVC/LMAN in birds) between the stereotypy of the motor code, the causal nature of a learned inverse model, and the temporal lag between sensory and motor responses of the mirror neurons it contains.

II-51. Timing of invariant object recognition in humans

Leyla Isik
Ethan Meyers
Joel Z Leibo
Tomaso Poggio

LISIK@MIT.EDU
EMEYERS@MIT.EDU
JZLEIBO@MIT.EDU
TP@AI.MIT.EDU

Massachusetts Institute of Technology

The human visual system can rapidly recognize objects despite transformations to their visual appearance, such as position in the visual field, size, and viewpoint. The precise timing and stages involved in invariant object recognition, however, are still poorly understood. Here we apply new methods in magnetoencephalography (MEG) decoding to measure the latencies of position- and size-invariant visual information in the ventral stream. With these methods we can read out the identity of objects based on subjects' MEG data as early as 60 ms, and do so invariantly to size and position at 125 ms and 150 ms, respectively. The temporal accuracy of MEG decoding reveals several interesting properties of invariant object recognition. First, we observe a clear delay between the initial identity decoding and invariant decoding that is consistent across subjects. Second, these visual signals have very fast dynamics, and change and move to a different sensor configuration within a 20-50 ms time window. Third, invariance develops in a sequential order, meaning that smaller transformations were decoded before larger transformations. This sequential development is consistent with a hierarchical, feed-forward visual model where

receptive fields pool at each successive visual layer to first create local invariance and then build invariance over a larger area. In conjunction with this timing data, preliminary source localization results suggest that invariant visual representations are being developed as neural signals move down ventral stream. This study provides novel MEG decoding methods as well as results directly comparing the dynamics of size- and position-invariance in the human visual system. Together, these new methods and their applications bring us closer to a computational understanding of invariant object recognition.

II-52. Object-vision models that better explain IT also categorize better, but all models fail at both

Seyed-Mahdi Khaligh-Razavi
Nikolaus Kriegeskorte
University of Cambridge

SEYED.KALIGHRAZAVI@MRC-CBU.CAM.AC.UK
NIKOLAUS.KRIEGESKORTE@MRC-CBU.CAM.AC.UK

The inferior temporal (IT) representation is thought to serve object recognition and has been shown to explain human patterns of categorization performance across different tasks. However, current computational models do not reach human performance levels and it is unclear to what extent their internal representational geometries match the IT representation. Here we investigate a wide range of computational models and test their categorization performance and their ability to account for the IT representational geometry. The models included some of the well-known neuroscientific object-recognition models (including HMAX and VisNet) along with several models from computer vision (including SIFT, GIST, and Self-similarity features). Some of the models were trained (either supervised or unsupervised), others did not require training. We compared the representational dissimilarity matrices (RDMs) of the model representations with the RDMs obtained from human IT (measured with fMRI) and monkey IT (measured with cell recording) for the same set of stimuli (not used in training the models). We found that the more similar a model's representational geometry was to IT, the better the model performed at categorization (e.g. for animate vs. inanimate). This is unsurprising because the IT RDM exhibits strong category clustering (within-category dissimilarities < between-category dissimilarities). However, even when only considering the within-category representational geometry, the models that better explained IT also tended to better categorize. This lends a new kind of support to the idea that understanding the IT representation is a good way to improve computer-vision performance. While many models explained significant variance within the IT representational dissimilarities, none of them explained more than about 20% of the non-noise variance. The best models were HMAX and a combination of features across all models. Overall, these results suggest that the crucial type of feature for explaining IT and performing object recognition has yet to be discovered.

II-53. Size-invariant shape coding in visual area V4

Yasmine El-Shamayleh^{1,2}
Anitha Pasupathy¹

YASMINE1@UW.EDU
PASUPAT@U.WASHINGTON.EDU

¹University of Washington

²Washington National Primate Research Center

How do we recognize objects across changes in retinal size? This fundamental capacity of biological visual systems is computationally challenging. Size-invariant object recognition is supported by neurons in IT cortex, which maintain their preferences for objects across changes in scale. However, the detailed mechanisms that establish invariance remain unclear. To investigate this at the single neuron level, we targeted cortical area V4, a critical stage of object processing and the foundation of IT responses. Importantly, we leverage our understanding of object representation in V4 and a candidate model of shape encoding that makes direct predictions for neuronal responses at different sizes. Many V4 neurons encode objects in terms of their component contour features; their selectivity can be modeled as preferences for contour curvature (convex/concave) at specific locations relative

to object center (Pasupathy and Connor, 2001). This model presumes that neurons encode absolute curvature, a variable that is inversely related to object size; e.g., the curvature of a circle halves as its radius doubles. Thus, a curvature-tuned neuron cannot be size-invariant. This is because a particular contour feature will have different curvatures at different scales. We exploit this key idea here to ask whether neurons are curvature-tuned or size-invariant. We characterized neurons in two awake-fixating primates using parametric shapes that sampled a range of curvature values. Stimuli were presented at 4–5 scales and at 8 rotations, all inside the neuron's receptive field. We found that most V4 neurons were size-invariant, maintaining their shape preferences across the range of scales sampled (~2 octaves). Only a few neurons were curvature-tuned, shifting their shape preferences systematically, as predicted by the model. Our results motivate a key refinement of the curvature model; V4 neurons encode contour characteristics relative to object size — a shape code that can support size-invariant object recognition.

II-54. A general theory of learning and memory with complex synapses

Subhaneil Lahiri
Surya Ganguli
Stanford University

SULAHIRI@STANFORD.EDU
SGANGULI@STANFORD.EDU

An incredible gulf separates theoretical models of synapses, often described solely by a single scalar value denoting the size of a postsynaptic potential, from the immense complexity of molecular signaling pathways underlying real synapses. To elucidate the functional contribution of such molecular complexity to learning and memory, it is essential to expand our theoretical conception of a synapse from a single scalar to an entire dynamical system with many internal molecular functional states. Moreover, theoretical considerations alone demand such an expansion; network models with scalar synapses assuming finite numbers of distinguishable synaptic strengths have strikingly limited memory capacity [Amit/Fusi92]. This raises the fundamental question, how does synaptic complexity give rise to memory? To address this, we develop new mathematical theorems elucidating the relationship between the structural organization and memory properties of complex synapses that are themselves molecular networks. We consider an extremely general class of models where memories are stored in a population of N synapses each with M internal molecular functional states, where potentiation and depression each induce an arbitrary network transition between states, parameterized by a pair of $M \times M$ stochastic transition matrices. The cascade model of [Fusi 05] for example is one member of this model class. We find that no molecular network can achieve a memory capacity that exceeds $N^{1/2} M$, or have a memory curve that exceeds a power law envelope with exponent -1 . Molecular networks achieving optimal capacity at any given time correspond to a simple linear chain of states, but have highly suboptimal memory at other times, and model independent trade-offs between storing proximal and distal memories necessitate synaptic complexity. Overall, we uncover general design principles governing the functional organization of complex molecular networks, and suggest new experimental observables in synaptic physiology, based on first passage time theory, that connect molecular complexity to memory.

II-55. Multi-step decision tasks for dissociating model-based and model-free learning in rodents.

Thomas Akam¹
Peter Dayan²
Rui Costa¹

THOMAS.AKAM@NEURO.FCHAMPALIMAUD.ORG
DAYAN@GATSBY.UCL.AC.UK
RUI.COSTA@NEURO.FCHAMPALIMAUD.ORG

¹Champalimaud Neuroscience Program

²University College, London

We have developed two multi-step decision tasks involving chains of actions, intended for use in rodent elec-

trophysiology. In both tasks an initial decision between a pair of actions (lever presses) makes available one of two further 'second link' actions (retractable levers), which in turn lead to the trial outcome (reward or timeout). The first task aims to dissociate model-based and model-free learning by revaluing the states reached as a consequence of the initial decision through experience with the second link actions under changed reward contingencies. Revaluation occurs using trials in which the animal chooses directly between the second link actions (Fig. 1a). We have shown that mice are able to use experience on these revaluation trials to guide subsequent choices on decisions between the first links (Fig. 1b). The second task, adapted from a recent design by Daw et al. (2011), uses probabilistic transitions between the initial decision and the second link states, such that each of the initially available actions has a normal transition which makes one of the second link actions available, and a rare transition which makes the other available (Fig. 2a). We have shown in mice that the effect of trial outcome (reward or timeout) on choice probability in the subsequent trial depends on whether the outcome followed a normal or rare transition (Fig. 2b), consistent with the use of a model-based strategy. One concern with this type of task is that sophisticated model free strategies may exist which can produce behaviour closely resembling model based control. Specifically, in our tasks, outcomes from second link actions on a given trial could be used as discriminative stimuli to guide subsequent choices between first links. We are working to identify how well we can distinguish between these possibilities by comparing different reinforcement learning models fitted to the behaviour.

II-56. Ghrelin modulates phasic dopamine evoked by food-reward via action in the lateral hypothalamus

Jackson Cone
Mitchell Roitman

University of Illinois at Chicago

JCONE2@UIC.EDU
MROITMAN@UIC.EDU

Brief (phasic) increases in dopamine (DA) in the nucleus accumbens (NAc), signal unexpected food reward and participate in learning, reinforcement and goal-directed behavior. However, the rewarding value of food depends on many factors, including physiological state. Ghrelin, a peptide secreted by the stomach, signals energy need and promotes food intake via actions in multiple brain regions. Indeed, ghrelin may promote food intake, at least in part, by influencing midbrain DA neurons. While microdialysis studies indicate that ghrelin increases DA concentration in the NAc, whether or not ghrelin influences phasic DA evoked by rewarding stimuli remains unknown. Here, we investigated whether ghrelin modulates phasic DA evoked by food reward using fast-scan cyclic voltammetry (FSCV) in awake-behaving rats. Rats were trained to retrieve sugar pellets that were unpredictably delivered. Following training, rats were surgically prepared for FSCV and a cannula was aimed at the lateral ventricle (LV). We recorded NAc DA before and after infusion of either ghrelin or vehicle into the LV of ad-libitum fed rats. We repeated this experiment in a second group of food-restricted rats but infused either the ghrelin receptor antagonist D-[Lys]-GHRP or vehicle into the LV. LV ghrelin significantly increased, while LV D-[Lys]-GHRP decreased, the magnitude of phasic DA evoked by food reward compared to animals that received LV vehicle. To begin determining a locus for ghrelin action, rats were prepared as above but cannulae were aimed at the lateral hypothalamus (LH). In ad-libitum fed rats, LH ghrelin, but not vehicle, significantly enhanced phasic DA signaling evoked by food reward. Our data demonstrate that ghrelin signaling in the LH is an important locus for physiological state to affect the neural circuitry underlying motivated behavior.

II-57. Dopamine modulates functional communication of the human striatum subdivision

Payam Piray
 Marieke van der Schaaf
 Ivan Toni
 Roshan Cools
 Radboud University Nijmegen

P.PIRAY@DONDEERS.RU.NL
 MARIEKE.VANDERSCHAAF@FCDONDEERS.RU.NL
 IVAN.TONI@FCDONDEERS.RU.NL
 ROSHAN.COOLS@GMAIL.COM

Many functions of the mammalian brain depend on information processing within the striatum and on interactions between distinct striatal subregions. Neuroanatomical data from non-human primates have suggested that interactions between striatal subregions might depend on dopaminergic afferents from the midbrain to the striatum, but the presence and the direction of those modulatory effects remain controversial (Haber et al., J Neurosci 2000; Bergman et al., Trends Neurosci 1998). Here, we address this issue by quantifying the effects of modulating dopaminergic receptors on resting-state interactions between observer-independent striatal subdivisions. We scanned 25 participants using fMRI after intake of a dopamine receptor agonist, a dopamine receptor antagonist and placebo. To identify observer-independent functional subdivisions within the striatum, we used clustering, a data-driven and multidimensional method that takes into account time-courses of all striatal voxels to decompose the striatum into clusters. First, we used clustering to identify functional subdivisions of the striatum. Second, to quantify functional segregation within the striatum, we calculated average between-cluster dissimilarity (BCD), the dissimilarity between voxels that grouped into different clusters. Thus, greater BCD reflects less communication between clusters and a higher degree of segregation within the striatum. We found that there is a stable effect of dopaminergic manipulation on BCD. In particular, we found that blockade of dopamine receptors with sulpiride enhanced functional segregation, while stimulation of dopamine receptors with bromocriptine reduced functional segregation. These findings highlight the effects of dopamine on communication between striatal subdivisions and might have implications for pathologies hypothesized to be related to striatal subregions' communication through the dopaminergic network, such as drug addiction (Belin and Everitt, Neuron 2008).

II-58. Volitional control by a person with tetraplegia of high gamma local field potentials (LFPs) recorded

Tomislav Milekovic¹
 Daniel Bacher¹
 Blaise Yvert¹
 Leigh Hochberg¹
 Emad Eskandar²
 Sydney Cash²
 John Donoghue¹

TOMISLAV_MILEKOVIC@BROWN.EDU
 DANIEL_BACHER@BROWN.EDU
 BLAISE_YVERT@BROWN.EDU
 LEIGH_HOCHBERG@BROWN.EDU
 EESKANDAR@PARTNERS.ORG
 SCASH@PARTNERS.ORG
 JOHN_DONOGHUE@BROWN.EDU

¹Brown University

²Harvard Medical School; MGH

Brain-computer interfaces (BCIs) hold great promise for restoring movement and communication for people with longstanding tetraplegia by enabling control of external devices, such as a computer cursor or a robotic arm, exclusively using neuronal activity. In most intracortical BCIs, control of the devices is achieved by directly decoding the user's intentions from spiking activity of neurons in the motor cortex. Local field potentials (LFPs) of people with tetraplegia have also been shown to contain signals that can be used to decode intentions. However, even with the optimal decoding algorithms, the accuracy and the complexity of the BCI control will depend on the degree to which the neuronal activity can be volitionally modulated. Here, we show that a person with tetraplegia can volitionally control the amplitudes in the gamma frequency band (55-192Hz) of the LFPs recorded by a single microelectrode, thereby demonstrating a human BCI based on high gamma LFPs. In our study (BrainGate2, IDE), a 66 year-old man with tetraplegia was implanted with a 96-channel microelectrode array (Blackrock Microsystems)

in the hand/arm area of his primary motor cortex. During the recording sessions, the participant interacted with a simple game shown on a computer monitor. His task was to position the cursor over a target which changed position every 10 seconds. LFPs, recorded by a single electrode of the array, were converted into a vertical cursor position by extracting the logarithm of the amplitude in the 55-192Hz band and then normalizing the values by a linear transform to fit them in the monitor workspace. We measured the participant's ability to control the cursor by calculating mean cursor-to-target distance. The participant's control of the cursor was up to 22 μ m1% better than random cursor control (session 1: 22 μ m1%; session 2: 5.9 μ m0.1%; $P < 10^{-6}$), as measured by randomly shuffling target positions.

II-59. Value-updating interaction among contexts in choice behaviors of rats

Akihiro Funamizu¹

Makoto Ito²

Kenji Doya²

Ryohei Kanzaki³

Hirokazu Takahashi³

FUNAMIZU@OIST.JP

ITO@OIST.JP

DOYA@OIST.JP

KANZAKI@RCAST.U-TOKYO.AC.JP

TAKAHASHI@I.U-TOKYO.AC.JP

¹JSPS research fellow

²OIST Graduate University

³University of Tokyo

To investigate how a decision making in one context affects a decision in another context, we analyzed rats' performance in a choice task consisting of a random sequence of fixed-reward and variable-reward trials. Only in fixed-reward trials, a light stimulus was presented. In the choice task, rats were asked to nose-poke to a left or right hole, and received a reward stochastically. While the reward probability was fixed in fixed-reward trials, it was varied among 4 settings after the choice frequency of more rewarding (i.e., optimal) hole reached 80% in variable-reward trials. In variable-reward trials, rats found an optimal choice earlier when the optimal choices of variable- and fixed-reward condition were same hole, compared when the optimal choices of two conditions were different. In addition, the choices in variable-reward trials were affected by the action and reward presentation of previous fixed-reward trial, indicating that the experiences in fixed-reward trials affect the behaviors in variable-reward trials. The choice behaviors were then analyzed with reinforcement learning models. A forgetting Q-learning model, where the action values in chosen and un-chosen options were updated and forgotten, respectively, fit to the behaviors in variable-reward trials, while a fixed-choice-probability model, where the choice probability was constant in all trials, fit to the behaviors in fixed-reward trials. This result suggests that the learning strategy depended on the reward context. An interactive value-updating model, in which the action values in variable-reward condition were updated with the experiences not only in variable- but also in fixed-reward trials, better fit to the whole choice sequences (i.e., random sequences of variable- and fixed-reward trials), compared to the models with independently updating action values in each condition. Thus, our results suggest that the experiences in one condition affect the value-updating in another.

II-60. Solving the secretary problem

Vincent Costa¹

Bruno Averbeck²

¹NIMH/NIH

²NIH

COSTAVD@MAIL.NIH.GOV

BRUNO.AVERBECK@NIH.GOV

The secretary problem is an optimal stopping problem that involves the sequential presentation of a list of items of a predetermined length. The goal of the task is to find the 'best' item in the list. Options from the list are presented one at a time. After each option is presented, participants decide whether to take or decline the current

option. If they decline the current option, they can never return to it and they must take one of the subsequent options. If they reach the end of the list they must take the last option. The payoff depends on the rank of the chosen option from the full list of choices, seen and unseen. We formulated the task as a series of purchases and interview situations that varied on a single parameter that the participants were told to optimize, (e.g. maximizing square footage when buying a house). We rewarded participants for finding any of the top 3 options (best = \$5, second best = \$3, third best = \$1). We also used list lengths of 8 and 12, and informed the participants of the list length before they searched. We formulated a Markov Decision Process (MDP) model to estimate the value of individual options as they arrived and ran 32 subjects on the task while carrying out fMRI. We used the MDP model to estimate subject performance relative to optimal. On average, subjects sampled fewer options than was optimal for a given list length. We then parameterized the model for individual subjects, and used it to generate value estimates for each option the participants saw. These value estimates were used as parametric regressors on the fMRI data. Value estimates correlated with activation in the insula. Taking an option versus continuing to search activated a ventral-striatal, orbitofrontal network.

II-61. Single neuron contributions to motor variability

Kris Chaisanguanthum
Helen Shen
Philip N. Sabes

CHAIANG@PHY.UCSF.EDU
HSHEN@PHY.UCSF.EDU
SABES@PHY.UCSF.EDU

University of California, San Francisco

Movements cannot be repeated perfectly; some amount of motor variation is inevitable. Understanding its variation is critical to understanding the way neural processing drives motor behavior. We study the relationship between motor variation and the patterns of neural activity in brain areas that play a central role in the transformation from perception to action—the dorsal premotor cortex (PMd) and primary motor cortex (M1)—to determine how variability in the response of neurons driving motor behavior manifests in those behaviors. We recorded from large populations of neurons in macaque PMd and M1 as animals performed a simple visually-cued center-out reaching task, and related the measured activity during motor preparation (and other epochs of the task) of these neurons to cued stimuli and the produced movements. In particular, we make the distinction between the way each neuron is tuned to represent the (visual) stimulus, i.e., the direction of the presented reach target ("population encoding"), versus how its activity is "tuned" to the movement itself, i.e., how the variation of that cell's activity predicts variations in motor behavior (which reflects the mechanics of how a stimulus estimate is made from the activity of these neurons, or, "population decoding"). We find that for many cells in these brain areas, the relationship between a cell's tuning for input versus its tuning for output is nontrivial. We then develop a simple model of stimulus encoding and downstream population decoding and, using this model, show what the relationship between stimulus and behavior tuning reveals about the identity and connectivity of the cells within the network. Our results have implications about the structure and function of the neural circuit implicated in motor planning, as well as about population encoding and decoding in general.

II-62. An origin for coarticulation of speech sequences in human sensory-motor cortex

Kristofer Bouchard^{1,2}
Keith Johnson²
Edward Chang³

KRIS@PHY.UCSF.EDU
KEITHJOHNSON@BERKELEY.EDU
CHANGED@NEUROSURG.UCSF.EDU

¹University of California, San Francisco

²UCB

³UCSF-UCB

The ability of humans to communicate through spoken language depends on the capacity to produce a large variety of precise movements in rapid sequence, making it among the most complicated sequential behaviors found in nature. During connected speech, phonemes are coarticulated, meaning that production is dependent on the surrounding phonemic sequence. Coarticulation is a central reason why connected speech is not simply the concatenation of discrete units, but instead reflects a smoothed trajectory through the phoneme sequence. For example, the formant structure of vowels fluctuates between utterances and is sensitive to the sequential context in which it appears. Here, we investigate how single-trial fluctuations of vowel formants are generated by speech sensory-motor cortex (vSMC), and examine if these cortical signals depend on the upcoming and preceding phoneme (anticipatory and carry-over coarticulation, respectively). To this end, we recorded neural activity from the surface of vSMC using high-density electrocorticography (ECoG) in neurosurgical patients during the production consonant-vowel (CV) syllables. We found that population decoding of spatial patterns of activity allowed for accurate prediction of particular vowel formants on a trial-by-trial basis. Decoders based on cortical activity at during vowel times could accurately predict a significant fraction of the variability within a given vowel. Interestingly, decoding performance of vowel formants extended well into the consonant phase. We show that a portion of carry-over coarticulation of vowel formants is related to immediately preceding cortical activity, demonstrating that the vSMC activity generating vowels depends on the preceding consonant. Likewise, significant decoding of vowel formants remained during the consonant phase after removing the effect of carry-over coarticulation, demonstrating that vSMC activity generating consonants depend on upcoming vowels. Together, these results demonstrate that cortical signals during the generation of phonemes reflect the surrounding sequential context, and therefore the vSMC activity generating phonemes is non-unitary.

II-63. Large-scale optical imaging reveals structured network output in isolated spinal cord

Timothy Machado^{1,2}
Liam Paninski¹
Thomas M. Jessell^{1,3}

TAM2138@COLUMBIA.EDU
LIAM@STAT.COLUMBIA.EDU
TMJ1@COLUMBIA.EDU

¹Columbia University

²Howard Hughes Medical Institute

³HHMI

Isolated neonatal mouse spinal cord contains intact recurrent neural circuits that can generate ordered patterns of periodic population activity (termed fictive locomotion) following experimentally controlled stimulation. Taking advantage of known genetic entry points into the circuit, ablation experiments have demonstrated that specific interneuronal subtypes are necessary to produce and maintain fictive locomotion. However, much remains uncharacterized including the precise structure of fictive locomotion at single-cell resolution, the amount of variance across groups of individual motor neurons that share common muscle targets, and the robustness of phase tuning to network perturbations. In this study, we measured motor neuron activity using large-scale, cellular resolution calcium imaging across hundreds of retrogradely identified motor neurons. Spike inference methods for phase estimation were developed, and were validated in each experiment using antidromic stimulation to generate data where ground-truth phase and spiking information were known. Neurons with the same muscle targets fired in phase with one another, while neurons innervating synergist muscles (quadriceps and ankle flexors) had consistently shifted burst times. Neurons innervating antagonist muscles (quadriceps and hamstrings) reliably fired out of phase. Notably, groups of motor neurons that fired asynchronously were found intermingled at each lumbar spinal segment, demonstrating that the recruitment of motor neurons during fictive locomotion is considerably more complex than simple rostrocaudal alternation. In order to assess the robustness of the observed phasic activity to frequency modulation, we lowered the frequency of fictive locomotion through adenosine application. We observed no change in the phase of individual neurons despite a drop in the rhythm frequency. These results reveal complexity in the specificity of motor activation patterns in isolated spinal circuits and set the stage for work that will examine the role of local circuit interneurons in generating motor activity.

II-64. Novel motion illusion evidence for perception-action coupling

Dongsung Huh
University College, London

HUH@GATSBY.UCL.AC.UK

The motor theory of perception proposes that there is a tight coupling between the processing of sensory input and the generation of motor output. Some of the strongest evidence for this comes from the visual perception of movement since subjects are particularly sensitive to biological motion. One of the simplest and most powerful examples is the speed illusion of single dot movements — the apparent fluctuation in the speed of a dot which is actually moving uniformly along an elliptical path. The motion only appears uniform if the dot's speed is actually modulated as the one-third power of the radius of curvature ($v(t) \propto r(t)^{1/3}$). This is exactly the relationship between speed and curvature that is evident when human subjects draw or trace ellipses (Viviani and Stucchi 1992). However, we have recently shown that this power law relationship changes when non-elliptical figures are drawn. Instead, there is a whole family of power-laws with exponents ranging between 0 and 2/3, depending on the shape of the movement paths (Huh and Sejnowski 2012). Here, we tested the perceptual counterpart of this generalized family. Subjects were shown single dots moving along various curved paths, and asked to alter the speed profile until the motion looked uniform. The path shape was determined such that the log-curvature would vary along the curve according to a single frequency of oscillation (Figure 1). Preliminary results agree with the prediction (Equation 1): the exponent of the power-law which the subjects perceive to have uniform speed is not constant, but decreases with the frequency of oscillation of the curvature. This result provides stronger evidence for the coupling between perception and action, and suggests exploiting other experimental results in visual motion to generate conjectures about movement regularities.

II-65. A brain-machine interface for control of medically-induced coma

Maryam Shanechi^{1,2}
Jessica Chemali²
Max Liberman²
Ken Solt²
Emery Brown¹

SHANECHI@MIT.EDU
JESSICA.CHEMALI@GMAIL.COM
MAX.Y.LIBERMAN@GMAIL.COM
KSOLT@PARTNERS.ORG
ENB@NEUROSTAT.MIT.EDU

¹Massachusetts Institute of Technology

²Massachusetts General Hospital

A medically-induced coma is a drug-induced state of profound brain inactivation and unconsciousness used to facilitate recovery following traumatic and hypoxic brain injuries and to treat epilepsy that is refractory to conventional drug therapies. The state of coma is maintained by administering an intravenous infusion of anesthetics (e.g., propofol) to target a pattern of burst suppression on the electroencephalogram (EEG). Burst suppression consists of bursts of electrical activity alternating with periods of isoelectricity called suppression. The coma may need to be maintained for several hours or days, a period considerably longer than any human operator can be expected to maintain tight control. Currently the anesthetic infusion rates are adjusted manually to control the burst suppression level. Here we present a brain-machine interface (BMI) for automatic control of medical coma that selects the real-time drug infusion rate based on EEG observations and can precisely control the burst suppression level in real time in rodents. To develop the BMI, we quantify the burst suppression level by introducing the concept of burst suppression probability (BSP) that defines the brain's instantaneous probability of being in the suppressed state. To characterize the effect of the anesthetic on BSP, we use a linear two-dimensional compartment model whose parameters we fit in experiments. The BMI consists of two main components: an estimator that computes the BSP from the EEG and a controller that uses this estimate as feedback to adjust the drug infusion rate and achieve a target BSP level. We derive a two-dimensional state-space algorithm to estimate the BSP in real-time from the EEG segmented into a binary time-series. We then derive a stochastic controller using both a linear-quadratic-regulator strategy and a model predictive control strategy. The BMI achieves precise control of time-varying target levels of burst suppression in individual rats in real time.

II-66. Spike-triggered local field potential as an indicator of the human epileptogenic zone

Beth Lopour
Itzhak Fried
Dario Ringach

University of California, Los Angeles

BETHLOPOUR@UCLA.EDU
IFRIED@MEDNET.UCLA.EDU
DARIO@UCLA.EDU

When medication is ineffective, resection of epileptic brain tissue can result in marked reduction or cessation of seizures. The goal of the surgery is to remove the minimal amount of tissue necessary to produce seizure freedom; accurate characterization of the epileptogenic zone is crucial for success. The surgical plan is formulated based on data from a variety of modalities, including electrophysiological recordings during seizures, which often necessitates prolonged inpatient hospitalization. Here we aim to analytically identify the epileptogenic zone through the analysis of interictal (between seizures) measurements of single-unit activity and local field potential (LFP) data. Previous work has examined changes in single-unit firing rates relative to interictal spikes or the initiation and spread of seizures. However, the heterogeneity of spike rate responses both inside and outside the seizure onset zone makes these data difficult to interpret. Our analysis, based on spontaneous activity, uses the spike-triggered LFP to characterize the relationship between single-unit and regional aggregate activity. Using depth electrodes bilaterally implanted in humans for the purposes of pre-surgical evaluation, we identify correlations between single-unit spike activity in one region and LFP activity (measured by the wavelet amplitude) in a second. This serves as a general measure of functional connectivity. Preliminary data from six patients suggest that regions in the epileptogenic zone, as identified clinically, exhibit the largest amplitude of spike-correlated LFP. This LFP activity correlates with multiple single-unit spike locations including the contralateral hemisphere, consistent with the idea that heterogeneous, distributed networks contribute to seizure generation. By comparing our predictions to the clinical outcome, we can further our understanding of epileptic networks in the human brain while helping to improve the outcomes of resective surgery and reducing the duration of inpatient hospitalization required for evaluation.

II-67. Directed communication-through-coherence during synchronous transients

Agostina Palmigiano^{1,2}
Annette Witt¹
Demian Battaglia¹
Theo Geisel¹

AGOS@NLD.DS.MPG.DE
ANNETTE@NLD.DS.MPG.DE
DEMIAN@NLD.DS.MPG.DE
GEISEL@NLD.DS.MPG.DE

¹MPI for Dynamics and Self-Organization

²Bernstein Center for Computational Neuroscien

The mechanisms that mediate fast, flexible and precisely targeted switching between communication pathways in the brain are not yet fully understood. The communication through coherence hypothesis states that dynamic coherence between oscillatory neural activity allows pliable pathway selectivity, allowing manifold functional connectivities to stem from fixed structural ones. In this work we investigate through spiking network models how small motifs of interacting local populations can set their collective oscillatory activity into multistable phase-locked patterns. Due to spontaneous symmetry breaking, such dynamic states display out-of-phase locking in which a hierarchy of phase-leading and lagging areas emerge, despite symmetric rules for structural connections. We show that the inter-areal information flow is determined by this dynamics. More specifically, we focus on systems consisting of two identically coupled randomly connected spiking networks of inhibitory and excitatory neurons with delays, heterogeneous parameters and realistic statistics. Each network can be tuned to obtain gamma band oscillations. The degree of synchronization smoothly varies across increasing local delayed inhibition. In these symmetrically coupled networks, we observe that out-of-phase locking is associated to anisotropic information flow with a dominant direction from leader to laggard areas, as revealed by a transfer entropy analysis of simu-

lated LFPs. Moreover, we show that the degree of synchrony of the ongoing oscillations regulates the time the system spends in a fixed laggard-leader configuration. Thus, for nearly asynchronous states, windows of directed communication appear as short transients, during which the effective information flow shows the same anisotropic properties as for strong synchrony. We finally explore how input modulations can be used to select or switch the dominant directionality of information flow. We hypothesize that similar dynamic mechanisms might underlie the flexible switching of selective attention or prioritized selection of alternative computations within the same network.

II-68. State dynamics of the epileptic brain and the influence of seizure focus

Sridevi Sarma¹
Sam Burns¹
Sabato Santaniello¹
William Stan Anderson²

SREE@JHU.EDU
SBURNS9@JHU.EDU
SSANTAN5@JHU.EDU
WANDERS5@JHMI.EDU

¹Johns Hopkins University
²JHH

Communication between specialized regions of the brain is a dynamic process that allows for different connections to accomplish different tasks. While the content of the interregional communications is complex, we hypothesize that the pattern of connectivity (i.e., which regions are communicating) may be described in a lower dimensional state-space that contains information about the brain functions. In epilepsy, seizures elicit changes in connectivity, whose pattern sheds insight into the nature of seizures and the seizure focus (minimal brain region generating seizures). We investigated the temporal evolution of the brain connectivity before, during, and after seizure by applying network-based analysis on continuous multi-day subdural electrographic recordings (ECoG) from 12 medically refractory epilepsy patients. In our analysis, each ECoG electrode was considered a node in a graph, and edges between pairs of nodes were weighted by their ECoG coherence in the beta frequency band. Then, the network was represented by a connectivity matrix, which was broken down using eigenvalue decomposition. The leading eigenvector was interpreted as a summary of the network connectivity structure, and was tracked over time and clustered to uncover a finite set of brain states. Across all patients, we found that (i) the network connectivity is structured and defines a finite set of brain states (2-3 during non-seizure periods, 2-7 during seizure), (ii) seizures are characterized by a significantly consistent progression of states, (iii) the focus is isolated from the surrounding regions at the seizure onset (signature of the focus) and becomes most connected in the network towards the seizure termination, and (iv) this particular seizure state and the corresponding network structure can be used to detect the focus with high specificity and sensitivity. Results suggest that a finite-dimensional state-space model may characterize the dynamics of the epileptic brain and the seizure focus, and may ultimately predict seizure onsets.

II-69. The role of adaptation in intrinsic dynamics of primary visual cortex

Yashar Ahmadian¹
Sandro Romani¹
Amiram Grinvald²
Misha Tsodyks²
Ken Miller¹

YA2005@COLUMBIA.EDU
SR2944@COLUMBIA.EDU
AMIRAM.GRINVALD@WEIZMANN.AC.IL
MTSODYKS@GMAIL.COM
KEN@NEUROTHEORY.COLUMBIA.EDU

¹Columbia University
²Weizmann Institute

Cortical spontaneous activity (SA) in the absence of sensory inputs exhibits a complex spatiotemporal structure. SA may play an important role in sensory processing; sensory stimuli interact with and are modulated by SA, while SA itself is shaped by past experience. A well-known instance is offered by SA in V1. Voltage sensitive

dye imaging (VSDI) experiments in anesthetized animals have shown that SA in V1 has stronger correlations with "orientation maps", i.e. activity patterns evoked in response to full-field oriented stimuli, than with control patterns of similar spatial structure. However, the detailed temporal statistics of SA have not been thoroughly studied. We analyzed VSDI data from cat V1. Using data from evoked orientation maps, we employ a temporal analysis of the SA dynamics in a low-dimensional "orientation space". This low-dimensional representation captures the most significant aspects of the data, and provides a valuable tool for the study of SA dynamics and a thorough comparison of different models for it. We found that the SA dynamics is characterized by a sequential progression through neighboring orientations, composed of intervals of smooth drift in orientation interrupted by switches to the orthogonal orientation. A common model explains SA in V1 as shaped by "Mexican hat" recurrent interactions between neurons of different preferred orientations. The SA dynamics in this model does not feature such transitions to the orthogonal orientation. However, we show that the same model endowed with adaptation, either neuronal or synaptic, can account for such switches. Importantly, we show that adaptation is crucial in making the model robust to non-uniformities in the distribution of preferred orientations, which bias the switching to occur typically between the most represented orientation and that orthogonal to it. Indeed the analyzed experimental data does show such a bias.

II-70. Strategies for optomotor control in free flying *Drosophila*

Leonardo Mesquita¹
Shiva Sinha²
Rob de Ruyter van Steveninck²

LGMESQUI@INDIANA.EDU
SRSINHA@INDIANA.EDU
DERUYTER@INDIANA.EDU

¹Indiana University

²Indiana University Bloomington

Studies of the insect optomotor response provide insight both into specific computations by the visual system and into strategies for behavior. Traditionally, fruit fly optomotor behavior was studied through torque meter experiments, whereas recently video tracking methods have been used. Here we present a novel non-video based opto-mechanical technique, which enables tracking of single insects over large volumes (~20000 cm³) at high spatial (~1mm) and temporal (~1ms) resolution for extended periods (>1 hour) of time. The method also allows us to record high resolution video and deliver behavioral conditioning heat shocks in free flight. Further, we control an external visual stimulus projected on a cylinder surrounding the flight arena. We revisit a modified version of a classical torque meter experiment, and record free flight behavior in the presence of a rotating visual stimulus. Whereas the torque meter limits behavior to one degree of freedom, here the fly can control its slip speed by selecting combinations of linear speed and radial position within the cylinder. For our analysis we record the 3D position of the fly, the angular position of the visual stimulus and video of all flight segments. The statistics of angular speed and radial position (both defined with respect to the center of the cylinder) show that the mode of the slip speed distribution equals the external stimulus velocity, indicating that the fly attempts to track the visual stimulus. To accomplish this, the fly combines two strategies, varying both its linear speed and its radial position: At increasing stimulus speeds, the fly tends to smaller flight radius and higher linear velocity. Interestingly, the slip speed distributions more or less scale with stimulus velocity, suggesting that the fly's tracking strategy approximates scale invariant behavior. Studies are underway to understand the implications of these observations in more detail.

II-71. Supervised cue calibration relies on the multisensory percept

Adam Zaidel
Wei Ji Ma
Dora Angelaki
Baylor College of Medicine

AJZAIDEL@GMAIL.COM
WJMA@BCM.EDU
ANGELAKI@CABERNET.CNS.BCM.EDU

Multisensory plasticity enables our senses to dynamically adapt to one another and to the environment. Although most pronounced during development, plasticity is now believed to be a normal capacity of the nervous system throughout our lifespan. But, despite its importance in normal and abnormal brain function, the rules governing multisensory plasticity are still a mystery. Multisensory calibration without external feedback ('unsupervised' calibration) functions to reduce or eliminate discrepancies between the cues, and is largely independent of cue reliability. But calibration naturally occurs with environmental feedback regarding cue accuracy ('supervised'). Hence to provide a comprehensive theory for multisensory calibration, simultaneous manipulation of both cue accuracy and cue reliability is required. Here we measured the combined influence of cue accuracy and cue reliability on supervised multisensory calibration, using discrepant visual and vestibular motion stimuli. Five monkeys were trained to perform a heading discrimination task in which they were required to report whether self-motion was to the right/left of straight ahead. Relative cue reliability was controlled by coherence of the visual stimulus (optic flow), and cue accuracy was controlled by external feedback. When the less reliable cue was also inaccurate, it alone was calibrated. However, when the more reliable cue was inaccurate, cues were yoked and calibrated together in the same direction. Strikingly, the less reliable cue shifted away from the direction indicated by the external feedback, thus becoming less accurate. A computational model in which supervised and unsupervised calibration work in parallel, where the former only has access to the multisensory percept, but the latter can calibrate cues individually, accounts for the observed behavior. This suggests that individual cue information is not accessible to the mechanism of supervised calibration. In combination, supervised and unsupervised calibration could ultimately achieve the optimal solution of both external accuracy and internal consistency.

II-72. Population decoding in rat barrel cortex

Ehsan Arabzadeh¹

Mehdi Adibi²

Colin WG Clifford³

EHSAN.ARABZADEH@ANU.EDU.AU

MEHDI@UNSW.EDU.AU

COLIN.CLIFFORD@SYDNEY.EDU.AU

¹Australian National University

²University of New South Wales

³University of Sydney

How do cortical neurons encode sensory stimuli, and how is the population response decoded by downstream neurons? Theoretical studies reveal that even weak correlations in activity distributed across large populations can significantly reduce the efficiency of coding. Such correlation in trial-to-trial response fluctuations, or noise correlation, has often been characterized for pairs of neurons. Here we investigate synergy and redundancy across multiple simultaneously recorded single neurons (up to 11) using various decoding procedures. We recorded the activity of barrel cortex neurons (73 single-units) using a 32-channel 4-shank probe, while applying sinusoidal vibrations (80Hz; 0 to 33um amplitude) to the corresponding whisker. Using the area under ROC (AUROC) from Signal Detection Theory, we compared the performance of multiple decoding approaches using neuronal variability, co-variability, and signal correlation. Within each population size, we compared the AUROC values for individual neurons, when their responses were simply pooled together, or when their responses were integrated after applying an optimum set of weights. Optimum weights were found by the analytical solution that maximized the average signal to noise ratio based on Fisher Linear Discriminant analysis. This gave us a biologically plausible decoder that integrates neuronal activity after applying different synaptic weights provided an optimal 'read out' of the sensory signal. Neurons exhibited a positive correlation in their trial-to-trial activity. This noise correlation was in the signal direction, and thus detrimental to coding efficiency. Whisker stimulation decreased noise correlation. At a population size of 8 single units, the optimum decoder achieved 96.8% (*pm*21%) improvement over pooling. Decorrelating neurons by trial shuffling revealed that, unlike pooling, the performance of the optimum decoder was minimally affected by noise correlation. AUROC on untrained trials was 0.97 that on trained trials: a remarkable degree of generalization.

II-73. Hemodynamic responses in the somatosensory cortex during locomotion

Bingxing Huo
Yurong Gao
Patrick Drew

BIH5103@PSU.EDU
YURONG.G@GMAIL.COM
PJD17@PSU.EDU

Pennsylvania State University

Increases in cerebral blood volume (CBV) are correlated with neural activity in anesthetized animals, but the effects of natural behavior on cerebral hemodynamics are not well understood. We quantified the spatial and temporal dynamics of CBV in the somatosensory cortex of awake, behaving mice using intrinsic optical signal (IOS) imaging. The mouse was head-fixed on top of a spherical treadmill. During bouts of voluntary locomotion, we observed a regionally specific decrease in reflectance, indicating a localized increase in CBV. Using the motion impulses as the input, we solved for a pixel-wise linear, time-invariant impulse response function that described the hemodynamic response (Silva et al., 2007) to locomotion. The impulse response at any location on the cortex could be represented by a weighted summation of two exponential decay functions: a fast (3.33-second time constant, 'arterial') and a slow (100-second time constant, 'venous') component (Kim and Kim, 2010; Drew et al., 2011). We found that with the appropriately chosen weights, the CBV response to locomotion was well fit ($R^2=0.5-0.8$) by the sum of the arterial and venous components. The goodness-of-fit was consistent across trials and animals. Pharmacologically occluding or blocking heart rate increases during locomotion did not alter the weight of the arterial component significantly, indicating that the cerebral hemodynamic response during locomotion was not of a peripheral origin. Using two-photon laser scanning microscopy (2PLSM), we quantified the microvascular basis of this impulse response function. The diameter change of an artery or vein during locomotion could be well fit by convolving the motion impulses with a fast or slow exponential decay kernel. The time constants of the kernels were consistent with those of the IOS data. Thus, the hemodynamic response in the somatosensory cortex can be accurately described with a simple linear convolution model, which has a clear interpretation in microvascular dynamics.

II-74. Rapid detection of odors based on fast, sampling-based inference

Agnieszka Grabska-Barwinska¹
Jeffrey Beck²
Alexandre Pouget³
Peter Latham¹

AGNIESZKA.GRABSKABARWINSKA@GMAIL.COM
JBECK@BCS.ROCHESTER.EDU
ALEX@CVS.ROCHESTER.EDU
PEL@GATSBY.UCL.AC.UK

¹University College, London

²University of Rochester

³University of Geneva

Sensory processing is hard because we have to infer high level quantities from low level features. In vision, for example, we have to infer (at the very least) which objects are present based on colors, shapes, textures etc. The problem is exacerbated by the fact that there are a very large number of a priori possible objects in the world, and some combination of them will be needed to explain the low level features. The latter observation means the search space is combinatorially large. Nevertheless, an almost universal feature of sensory processing is that it is fast: organisms can typically make sense of a scene, be it visual, auditory, olfactory, etc., in a few hundred milliseconds. Here we propose an algorithm that can, in principle, search a combinatorially large space of objects in parallel. The algorithm uses Gibbs sampling to estimate which objects are present, and a stochastic version of maximum a posteriori estimation to fill in the low level features. We formulate the algorithm so that networks of biologically plausible neurons could implement it in parallel and asynchronously. This is in contrast to the standard implementation on a serial computer, which is slow and often scales poorly with problem size. We apply the algorithm to olfaction, a sensory modality without the structure of, say, vision or audition, but still retaining some of the hard features (inferring combinations of odors based on a noisy and compressed signal).

As expected, given the formulation, the algorithm is indeed fast and accurate when there are a large number of odors present.

II-75. Neuronal nonlinearity explains greater visual spatial resolution for dark than for light stimuli

Jens Kremkow
 Jianzhong Jin
 Stanley Jose Komban
 Yushi Wang
 Reza Lashgari
 Michael Jansen
 Xiaobing Li
 Qasim Zaidi
 Jose-Manuel Alonso
 SUNY-Optometry

JENS@KREMKOW.INFO
 JJIN@SUNYOPT.EDU
 JKOMBAN@SUNYOPT.EDU
 YUSHIWANG@SUNYOPT.EDU
 RLASHGARI@SUNYOPT.EDU
 MJANSEN@SUNYOPT.EDU
 XLI@SUNYOPT.EDU
 QZ@SUNYOPT.EDU
 JALONSO@SUNYOPT.EDU

Astronomers and physicists noticed centuries ago that visual spatial resolution is higher for dark than light stimuli, but the neuronal mechanisms for this perceptual asymmetry remain undetermined. We investigated the neuronal mechanisms by recording extracellular responses of single ON- and OFF-center cells in the visual thalamus (LGN) and multi-unit activity in the visual cortex (V1) of cat. We found that receptive fields of ON-center cells were larger than receptive fields of OFF-center cells when mapped on binary backgrounds (light targets on dark backgrounds and dark targets on light backgrounds). Similar differences were found in V1 multi-unit activity. Strikingly, when receptive fields were mapped on gray backgrounds, these differences disappeared in the LGN and were slightly reversed in V1. Thus, the difference in spatial resolution reported perceptually corresponds to neuronal differences in the LGN and V1. Further, the difference in spatial resolution is not constant and changes dynamically with the background luminance. We hypothesized that a nonlinear encoding of luminance increments and decrements could explain these differences. We found that OFF-center cells increase their responses roughly linearly with luminance decrements, independent of the background luminance. In marked contrast, ON-center cells saturate their responses with small increases in luminance and require bright backgrounds to approach the linearity of the OFF-center cells. V1 neurons showed qualitatively the same behavior. Although the integration of lights becomes more linear on gray backgrounds, a pairwise comparison in V1 neurons showed that responses to increments saturate earlier than responses to decrements. We show that this nonlinearity can explain the larger receptive fields and lower spatial resolution of ON channel cells: receptive fields are more blurred in the ON- than OFF- channel. Preliminary data from local field potential recordings in awake primate support the findings, suggesting a general principle of how ON and OFF channels process visual information.

II-76. Encoding of natural scene statistics in the primary visual cortex of the mouse

Emmanouil Froudarakis¹
 Philipp Berens²
 R. James Cotton¹
 Alexander S. Ecker³
 Peter Saggau¹
 Matthias Bethge⁴
 Andreas Tolias¹

EFROUD@CNS.BCM.EDU
 PHILIPP.BERENS@UNI-TUEBINGEN.DE
 RCOTTON@CNS.BCM.EDU
 ALEXANDER.ECKER@UNI-TUEBINGEN.DE
 PETER@CNS.BCM.EDU
 MBETHGE@TUEBINGEN.MPG.DE
 ATOLIAS@CNS.BCM.EDU

¹Baylor College of Medicine
²BCCN

³University of Tübingen

⁴Max Planck Institute

The visual system has evolved to process ecologically relevant information in the organism's natural environment, and thus it is believed to have adapted to its statistical properties. The most informative components of natural stimuli lie in their higher order statistical structure. If the primary visual cortex has indeed adapted to this higher order structure — as has been posited by theoretical studies over the last 20 years — neural responses to stimuli which differ in their statistical structure from natural scenes should exhibit pronounced deviations from responses to natural scenes. Theoretical studies argue for a sparse code for natural scenes, where only a few neurons need to be active simultaneously in order to encode visual information. However, it has been difficult to assess the sparseness of the neural representation directly and measure the 'population sparseness' in neural populations. Here we use 3D random access and conventional 2D two-photon imaging in mice to record populations of hundreds of neurons while presenting natural movies and movies where the higher order structure had been removed (phase scrambled). This technique allows assessing directly how sparse the representation of natural scenes in V1 really is and how this impacts the functional properties of the population code. First, we show that a decoder trying to discriminate between neural responses to different movie segments performs better for natural movies than for phase scrambled ones (nearest-neighbor classifier). Second, we show that this decoding accuracy improvement could be mainly explained through an increase in the sparseness of the neuronal representation. Finally, to explain the link between population sparseness and classification accuracy, we provide a simple geometrical interpretation. Our results demonstrate that the higher order correlations of natural scenes lead to a sparser neural representation in the primary visual cortex of mice and that this sparse representation improves the population read-out.

II-77. Recording the entire visual representation along the vertical pathway in the mammalian retina

Philipp Berens¹

Tom Baden^{1,2}

Matthias Bethge³

Thomas Euler⁴

PHILIPP.BERENS@UNI-TUEBINGEN.DE

THOMAS.BADEN@UNI-TUEBINGEN.DE

MBETHGE@TUEBINGEN.MPG.DE

THOMAS.EULER@UNI-TUEBINGEN.DE

¹BCCN

²CIN

³Max Planck Institute

⁴BCCN & CIN

In the retina, the stream of incoming visual information is split into multiple parallel information channels, represented by different kinds of photoreceptors (PRs), bipolar (BCs) and ganglion cells (RGCs). Morphologically, 10-12 different BC and about 20 different RGC types have been described. Here, we record from all cells in the vertical cone pathway, including all PR, BC and RGC types, using 2P imaging in the mouse retina. We show that BCs and RGCs can be clustered into functionally defined classes based on their Ca²⁺-responses to simple light stimuli. For example, we find 8 functional BC types, which match anatomical types and project to the inner retina in an organized manner according to their response kinetics. The fastest BC types generate clear all-or-nothing spikes. In addition, we find more than 15 functional RGC types, including classic ON- and OFF as well as transient and sustained types. We verify the functional clustering using anatomical data. This dataset allows us to study the computations performed along the vertical pathway in the mammalian retina and to obtain a complete sample of the information the retina sends to the brain.

II-78. Spontaneous emergence of simple and complex receptive fields in a spiking model of V1

Eugene Izhikevich¹

Filip Piekniewski¹

Jayram Moorkanikara Nageswaran¹

Csaba Petre¹

Micah Richert¹

Sach Sokol²

Philip Meier¹

Marius Buibas¹

Dimitry Fisher¹

Botond Szatmary¹

¹Brain Corporation

²Johns Hopkins University

EUGENE.IZHKEVICH@BRAININCORPORATION.COM

PIEKNIIEWSKI@BRAININCORPORATION.COM

NAGESWARAN@BRAININCORPORATION.COM

CSABA.PETRE@BRAININCORPORATION.COM

RICHERT@BRAININCORPORATION.COM

SSOKOL3@JHU.EDU

MEIER@BRAININCORP.COM

BUIBAS@BRAININCORPORATION.COM

FISHER@BRAININCORPORATION.COM

Brain Corporation is engaged in a multi-year project to build a spiking model of vision, paying special attention to the anatomy and physiology of the mammalian visual system. While it is relatively easy to hand-tune V1 to get simple and complex cells, it is not clear how to arrange connectivity in other cortical areas to get appropriate receptive fields, or what the appropriate receptive fields even should be. Instead of pre-wiring cortical connectivity according to a computational theory of how vision should work, we start with a generic ‘tabula rasa’ spiking model having multiple cortical layers and neuronal types (single-compartment RS, FS, LTS cells). The goal is to find the anatomical and physiological parameters so that the appropriate connectivity emerges through STDP and visual experience. Since we know exactly what kind of receptive fields and visual responses are in V1, we build a smaller model of retina-LGN-V1 pathway and tune the STDP parameters so that the expected responses emerge. Once we trust what we see in V1, we are ready to copy and paste the cortical model to implement V2, V3, V4, and IT areas with the hope that useful connectivity, receptive fields, and visual responses emerge. Our large-scale simulations of the spiking model of the visual system show spontaneous emergence of simple and complex cells, orientation domains, end-stopping receptive fields, extra-classical receptive fields with tuned surround suppression, color opponency that depends on the eccentricity of the receptive field, contrast invariance, and many other features that are routinely recorded in V1. Since the visual model exhibits micro- and full saccades, we observe perceptual behavior, such as the emergence of bottom-up (pop-out) attention. The model underscores the importance of spike- timing dynamics, inhibition, saccadic mechanism, and it imposes important restrictions on the possible types of STDP to model early visual processing.

II-79. Eye’s imaging process explains ganglion cells anisotropies

Daniela Pamplona^{1,2}

Jochen Triesch¹

Constantin A Rothkopf^{1,3}

PAMPLONA@FIAS.UNI-FRANKFURT.DE

TRIESCH@FIAS.UNI-FRANKFURT.DE

ROTHKOPF@FIAS.UNI-FRANKFURT.DE

¹Frankfurt Institute for Advanced Studies

²Johann Wolfgang Goethe University

³University of Osnabrueck

While the statistical regularities of the natural environment have been shown to determine properties of neurons in sensory systems, much less work has considered the influence of the imaging system itself. Here we use a model of the geometric and optical transformations that shape the local input signal statistics to the visual system across the visual field. Under this model, we have recently shown that the second order intensity statistics of naturalistic images vary systematically with retinal position [1]. In the present study, we investigate the consequences of the imaging process on the properties of retinal ganglion cells according to a generative model encompassing several previous approaches[2,3,4]. First, we generated artificial naturalistic scenes and quantified the local correlational

structure under perspective projection onto a spherical surface and optical blur. These distributions show the strong influence of the imaging process on the image statistics as a function of eccentricity and radial distance from the center of projection. Model ganglion cells were computed with one generative model that generalizes the optimal whitening filter [2], the encoding and decoding with additional channel noise [3], and a metabolic constraint [4] models. We quantified the parameters of the resulting model ganglion cells and compared them to previous empirical studies including those showing radial dependencies of orientation and size on the eccentricity. We conclude by providing a detailed quantitative analysis of model retinal ganglion cells across the visual field. Our results agree with previous empirical data reporting anisotropies in retinal ganglion cells' receptive fields and thereby provide a functional explanation of these properties in terms of optimal coding of sensory stimuli [5,6]. [1] Pamplona et al (under review) [2] Dong and Atick, 1995 [3] Doi, 2006 [4] Vincent and Baddeley, 2003 [5] Croner and Kaplan, 1995

II-80. Circuit mechanisms revealed by spike-timing correlations in macaque area MT

Xin Huang

XHUANG43@WISC.EDU

University of Wisconsin, Madison

Interactions among intricately-connected neurons in cortical circuits help to shape the activity of neuronal populations. How these neural circuits operate and influence the coding of sensory information is not well understood. We analyzed neuron-neuron correlations to uncover the dynamic interactions between excitatory and inhibitory neurons in the visual cortex. We recorded simultaneously from pairs of motion-sensitive neurons in extrastriate area MT of macaque monkeys and used cross-correlations in the timing of spikes between neurons to gain insights into cortical circuitry. We characterized the time course and stimulus-dependency of the cross-correlogram (CCG) for each pair of neurons and of the auto-correlogram (ACG) of the individual neurons. For some neuron pairs, the CCG showed negative flanks that emerged next to the central peak during stimulus-driven responses. Similar negative flanks appeared in the ACG of many neurons. Negative flanks were most prevalent and deepest when the neurons were driven to high rates by visual stimuli that moved in the neuron's preferred directions. The temporal development of the negative flanks in the CCG coincided with a parallel, modest reduction of the noise correlation between the spike counts of the neurons. Computational analysis of a model cortical circuit suggested that negative flanks in the CCG arise from the excitation-triggered mutual cross-inhibition between pairs of excitatory neurons. Intracortical recurrent inhibition and after-hyperpolarization caused by intrinsic, outward currents such as Ca²⁺-activated K current of small conductance (SK channels) both can contribute to the negative flanks in the ACG. In the model circuit, stronger intra-cortical inhibition helped to maintain the temporal precision between the spike trains of pairs of neurons and led to weaker noise correlations. Our results suggest a neural circuit architecture that can leverage activity-dependent intracortical inhibition to adaptively modulate both the synchrony of spike timing and the correlations in response variability.

II-81. Hierarchical shape processing and position tolerance in rat lateral visual cortex

Ben Vermaercke

Gert Van den Bergh

Florian Gerich

Hans P. Op de Beeck

KU Leuven

BEN.VERMAERCKE@GMAIL.COM

GERT.VANDENBERGH@PPW.KULEUVEN.BE

FLORIAN.GERICH@PPW.KULEUVEN.BE

HANS.OPDEBEECK@PPW.KULEUVEN.BE

Recent studies have revealed a surprising degree of functional specialization in rodent visual cortex. However, these studies fall short of establishing a functional hierarchy. We designed a study in rats that targets two hall-

marks of the hierarchical object vision pathway in primates: higher tolerance for image transformations and selectivity for behaviorally relevant dimensions. We targeted five visual areas from primary visual cortex (V1) over areas LM, LI, LL, up to lateral occipito-temporal cortex (TO). We examined the responses of single neurons in these regions to six simple shapes used previously to probe monkey anterior infero-temporal cortex. These shapes were slowly jittering around a particular position of the stimulus display during 4s per presentation. After delineating the receptive field (RF) of each neuron, we presented the six shapes around one (Exp. A; N=299 cells) or two (Exp. B; N=258 cells) positions within the RF. First, we quantified the selectivity of populations of neurons in each visual area using all data from Exp. A plus the most responsive position of Exp. B. Overall discrimination performance was highly significant in all visual areas, although it decreased from V1 to higher visual areas. Neighboring areas were correlated with respect to which shape pairs were easiest to discriminate. This correlation suggests that the representation of shape transforms gradually across areas. In addition, we found evidence for an increase in position tolerance along the five areas. In TO, the preference for different shapes at one position was most closely related to the shape preference at another position. Finally, we found strong correlations between TO selectivity and behavioral performance of rats in a discrimination task. These findings demonstrate that the functional specialization in lateral rodent visual cortex reflects a processing hierarchy resulting in the emergence of tolerance and complex selectivity.

II-82. Anesthesia amplifies visual responses by suppressing cortical state dynamics

Flavio Frohlich
Kristin Sellers
Davis Bennett

UNC - Chapel Hill

FLAVIO_FROHLICH@MED.UNC.EDU
KKS34@EMAIL.UNC.EDU
DAVIS.V.BENNETT@GMAIL.COM

Anesthesia is widely used in science and medicine, but we know surprisingly little about how anesthesia disrupts brain function. The effects of anesthesia on synapses and single cells have been well characterized, but how anesthesia affects cortical network dynamics is not yet understood. To investigate the impact of anesthesia on functional cortical networks, we measured the local field potential (LFP) in primary visual cortex (V1) and prefrontal cortex (PFC) of awake and anesthetized ferrets in response to artificial and naturalistic visual stimuli. We found that awake cortex exhibits visual stimulation-induced transitions between states dominated by low (delta and alpha band) or high (gamma band) frequency oscillations in the LFP, respectively. Stimulus-driven spiking in awake V1 exhibited low rates but high temporal precision. In contrast, animals anesthetized with varying concentrations of isoflurane (0.5%, 0.75%, 1.0%) lacked frequency specific up- and down-regulation of spectral power and instead displayed increased power across all frequencies, paired with massively amplified and prolonged spiking response to visual input. PFC demonstrated stimulus-induced phase locking at the onset and offset of the visual stimulus in the awake animal. Anesthesia abolished phase-resetting to visual input in PFC, suggesting interruption of long-range projections and a break-down of effective information integration across cortical areas. These results suggest that awake and anesthetized cortex exhibit distinct processing modes, in which state-defining network dynamics dominate neural activity in the awake animal but loss of these dynamics under anesthesia opens the gate to massive and unrefined sensory responses. We found that the effects of anesthesia are more complex than the commonly assumed reduction in activity levels and increase in LFP low frequency power. Anesthesia combined with sophisticated analysis strategies of cortical LFP recordings may open new avenues for understanding the mechanisms that underlie higher-order cognition and ultimately consciousness.

II-83. Computational models of contour detection: role of lateral connections, inhibition and normalization

David A Mély
Thomas Serre
Brown University

DAVID_MELY@BROWN.EDU
THOMAS_SERRE@BROWN.EDU

Contour integration is an important function thought to subtend many visual tasks. One generic computation that is common to several models of contour integration (Li, 1998; Ross et al., 2000, Ben-Shahar & Zucker, 2004) is based on the notion of ‘association fields’ whereby lateral connections between V1-like edge detectors favor patterns of co-activations that are consistent with natural image statistics (Geisler et al., 2001). Individual models, however, differ in the specific operations used (e.g., tuned vs. untuned inhibition, subtractive vs. divisive normalization, etc.). Such biological models are seldom evaluated on large datasets of natural images, which render their comparison difficult and limit our understanding of the underlying computations, and in particular, the role of inhibition. Here we implemented a standard V1 model based on populations of oriented simple and complex cells with parameters constrained by electrophysiology data. In addition, we modeled two different local inhibitory circuits, to assess the extent to which different types of inhibition drive performance in a contour detection task on a set of natural scenes. We used a standard computer vision dataset of curve fragments previously used to evaluate the performance of algorithms for the bottom-up detection of contours in natural scenes (Guo & Kimia, 2012). After optimizing the model parameters for performance, we find that the resulting model balances subtractive and divisive normalization and competes with state-of-the-art computer vision systems, notably the Pb algorithm by Martin et al. (2004). Overall these results suggest that different types of inhibition play distinct and complementary roles in regulating the activity of simple and complex cells in the primary visual cortex.

II-84. Is there a critical area size for the transition from interspersed to columnar V1 architecture?

Wolfgang Keil^{1,2}
Fred Wolf³
Juan-Daniel Florez-Weidinger⁴
Matthias Kaschube⁵
Michael Schnabel⁶
David M. Coppola⁷
Siegrid Löwel⁸
Leonard E. White⁹

WOLFGANG@NLD.DS.MPG.DE
FRED@NLD.DS.MPG.DE
CHEPE@NLD.DS.MPG.DE
KASCHUBE@FIAS.UNI-FRANKFURT.DE
M-SCHNABEL@NORTHWESTERN.EDU
DCOPPOLA10@VERIZON.NET
SIEGRID.LOEWEL@BIOLOGIE.UNI-GOETTINGEN.DE
LEN.WHITE@DUKE.EDU

¹MPI for Dynamics & Self-organization

²BCCN Göttingen

³MPI for Dynamics and Self-Organisation

⁴MPI-DS Goettingen, BCCN Goettingen

⁵Frankfurt Institute for Advanced Studies

⁶Northwestern University

⁷Randolph-Macon College

⁸Northwestern

⁹Duke University

Response characteristics of orientation-tuned neurons in the visual cortex appear to be similar in mammalian lineages widely separated in evolution. The spatial arrangement of preferences across the cortex, however, shows qualitative differences. While in primates and carnivores orientation preferences form orientation preference maps (OPMs), in rodents they are spatially interspersed. Previously, we showed that OPMs in several primate and carnivore species realize a single common design. Both designs can be explained by activity-dependent circuit self-organization and available evidence indicates that both functional organizations are insensitive to V1 size.

However, optimization principles for V1 architecture predict that V1 size could become a constraint in the spatial arrangement of orientation preferences; e.g., in carnivores and primates with V1 areas smaller than some minimal size. Here, we investigate whether optimization principles also predict an upper critical size above which a columnar design becomes preferable to an interspersed layout. We examined models in which cortical organization is assumed to optimize a composite cost function that penalizes reductions in stimulus coverage and excessive wiring length. Since interspersed layouts exhibit near optimal coverage and wiring length cost decreases with area size, interspersed layouts form robust optima for small area sizes. With increasing V1 size, the relative advantage in coverage of an interspersed over a columnar organization decreases. Given that even in interspersed layouts neurons of similar orientation preference are preferentially connected, a columnar layout is expected to reduce wiring cost relative to interspersed layouts with the reduction in wiring cost per unit area of cortex essentially constant. Thus, under a wide range of conditions there exists a critical area size above which a columnar organization is advantageous compared to an interspersed arrangement. The predicted transition from interspersed to columnar cortical design could be tested by examining preference layouts in large rodents, such as the capybara.

II-85. Modeling cortical responses to mixture stimuli reveals origins of orientation tuning variation

Robbe Goris
Eero Simoncelli
Anthony Movshon
New York University

ROBBE.GORIS@NYU.EDU
EERO.SIMONCELLI@NYU.EDU
MOVSHON@NYU.EDU

Neurons in primary visual cortex are typically selective for the orientation of bars and gratings. But there is considerable variation in the sharpness of orientation tuning. In the feedforward model of neuronal selectivity, simple cells in visual cortex inherit their basic stimulus preference from a linear filtering stage. The subsequent nonlinear rectification stage further sharpens tuning. It is not known whether variation in conventionally measured tuning widths primarily originates in the linear or nonlinear stage. To address this question, we performed a superposition experiment that allowed us to estimate the contribution of both stages. We studied the orientation selectivity of cortical neurons in anesthetized macaques using stimuli that varied in their orientation composition – our stimuli contained multiple incoherently-drifting sinusoidal gratings whose orientations were spaced at 20 deg intervals and whose orientation-dependent contrasts were chosen to match a Gaussian profile (standard deviation 0-55 deg). We fit the data with an LN cascade model with the linear stage parameterized for orientation selectivity and found that the model successfully accounts for the dependency of response gain and tuning bandwidth on stimulus bandwidth. Analysis of the model parameters revealed that both the linear and nonlinear stage contribute significantly to the variability in orientation selectivity as measured with conventional gratings. However, a larger portion of this variability originates in the linear processing stage.

II-86. Psychophysical evidence for a sampling-based representation of uncertainty in low-level vision

Marjena Popovic¹
Ralf Haefner¹
Mate Lengyel²
Jozsef Fiser³

MARJENAP@BRANDEIS.EDU
RALF@BRANDEIS.EDU
M.LENGYEL@ENG.CAM.AC.UK
FISER@BRANDEIS.EDU

¹Brandeis University

²University of Cambridge

³Central European University

Human and animal studies suggest that human perception can be interpreted as probabilistic inference that relies

on representations of uncertainty about sensory stimuli suitable for statistically optimal decision-making and learning. It has been proposed recently that the way the brain implements probabilistic inference is by drawing samples from the posterior probability distribution, where each sample consists of instantaneous activity of a population of neurons (Fiser et al, 2010). However, there is no experimental evidence thus far, showing that an internal representation of uncertainty can extend to low-level sensory attributes, nor that humans use sampling-based representations in perceptual judgment tasks. To address these questions, we created an orientation-matching task in which we measured both subjects' performance and their level of uncertainty as they matched orientation of a randomly chosen element of the previously presented stimulus. Stimuli consisted of 2-7 differently oriented line segments shown spaced evenly on a circle extending 2 degrees of the visual field. In response to the first question, we found that subjects' performance and subjective report of uncertainty were significantly correlated ($r=0.37$, $p<.001$) and that this correlation was independent of the number of oriented line segments shown. To address the second question, we varied the stimulus presentation time trial-to-trial to influence the number of samples available before making a judgment. Since samples are drawn sequentially, the prediction of the sampling-based representations is that precision of representing uncertainty will depend on the time available independent of the recorded performance. We found that decreasing the presentation time results in a significant decrease of the error-uncertainty correlation ($p<0.05$) while the performance levels remain constant. Thus, limiting the presentation time influences the reliability of uncertainty representation specifically, in agreement with sampling-based representations of uncertainty in the cortex, and in contrast with the predictions of other probabilistic representations.

II-87. Effects of attention on spatio-temporal correlations across layers of a single column in area V4

Tatiana Engel
Nicholas Steinmetz
Tirin Moore
Kwabena Boahen
Stanford University

TATIANA.ENGEL@STANFORD.EDU
NICK.STEINMETZ@GMAIL.COM
TIRIN.MOORE@GMAIL.COM
BOAHEN@STANFORD.EDU

The likely source of attentional modulation in the visual cortex are top-down inputs from fronto-parietal areas, but little is known about how top-down and sensory signals are integrated across layers within the local microcircuit of a cortical column. We aimed to elucidate the spatio-temporal pattern of interactions between cortical layers and to establish how attention affects correlations within a column of the visual cortex. To this end, we analyzed spiking activity simultaneously recorded from 16 channels across different layers of the cortex in the area V4 of monkeys engaged in an attentional task. To record from a single column, the electrode arrays were inserted into V4 perpendicularly to the cortical surface, confirmed by measured receptive field alignment. The laminar positions of the recording sites were estimated using current source density analysis, which allowed identification of four distinct functional layers and provided means to align data from multiple penetrations. To reliably estimate spike correlations in the presence of non-stationarities caused by visual stimulation and slow trial-to-trial rate fluctuations, we employed the resampling technique known as spike jitter. This technique filters out spurious correlations due to non-stationarity of firing rates and allows exact statistical significance testing of observed spike correlations. Previous applications of this method required numerical bootstrapping of original spike trains, which involves extensive computations and is feasible only for small data sets. To overcome this limitation, we derived exact analytical solution for the distribution of cross-correlation values at each time lag under the null-hypothesis of independent spike trains. This analytical solution enabled us to rigorously assess significance of spike correlations in our large database. We then estimated the effective coupling between cortical layers and mapped it through the cortical depth. This analysis detected differential changes in inter-laminar interactions during attention, even so the firing rate modulation was similar across layers.

II-88. Population codes for topography in the zebrafish optic tectum

Lilach Avitan
Zac Pujuc
Ethan Scott
Geoffrey Goodhill

University of Queensland, Australia

L.AVITAN@UQ.EDU.AU
Z.PUJIC@UQ.EDU.AU
ETHAN.SCOTT@UQ.EDU.AU
G.GOODHILL@UQ.EDU.AU

The visual system has been an attractive target for studying neural coding. However, this has so far been mostly in the context of problems such as decoding edge orientation from the activity of populations of edge-selective neurons in V1. Surprisingly, there has been little quantitative investigation of population coding in the topographic representation of visual space, perhaps because topography is traditionally thought of as a place code rather than a population code. To address this we perform functional imaging of topographic representations in the zebrafish optic tectum, a model system which permits non-invasive imaging of neural activity. Stimuli placed in different positions on an LCD screen, covering different areas of the zebrafish visual field, are presented to zebrafish larvae while performing confocal calcium imaging of tectal neurons loaded with fluorescent calcium indicator (OGB-AM1). Using a Bayesian framework we decode the visual topographic information from a large population of tectal cells, to examine the extent to which the spatial information in the stimulus is preserved in the tectum, and the role this plays in decoding.

II-89. Robust estimation for neural state-space models

Lars Buesing¹
Jakob H Macke^{2,3}
Maneesh Sahani¹

LARS@GATSBY.UCL.AC.UK
JAKOB.MACKE@GMAIL.COM
MANEESH@GATSBY.UCL.AC.UK

¹University College, London

²Max Planck Institute Tübingen

³Bernstein Center Tübingen

Neurons within cortical populations are tightly coupled into collective dynamical systems that code and compute cooperatively. This dynamical coupling leads to firing variability which is structured across both neurons and time, and which can be described by statistical models where a latent low-dimensional ('state-space') stochastic process represents the common effect of network activity on the recorded cells. However, such state-space models can be challenging to fit to experimental data. The discrete nature of spiking leads to non-linear and non-Gaussian models, necessitating approximations during model estimation that may be computationally intensive or error-prone. Furthermore, the likelihood function—the quality of fit as a function of model parameters—may have multiple maxima, making it difficult to find the overall best model amongst many locally-optimal ones. We present an algorithm which improves the efficiency and robustness of estimation for statistical models in which a latent stochastic linear dynamical system (LDS) drives generalised-linear representations of individual cells. Our algorithm is based on an engineering approach called subspace identification (SSID). SSID was developed to estimate LDS models of Gaussian variables and works by identifying low-dimensional structure in the matrix of covariances between anisochronic measurements. It yields a unique and statistically consistent estimate at relatively little cost. We have extended SSID to the generalised-linear setting. The extended SSID learns a good model of neural population activity. On large simulated data sets with Poisson spike-counts, the algorithm recovers the correct parameters rapidly, without iteration or approximation. On multi-electrode cortical recordings it provides an effective initialisation for conventional maximum-likelihood estimation, avoiding poor local optima and substantially speeding convergence. Thus the new approach promises to render state-space methods with non-Gaussian observations far more practicable.

II-90. A design procedure for hierarchical neural control.

Gregory Wayne
Larry Abbott
Columbia University

GDW2104@COLUMBIA.EDU
LFA2103@COLUMBIA.EDU

The control of movement is the final common pathway for all behavior. It is frequently proposed, often conceptually rather than mechanistically, that the nervous system produces behavioral output by a hierarchy of modular circuits, where higher circuits perform more abstract roles in governing motor generation. The absence of an obvious way to provide instructive feedback to higher, more behaviorally distal circuits has thwarted progress in hierarchical neural control. We design such hierarchies by a bottom-up construction where each controller's objective can be parametrically varied by contextual inputs. A higher level modulates a lower levels context inputs and models the effect of doing so. We build the hierarchy one level at a time. At each new level, we follow a three-step process: 1) we train a forward model to predict the result of passing a motor command to the level below; 2) we design a cost function that specifies the task the given level must perform; 3) we use the forward model to find optimal motor plans and build a controller that memorizes the optimal plan to apply in each context. A higher-level controller's motor command changes the contextual information and cost function of the level below. Trained controllers issue commands without experimenting on forward models. We demonstrate our method by constructing a two-level neural network that uses realistic sensors to solve the challenging sensorimotor task of driving a simulated semi-truck in reverse around stochastically drifting obstacles. Only the lower-level controller can drive, while only the higher-level controller senses a high-dimensional depth map of the obstacles. Although our particular application is not biologically grounded, we argue this work provides two significant ideas for understanding biological motor control: forward models can simulate not only the body but also lower controllers, and control circuits can propagate objectives as well as motor commands.

II-91. Robust learning of low dimensional dynamics from large neural ensembles

David Pfau
Eftychios A. Pnevmatikakis
Liam Paninski
Columbia University

DBP2112@COLUMBIA.EDU
EFTYCHIOS@STAT.COLUMBIA.EDU
LIAM@STAT.COLUMBIA.EDU

Activity in large neural populations can often be modeled in terms of simpler low-dimensional dynamics. The advent of large-scale neural recording techniques has motivated the development of algorithms for dimensionality reduction and latent dynamics estimation from single-trial data. However, many previous algorithms are slow or suffer from local optima. In this work we present a novel, more robust approach to simultaneous dimensionality reduction and linear dynamical system estimation from single-trial population spiking data. Our approach is based on recent work in control theory and convex optimization and can be broken into two stages. In the first stage, an instantaneous firing rate is estimated directly from spike counts by use of convex nuclear norm penalized regression; this trades off between the likelihood of the data and the dimensionality of the latent dynamics. In the second stage, the firing rates are used to estimate the parameters of a linear dynamical system via standard regression-based approaches. The full model can be viewed as a low-dimensional latent linear system with point process outputs. Because this approach is based on a convex optimization, it is not prone to local optima (unlike previous Expectation-Maximization approaches). Explicit assumptions on the state noise (e.g., Gaussianity) are avoided. Finally, this approach can incorporate any output model with a smooth concave loglikelihood, including generalized linear models, maximum-entropy models, dichotomized Gaussian models, and other population spike train models. We believe that the generality and broad applicability of our method makes it suitable for a wide variety of large-scale neuroscience problems.

II-92. Beyond GLMs: a generative mixture modeling approach to neural system identification

Lucas Theis¹
 Dan Arnstein²
 André Maia Chagas¹
 Cornelius Schwarz²
 Matthias Bethge³

LUCAS@BETHGELAB.ORG
 DANARNSTEIN@GMAIL.COM
 ANDRE.CHAGAS@KLINIKUM.UNI-TUEBINGEN.DE
 CORNELIUS.SCHWARZ@UNI-TUEBINGEN.DE
 MBETHGE@TUEBINGEN.MPG.DE

¹Centre for Integrative Neuroscience

²University of Tübingen

³Max Planck Institute

One of the principle goals of sensory systems neuroscience is to characterize the relationship between external stimuli and neuronal responses. A popular choice for modeling the responses of neurons is the generalized linear model (GLM). However, due to its inherent linearity, choosing a set of nonlinear features is often crucial but can be difficult in practice if the stimulus dimensionality is high or if the stimulus-response dependencies are complex. We derived a more flexible neuron model which is able to automatically extract highly nonlinear stimulus-response relationships from data. We start out by representing intuitive and well understood distributions such as the spike-triggered and inter-spike interval distributions using nonparametric models. For instance, we use *mixtures of Gaussians* to represent spike-triggered distributions which allows for complex stimulus dependencies such as those of cells with multiple preferred stimuli. A simple application of Bayes' rule allows us to turn these distributions into a model of the neuron's response, which we dub spike-triggered mixture model (STM). The superior representational power of the STM can be demonstrated by fitting it to data generated by a GLM and vice versa. While the STM is able to reproduce the behavior of the GLM, the opposite is not the case. We also apply our model to single-cell recordings of primary afferents of the rat's whisker system and find quantitatively and qualitatively that it is able to better reproduce the cells' behavior than the GLM. In particular, we obtain much higher estimates of the cells' mutual information rates.

II-93. Low-rank connectivity induces firing rate fluctuations in a chaotic spiking model

Brian DePasquale
 Mark M Churchland
 Larry Abbott

BRIAN.DEPASQUALE@GMAIL.COM
 MC3502@COLUMBIA.EDU
 LFA2103@COLUMBIA.EDU

Columbia University

Balanced networks of spiking neurons with strong synapses (van Vreeswijk & Sompolinsky, 1996) robustly reproduce the irregular spiking activity of cortical neurons. However, on their own, such models fail to capture observed slower, low-dimensional firing-rate fluctuations (Litwin-Kumar & Doiron, 2012; Churchland & Abbott, 2012) that presumably represent the dominant neural signals of cortical computation and representation. We show that slow firing-rate dynamics can be generated in a network of 1000-5000 spiking neurons by adding a low-rank (LR) connectivity structure to a randomly connected network. The added rank-R connectivity matrix contains $R/2$ pairs of complex eigenvalues. The addition of low-rank connectivity preserves the irregular spiking generated by the random, full-rank connectivity while generating slow firing-rate fluctuations. These fluctuations, analyzed by PCA, are many times slower than the membrane and synaptic time constants. For small R , the firing rates fluctuate periodically, whereas for larger R , the fluctuations appear chaotic. The average power spectrum of the first 30 principal components of network activity show significant power over a continuous range of low frequencies not displayed in a random network, indicative of the slow-rate fluctuations apparent in the network activity. Our approach combines a full-rank, random connectivity matrix — critical for generating chaotic spiking dynamics — with low-rank connectivity that induces firing-rate fluctuations similar to those seen in chaotic firing-rate networks. A description of this spiking network in terms of continuous, rate-like variables can be derived from the low-rank

synaptic connectivity. In this model, iterative network training algorithms (Sussillo & Abbott, 2009) can be applied to harness spiking network activity for pattern generation and various other tasks.

II-94. Got a moment or two? Neural models and linear dimensionality reduction

Il M Park
Evan Archer
Nicholas Priebe
Jonathan W Pillow

University of Texas at Austin

MEMMING@AUSTIN.UTEXAS.EDU
EVANARCHER@GMAIL.COM
NICHOLAS@MAIL.UTEXAS.EDU
PILLOW@MAIL.UTEXAS.EDU

A popular approach for investigating the neural code is via dimensionality reduction (DR): identifying a low-dimensional subspace of stimuli that modulate a neuron's response. The two most popular DR methods for spike train response involve first and second moments of the spike-triggered stimulus distribution: the spike-triggered average (STA) and the eigenvectors of the spike-triggered covariance (STC). In many cases, these methods provide a set of filters which span the space to which a neuron is sensitive. However, their efficacy depends upon the choice of the stimulus distribution. It is well known that for radially symmetric stimuli, STA is a consistent estimator of the filter in the LNP model. Recently, Park and Pillow proposed an analogous model-based interpretation of both STA and STC analysis based on a quantity called the expected log-likelihood (ELL). Here, building upon the previous work, we present a novel model class—the generalized quadratic model (GQM)—which bridges a conceptual and methodological gap between moment-based dimensionality reduction on one hand and likelihood-based generative models on the other. The resulting theory generalizes spike-triggered covariance analysis to both analog and binary response data, and provides a framework enabling us to derive asymptotically-optimal moment-based estimators for a variety of non-Gaussian stimulus distributions. This extends prior work on the conditions of validity for moment-based estimators and associated dimensionality reduction techniques. The GQM is also a probabilistic model of neural responses, and as such generalizes several widely-used models including the LNP, the GLM, and the 2nd-order Volterra series. We apply these methods to simulated and real neural data from retina (spiking) and V1 (membrane potential).

II-95. Spike train entropy-rate estimation using hierarchical Dirichlet process priors

Karin Knudson
Jonathan W Pillow

University of Texas at Austin

KKNUDSON@MATH.UTEXAS.EDU
PILLOW@MAIL.UTEXAS.EDU

Entropy rate quantifies the average uncertainty over a new symbol given the history of a stochastic process. For spiking neurons, the entropy rate places an upper bound on the rate at which the spike train can convey stimulus information, and a large literature has focused on the problem of estimating entropy rate from recorded spike train data. Here we present a Bayes Least Squares entropy rate estimator for binary spike trains using Hierarchical Dirichlet Process (HDP) priors. Our estimator leverages the fact that the entropy rate of an ergodic Markov Chain with known transition probabilities can be calculated analytically, and many stochastic processes that are non-Markovian can still be well approximated by Markov processes of sufficient depth. Choosing an appropriate depth of Markov model presents challenges due to possibly long time dependencies and short data sequences: a deeper model can better account for long time-dependencies, but is more difficult to infer with limited data. Our approach mitigates this difficulty by using a hierarchical prior to share statistical power across Markov Chains of different depths. For binary data, the HDP [Teh et al, 2006] reduces to a hierarchy of Beta priors, where the prior probability over g , the probability of the next symbol given a long history, is a Beta distribution centered on

the probability of that symbol given a history one symbol shorter. Thus, the posterior over symbols given a long Markov depth is "smoothed" by the probability over symbols given a shorter depth. We compare our method to existing estimators [Lempel & Ziv 1976, Nemenmen et al 2001], and evaluate the results on both simulated and real neural spike train data. Our results show that tools from modern Bayesian nonparametric statistics hold great promise for extracting the structure of neural spike trains despite the challenges of limited data.

II-96. The neural ring: an algebraic tool for analyzing the intrinsic structure of neural codes

Nora Youngs
Alan Veliz-Cuba
Vladimir Itskov
Carina Curto

S-NYOUNGS1@MATH.UNL.EDU
AVELIZ-CUBA2@UNL.EDU
VLADIMIR.ITSKOV@UNL.EDU
CCURTO2@MATH.UNL.EDU

University of Nebraska - Lincoln

Much has been learned about the coding properties of individual neurons by investigating stimulus-response functions, which relate neural activity to external stimuli. This is perhaps best understood in the context of receptive fields, defined broadly to include orientation tuning curves and place fields. If neurons fire according to their receptive fields, it is believed, then the corresponding stimuli are being represented in the brain. This story becomes more complicated, however, when one considers that the brain does not have access to stimulus-response functions, such as receptive fields, and instead only 'sees' neural activity. In hippocampus, for example, the brain must constantly form representations of novel environments without access to the changing place fields. What can be inferred about the represented stimulus space from neural activity alone? On the surface, it seems that nothing can be learned without the aid of 'dictionaries' that lend meaning to neural activity. Nevertheless, if one assumes that the responses of a neural population arise from convex receptive fields, it turns out that a great deal can be learned about the underlying stimulus space, even if the receptive fields themselves are unknown. This idea was first explored in (Curto & Itskov, 2008), where we found that topological features of an animal's environment can be inferred from hippocampal place cell activity — without the aid of place fields. In this work, we generalize that analysis by introducing a novel algebraic object, the neural ring, that fully captures the combinatorial data intrinsic to a receptive field code. Assuming that neurons have convex (but otherwise unknown) receptive fields, the neural ring enables us to algorithmically extract the detailed structure of the arrangement of receptive fields, which in turn enables a wide array of stimulus space features to be 'read off' directly from the neural code.

II-97. Flexible probabilistic approach for measuring model learning independent of behavioral improvement.

Nathaniel Powell
Paul Schrater
University of Minnesota

POWEL180@UMN.EDU
SCHRATER@GMAIL.COM

Estimating learning rates can be difficult in many tasks, particularly estimating the amount subject's have learned about underlying processes that don't directly mediate their performance. Classically, evidence of learning on a task was assessed by a subject's improvement in either accuracy or reaction time (or both). However, in some tasks an accuracy measure may be inappropriate, because subjects may improve their accuracy without learning anything about the underlying process generating the data they are attempting to track. Indeed, they may in some cases succeed in spite of having learned less about the underlying model. In other cases, model learning is assumed to proceed according to some learning theory such as reinforcement learning or Bayesian update, and the subject's learning rate is assessed based on fit to the parameters of these learning models. However, if the actual learning process being assessed does not match the theory, these measurements will be faulty. A

more ideal approach would allow us to measure the rate of model learning independently of performance based measures or assumptions of a particular learning theory. Our approach combines two ideas. First we can express the start points and likely endpoints of learning in terms of parametric generative models that can predict subject's responses at the beginning and end of the learning process. Second, we introduce a set of latent variables that represent whether the subject used a given model for each trial, a kind of ownership probability. We show how learning rates can be directly expressed in terms of the sequence of these ownership probabilities.

II-98. Fast missing mass approximation for the partition function of stimulus driven Ising models

Robert Haslinger¹
Demba Ba¹
Ziv Williams^{2,3}
Gordon Pipa⁴

ROB.HASLINGER@GMAIL.COM
DEMBA@MIT.EDU
ZWILLIAMS@PARTNERS.ORG
GPIPA@UOS.DE

¹Massachusetts Institute of Technology

²Harvard Medical School

³MGH

⁴University of Osnabruck

Ising models are routinely used to quantify the second order, functional structure of neural populations. With some recent exceptions, they generally do not include the influence of time varying stimulus drive to the population. Inclusion of stimulus drive carries a heavy computational burden because the partition function becomes stimulus dependent and must be separately calculated for all unique stimuli observed. Naive summations and/or Monte Carlo techniques may require hours of computation time. Here we present an extremely fast, yet simply implemented, method capable of approximating the stimulus dependent partition function in minutes or less. Noting that the most probable spike patterns (which are few) occur in the training data, we sum partition function terms corresponding to those patterns explicitly. We then approximate the sum over the remainder of the patterns (which are improbable, but many) by casting it in terms of the stimulus varying missing mass (total stimulus dependent probability of all patterns not observed in the training data). We use a product of conditioned logistic regression models to approximate how the missing mass is modulated by the stimulus. This method has complexity of roughly $O(LN^2)$ where N is the number of neurons and L the data length, contrasting with the $O(L^2N)$ complexity of other methods. Using multiple unit recordings from rat hippocampus and macaque DLPFC we demonstrate our method can approximate the stimulus driven partition function as or more accurately than Monte Carlo methods but requiring 2 or more orders of magnitude less computation time. This advance allows stimuli to be easily included in Ising models making them suitable for studying population based stimulus encoding.

II-99. Diversity of timescales in network activity

Rishidev Chaudhuri¹
Alberto Bernacchia¹
Xiao-Jing Wang^{1,2}

RISHIDEV.CHAUDHURI@YALE.EDU
ALBERTO.BERNACCHIA@YALE.EDU
XJWANG@YALE.EDU

¹Yale University

²New York University

The architecture of a biological network crucially shapes the functions it can subserve, but the relationship between a network's structure and its dynamical behavior is complicated and remains poorly understood. To shed insight into this important question, we studied model networks endowed with a connectivity that falls off with distance, as is observed in the cerebral cortex of the mammalian brain. We examined conditions under which local connectivity leads to neural activity that exhibits a separation of timescales, so that different parts of the network

respond to inputs with disparate temporal dynamics (some very fast, others much slower). In a linear network, this problem is formulated in terms of eigenvectors of the connection matrix, which determine characteristic activity patterns. A separation of timescales can be robustly achieved if different eigenvectors are localized to different parts of the network. We developed a framework to predict localized eigenvectors for classes of one-dimensional networks. Notably, spatially local connectivity by itself is insufficient to separate response timescales via localized eigenvectors. However, localization of time scales can be realized by heterogeneity, wherein the connectivity profile varies across nodes. Furthermore, provided that the connectivity profile changes smoothly with position, similar timescales are contiguous within a network. We contrast this with the well-known example of localized patterns induced by disorder, in which timescales are randomly distributed across the network. Our theory is in excellent agreement with numerical results and generalizes to predict the shapes of only partially localized eigenvectors. Beyond neural dynamics, our framework is general and applicable to the broader field of complex networks.

II-100. Modeling inter-neuron inhibition with determinantal point processes

Jasper Snoek¹
 Ryan P Adams²
 Richard Zemel¹

JASPER@CS.TORONTO.EDU
 RPA@SEAS.HARVARD.EDU
 ZEMEL@CS.TORONTO.EDU

¹University of Toronto

²Harvard University

Point processes on the real line have proven to be appropriate models of single-neuron neural spiking data. Renewal processes in particular model well the refractory period of a neuron through an explicit holding time in the form of a hazard function. Using these in the context of a generalized linear model can provide a stimulus dependent instantaneous spike rate. These models, however, do not incorporate the complex interactions between pairs or sets of neurons. Neurons within a retinotopic map may, for example, inhibit the spiking of other neurons with overlapping receptive fields. This phenomenon can not be captured in the context of a point process on a single dimensional space. We propose to model inter-neuron inhibition using a determinantal point process (DPP). Specifically, we model intra-neuron stimulus induced spiking using a generalized linear model with a Poisson process output. The Poisson spike rate of each neuron is used to indicate a preference for spiking behavior, while pairwise inhibition is introduced to model competition between neurons. The output of the generalized linear model in our approach is analogous to a unary potential in a Markov random field while the DPP captures pairwise interaction. Although inhibitory, i.e. negative, pairwise potentials render the use of Markov random fields intractable in general, the DPP provides a more tractable and elegant model of pairwise inhibition. Given neural spiking data from a collection of neurons and corresponding stimuli, we learn a latent embedding of neurons such that nearby neurons in the latent space inhibit one another as enforced by a DPP over the covariance between latent embeddings. Not only does this overcome a modeling shortcoming of standard point processes applied to spiking data but it provides an interpretable model for studying the inhibitive and competitive properties of sets of neurons.

III-1. Frontal neurons enable retrieval of memories over widely varying temporal scales

Ziv Williams^{1,2}

ZWILLIAMS@PARTNERS.ORG

¹Harvard Medical School

²MGH

When recalling the name of a person we met earlier during the day or when recalling their name weeks later, we are drawing on broadly distinct memory stores within the brain. After a memory has been formed, for ex-

ample, it is principally held by short-term neural assembly dynamics and synaptic changes within areas such as the hippocampus and entorhinal cortex. As the memory consolidates, however, it becomes represented by progressively distinct neuronal processes and associative neocortical areas that are largely independent of hippocampal function. The ability to track and retrieve memories from these different storage resources is necessary because it allows us to draw on broadly distinct neuronal mechanisms and anatomical areas responsible for storing information over both short and long time scales. The neural mechanism that enables the memory system to appropriately access these alternate memory stores during retrieval, however, is largely unknown. Here, we find that Brodmann area 45 of the ventrolateral prefrontal cortex (VLPFC) plays an essential role in this process. We observe that certain neurons in the VLPFC of primates performing an associative recall task responded only when transitioning between retrieval of memories held in long- versus short-term memory storage. Moreover, focal reversible inactivation of the VLPFC markedly degraded recall performance during this transition. However, performance was not affected at any other time during recall or when recalling separately trained consolidated memories, even though identically trained, indicating that this area was not essential for retrieval of the associations themselves. These findings define a novel regulatory mechanism that enables the memory system to dynamically access short- versus long-term memory storage during retrieval, thus allowing learned items to be tracked and targeted for retrieval over widely varying temporal scales.

III-2. The neuronal input channel switched by attention reflects Routing by Coherence

Iris Grothe¹
 David Rotermund²
 Simon Neitzel²
 Sunita Mandon³
 Udo Ernst³
 Andreas Kreiter²
 Klaus Richard Pawelzik³

IRISGROTHE@GMAIL.COM
 DAVROT@NEURO.UNI-BREMEN.DE
 SIMONNEITZEL@GOOGLEMAIL.COM
 MANDON@BRAIN.UNI-BREMEN.DE
 UDO@NEURO.UNI-BREMEN.DE
 KREITER@BRAIN.UNI-BREMEN.DE
 PAWELZIK@NEURO.UNI-BREMEN.DE

¹Strüngmann Institute

²Center for Cognitive Sciences

³University of Bremen

In natural environments neurons with large receptive fields incessantly receive inputs from different sources most of which are behaviourally irrelevant. Attention is believed to control the inputs in a way that signals from relevant sources can be selectively processed. The available evidence for such signal routing in the visual cortex, however, is mostly indirect. In particular, it was not yet shown, that attention de facto switches among specific channels. A detailed characterization of the channels' dynamical properties would furthermore provide important constraints for the neuronal mechanisms underlying attentional signal selection in cortex, which are not known. Here we establish a new experimental paradigm for investigating the channels' transmission properties for signals from different locations within the receptive field of area V4 neurons. We superimposed behaviourally irrelevant broad band contrast modulations on two visual objects during attention dependent shape tracking. We used a suitably normalized spectral coherence measure (NSC) to simultaneously characterize the transmission of the superimposed components towards local field potentials (LFPs) in areas V1 or V4. Also, we identify intrinsic activity components using the NSC between V1 and V4 LFPs. We found that attention gates the channels for visual signals (VS) towards V4 which are band limited at ~25 Hz. Contributions originating from the non-attended object are strongly suppressed. While this gating is absent for the transmission from VS to V1, it is also seen in the V1-V4 connection. Here, additional coherence in the gamma band (~60Hz) is observed, which co-varies with task conditions and attentional demand. These results are not only consistent with the hypothesis that gamma subserves gating by attention, but can be reproduced in detail by a minimal model implementing the idea of Routing By Coherence.

III-3. Optimal speed estimation in natural image movies

Johannes Burge
Wilson Geisler
University of Texas at Austin

JBURGE@MAIL.CPS.UTEXAS.EDU
GEISLER@PSY.UTEXAS.EDU

The neural computations underlying selective perceptual invariance possible are enormously complex. Many studies of neural encoding-decoding assume neurons with invariant tuning functions. Except for neural noise, these neurons give identical responses to all stimuli having the same value of the relevant dimension. This cannot occur under natural viewing conditions. Variation in natural signals along irrelevant dimensions inevitably cause the responses to vary. Here, we use a novel task-specific encoding-decoding framework that specifies how to encode task-relevant information and process it to construct neurons that are largely invariant to irrelevant natural image variation. We use the framework to estimate retinal image speed from photoreceptor responses to natural image movies. The space-time receptive fields (RFs) that optimally encode information relevant for estimating speed are direction selective but, interestingly, they are not speed-tuned. Appropriate non-linear combination of the RF responses yields a new population of neurons that are speed tuned and are (largely) invariant to irrelevant stimulus dimensions. These neurons represent the log-likelihood of speed and have tuning curves that are approximately log-Gaussian (LL) in shape. MAP decoding yields unbiased speed estimates over a wide range. The optimal space-time RFs and the speed-tuned LL neurons share many properties with neurons in cortex. A majority of motion sensitive neurons in V1 and MT are direction, but not speed selective. Roughly 25% of V1 and MT neurons are speed tuned (Priebe, Lisberger, Movshon, 2006), with tuning curves that are log-Gaussian in shape (Nover, Anderson, DeAngelis, 2005). Critically, the optimal space-time RFs and the optimal speed tuning curves in our analysis were not arbitrarily chosen to match the properties of neurophysiological RFs. Rather, they emerge from a task-specific analysis of natural signals. We find it remarkable that an ideal-observer analysis, with appropriate biological constraints and zero free parameters, predicts the dominant neurophysiological features of speed processing.

III-4. Dynamic calibration of the influence of priors and sensory input for perceptual estimation

Kamesh Krishnamurthy
Matthew Nassar
Shilpa Sarode
Joshua Gold
University of Pennsylvania

KAMESH@MAIL.MED.UPENN.EDU
MATTNASSAR@GMAIL.COM
MOONLIGHTSHINE90@GMAIL.COM
JIGOLD@MAIL.MED.UPENN.EDU

Perception of an ambiguous sensory stimulus often depends on both sensory information and prior expectations about the presence or identity of the stimulus. Numerous studies support the idea that the brain follows Bayesian inference for a broad range of perceptual problems by combining each source of information in proportion to its reliability. However, this idea has typically been tested under conditions in which the reliability of prior expectations is relatively stable, which is often not the case in dynamic, real-world environments. More realistically, the underlying statistics of the environment can undergo changes that alter the reliability of prior expectations. Here we tested how abrupt, unexpected changes in the reliability of expectations about a stimulus affected the extent to which those expectations influence perception. We used a novel sound-localization task to measure the influence of dynamic prior expectations on perception. We tested how dynamic changes in these expectations affected localization reports. The location of a virtual sound source was varied randomly from trial-to-trial about a mean value, and on certain, randomly chosen trials, the location of the mean itself changed abruptly. On each trial, the subjects indicated both their prior expectation about the location of the sound before listening to the sound and the perceived location of the sound afterwards. We found that: 1) following a change-point in the virtual location of the sound source, when the prior was least reliable, the prior had the weakest influence on perceived location, and 2) on subsequent trials, both the reliability of the prior and its influence on perceived location increased steadily.

These effects are consistent with an ideal-observer model describing the relative influence of priors and sensory evidence on perception in this environment. The results indicate that perception reflects a process of Bayesian inference that undergoes ongoing regulation in a changing environment.

III-5. Constraining a Bayesian model of orientation perception with efficient coding

Xuexin Wei
Alan Stocker

WEIXXPKU@GMAIL.COM
ASTOCKER@SAS.UPENN.EDU

University of Pennsylvania

A common challenge for Bayesian models of perceptual behavior is the fact that the two fundamental components of a Bayesian model, the prior distribution and the likelihood function, are formally unconstrained. Here we argue that a neural system that emulates Bayesian inference naturally imposes constraints by the way sensory information is represented in populations of neurons. More specifically, we show how an efficient coding principle can constrain both the likelihood and the prior based on the underlying stimulus distribution. We apply this idea to the perception of visual orientation and formulate an encoding-decoding model based on a prior distribution that reflects the statistics of visual orientations in natural scenes. At the behavioral level, our model predicts that perceived orientations are biased away from the cardinal orientations where the prior distribution peaks. Such biases are seemingly at odds with the traditional Bayesian view that the prior always biases a percept towards the prior peaks. Yet they are in perfect agreement with recent studies that report perceptual biases toward the oblique orientations. Our model also correctly predicts the reported relative biases toward the cardinal orientation when comparing the perceived orientation of a stimulus with low versus a stimulus with high external noise. The model is able to account for both types of biases because it assumes an efficient neural representation that typically generates asymmetric likelihoods with heavier tails away from the prior peaks. At the neurophysiological level, the proposed efficient coding principle predicts neural tuning characteristics that match many aspects of the known orientation tuning properties of neurons in primary visual cortex. Our results suggest that efficient coding provides a promising constraint for Bayesian models of perceptual inference, and might explain perceptual behaviors that are otherwise difficult to reconcile with traditional Bayesian approaches.

III-6. Some work and some play: a normative, microscopic approach to allocating time between work & leisure

Ritwik Niyogi¹
Yannick Breton²
Rebecca Solomon³
Kent Conover²
Peter Shizgal²
Peter Dayan¹

RITWIK.NIYOGI@GATSBY.UCL.AC.UK
YBRETON@LIVE.CONCORDIA.CA
RB.SOLOMON@HOTMAIL.COM
KENT.CONOVER@GMAIL.COM
PETER.SHIZGAL@CONCORDIA.CA
DAYAN@GATSBY.UCL.AC.UK

¹University College, London
²CSBN, Concordia University
³Concordia University

When suitably freed, subjects generally elect to divide their time between work and leisure, rather than consuming just one. Existing accounts of this division have largely characterised behaviour at a macroscopic level of detail, ignoring the wealth of regularities present in the microscopically fine temporal patterning of responding. Recently, as a substantial extension to Niv et al (2007)'s work on the vigour of responding, we posited a general optimal control framework in which subjects make stochastic, approximately-optimizing microscopic choices: whether and when to work or engage in leisure, and how long to do so for (Niyogi et al, CoSyNe 2012). They can gain

benefits from both work and leisure, but pay an automatic opportunity cost for the time they allocate. This framework spans computational questions, such as the optimality of behaviour; algorithmic issues such as the nature of the response policy; and provides a detailed tie to the neural mechanisms of free operant choice. Here, (a) we show that the critical construct for fitting behaviour is the function quantifying the benefit of leisure as a function of leisure duration; (b) we reverse engineer the empirical benefit of leisure functions in a number of rats, demonstrating that they are non-linear; and (c) we explore the formal relationships between this account and the traditional macroscopic accounts.

III-7. Optimally fuzzy memory

Karthik Shankar

SHANKARK@BU.EDU

Center for Memory and Brain

Any system with the ability to learn patterns from a temporally structured sequence of stimuli and predict the subsequent stimulus at each moment, should have a memory representing the stimuli from the recent past. If the external environment generating the sequence has a fixed scale, the memory can simply be a shift register—a moving window extending into the past, like traditional models of short term memory. However, this is inappropriate when the statistics to be learned involves long-range correlations as found in many physical environments. We argue in favor of a scale-free fuzzy memory which optimally sacrifices the accuracy of information representation in favor of capacity to represent long time scales. Such a memory system can be easily implemented as a neural network with bands of excitatory and inhibitory connections from a population of leaky integrator cells. Given the ubiquitous long-range correlated fluctuations in the natural world, it would have been evolutionarily adaptive for humans and animals to have evolved with such a memory system. This is indeed consistent with a variety of behavioral findings from timing and memory tasks. The neurobiological underpinnings of this memory system can be ascribed to the medial temporal lobe. Though it has been conventionally accepted that the hippocampus and the medial temporal lobe hold a spatio-temporal map of our current location in space and time, there is no unified model explaining the origins of both time-cells and place-cells. Here we show that the network model of the fuzzy memory system provides such a unified framework. Moreover, to emphasize the importance of the fuzzy memory in artificial intelligence, we show that the fuzzy memory outperforms a shift register in learning and forecasting some simple time series data.

III-8. Sparse Bayesian inference and experimental design for synaptic weights and locations

Ari Pakman

Carl Smith

Liam Paninski

Columbia University

ARIPAKMAN@GMAIL.COM

CAS2207@COLUMBIA.EDU

LIAM@STAT.COLUMBIA.EDU

A detailed understanding of the organization of local neural circuits requires determining not only the connectivity pattern in a neural population, but also the exact location and strength of synaptic interactions in the dendritic tree. In a previous work, we showed how to approach this problem by combining the ability to stimulate individual presynaptic neurons with simultaneous imaging of postsynaptic neurons at subcellular resolution. This work extends our previous results in two directions. On the one hand we revisit the inference method used to extract the locations and strengths of the synaptic weights from the observed data. Instead of the maximum a posteriori (MAP) solution of a state-space model with an L1 prior (the Lasso model), in this work we also obtain confidence intervals by adopting a fully Bayesian approach. In particular, we compare the results of several popular sparsity-inducing priors for the synaptic weights: the Bayesian Lasso, the Horseshoe and the Spike-and Slab. Particular emphasis is placed on the constraint imposed by Dale's law, which states that the synaptic weights

have a definite sign, thus leading to truncated probability distributions. Equipped with the full posterior distribution of the synaptic weights, our second contribution explores optimal experimental design. We extend the type of voltage measurements from localized observations to linear combinations of voltages across several locations with random coefficients. This setting corresponds to a "compressed sensing" sampling scheme, which yields an impressive reduction in the number of measurements required to infer the synaptic weights. In particular, we show how to choose the correlation among the random coefficients to offset the correlation between successive measurements imposed by the neuron dynamics. We illustrate our results on simulated measurements in toy and real neurons.

III-9. Inferring functional connectivity with priors on network topology

Scott Linderman¹
Ryan P Adams^{2,1}

SLINDERMAN@SEAS.HARVARD.EDU
RPA@SEAS.HARVARD.EDU

¹Harvard University

²School of Engineering and Applied Sciences

Neural computation is manifested in functional relationships between neurons. Reverse engineering this computation can be approached from the bottom up, looking for patterns in neural recordings, and from the top down, hypothesizing models of functional connectivity that could explain observed phenomena. We construct a Bayesian approach to modeling the latent functional network, using data from observed spike trains and representing top-down hypotheses via random graph models. We build on the GLM framework, extending the model to explicitly incorporate prior distributions on the topology and weight of the functional connections, as well as a saturating nonlinear link function. Since many interesting priors lack convexity, we depart from the traditional MLE and MAP approaches, opting instead for a Gibbs sampling algorithm with an efficient GPU implementation. Our method is demonstrated on a synthetic spike train generated from an integrate-and-fire network with clusters of highly-connected excitatory neurons suppressing one another through an inhibitory pathway, thereby forming stable attractor states. A prior for a densely connected functional network, as implicitly assumed in the GLM, claims direct inhibition between the two clusters due to their negative correlation, whereas a prior biased toward blocks of functionally similar neurons recovers the inhibitory pathway. This is but a proof of concept for a very general Bayesian framework. We can build priors for spatially proximal connectivity or time-varying weights, models for hierarchical or distributed representations of information among neurons, or models where neurons represent states in complex dynamical systems. More importantly, we have an algorithm for efficiently performing inference given spike train recordings, and for comparing models in a principled framework.

III-10. Recurrent generalized linear models with correlated Poisson observations

Marius Pachitariu
Lars Buesing
Maneesh Sahani
University College, London

MARIUS@GATSBY.UCL.AC.UK
LARS@GATSBY.UCL.AC.UK
MANEESH@GATSBY.UCL.AC.UK

We introduce the Recurrent Generalized Linear Model (R-GLM), an extension of GLMs based on a compact representation of the spiking history through a linear recurrent neural network. R-GLMs match the predictive likelihood of Linear Dynamical Systems (LDS) with linear-Gaussian observations. We also address a disadvantage of GLMs, including the R-GLM, that they cannot model instantaneous correlations. The LDS however allows for extra correlated variability through the new innovation in the latent space. To improve GLMs we introduce a class of correlated output distributions which can be used with any type of multivariate data - binary, counts or continuous. The correlated Bernoulli distribution matches the predictive likelihood of Ising models for static binarized

spike data. The correlated Poisson distribution offers significant improvements in predictive likelihood for GLMs and R-GLMs. We evaluate the performance of the models on a dataset recorded from a Utah array implanted into motor areas of a macaque monkey during a delayed reaching task. We report that the R-GLM consistently finds long timescales (of up to several seconds) of correlated activity similar to those found by LDS and longer than the timescales learnt by standard GLMs (up to 400 ms). Like all GLMs, the proposed model can be used with any link function and any output distribution. This is unlike models based on LDS which require careful approximations to be trained with Poisson outputs.

III-11. Distinct coherent ensembles reflect working memory processes in primate PFC

David Markowitz
Bijan Pesaran
New York University

DAVID.MARKOWITZ@NYU.EDU
BP31@NYU.EDU

Coherent ensemble spiking activity has been linked with attention, coordination and decision-making processes. Computational modeling studies have suggested that coherent activity may be an important property of working memory (WM) networks, as well. Although recent experimental work has identified beta (11-30 Hz) frequency LFP oscillations as a prominent signature of WM in prefrontal cortex (PFC), the origin of these oscillations and their relationship to spiking activity in PFC during WM remain unknown. To address this question, we chronically implanted two monkeys with a 32-channel movable electrode array microdrive over PFC and recorded from ensembles of isolated units at each depth of cortex while the animals performed memory- and visually-guided delayed saccades. We find that distinct cellular response classes are revealed when memory-tuned units (391/706 cells) are grouped by recording depth (above/below 1 mm) and the presence of significant beta frequency spike-field coherency. Superficial coherent neurons (n=74) respond in a memory-specific manner, with persistently elevated rates during memory delays and baseline firing rates during visual delays. Deep coherent neurons (n=75) do not respond in a memory-specific manner, and exhibit the same persistently elevated rates during both memory and visual delays. These findings suggest that populations of superficial beta-coherent neurons specifically support memory maintenance through cross-columnar interactions, consistent with theoretical predictions. To determine if such interactions lead to firing rate covariation within coherent ensembles, we estimated noise correlations during the memory delay for n=77 superficial coherent pairs separated by 1.5 – 8 mm, and find that this is not significantly different from zero (0.01 *pm* 0.01 s.e.m.). Based on these findings, we propose a computational framework in which working memory processes are supported by distinct ensembles of coherent neurons at superficial depths (< 1 mm) within the prefrontal cortical sheet.

III-12. Dimensionality, dynamics, and correlations in the motor cortical substrate for reaching

Peiran Gao¹
Eric Trautmann¹
Byron Yu²
Gopal Santhanam
Stephen Ryu¹
Krishna Shenoy¹
Surya Ganguli¹

PRGAO@STANFORD.EDU
ETRAUT@STANFORD.EDU
BYRONYU@CMU.EDU
GOPAL@NERUR.COM
SEOULMANMD@GMAIL.COM
SHENOY@STANFORD.EDU
SGANGULI@STANFORD.EDU

¹Stanford University

²Carnegie Mellon University

How coherent motor behavior emerges from large populations of neurons constitutes a fundamental question in

neuroscience. Dimensionality reduction and correlation analysis are often used to address this question. Interestingly, neuronal data's effective dimensionality is often far smaller than the number of recorded neurons [Yu 2009], and recent work has shown that low-dimensional neuronal dynamics exhibits rotational structure [Churchland 2012]. Moreover, despite weak pairwise correlations, one can accurately predict a neuron's spiking from only $O(100)$ other neurons' activities [Truccolo 2005]. These various low dimensional views leave open several important questions: What determines the effective dimensionality of neuronal activity? What creates rotational structure? And how do spatiotemporal correlations enable spiking prediction? By analyzing 109 simultaneously recorded PMd neurons from monkeys performing an eight direction delayed reach task [Yu 2007], we find a simple view that answers these questions. Using Gaussian mixture models fitted to trial averaged activity, we find that most neurons exhibit a sharp, monophasic activity peak during movement [but see Churchland...2007]. Each peak's timing, but not amplitude, is largely independent of reach angle. This sparse wave of neural activity comprises a nonlinear manifold, which does not lie within a single low dimensional linear space, and evolves through different dimensions over time. We show empirically and analytically that: (a) the dimensionality of the smallest linear subspace containing this manifold is near an upper bound estimated by task complexity and network correlation time; (b) when projected onto a lower dimension, this manifold exhibits rotational structure; (c) inter-trial variability concentrates along the manifold; and (d) sparsity of activity underlies our ability to predict single neuron behavior from the ensemble. This work unifies and explains, through a single high-dimensional perspective, disparate phenomena previously viewed through different low-dimensional lenses, and suggests new analysis methods for finding rich neural structures that may be missed by time-independent dimensionality reduction.

III-13. Optimal neural tuning for arbitrary stimulus priors with Gaussian input noise

Zhuo Wang¹
 Alan Stocker¹
 Haim Sompolinsky²
 Daniel Lee¹

WANGZHUO@SAS.UPENN.EDU
 ASTOCKER@SAS.UPENN.EDU
 HAIM@FIZ.HUJI.AC.IL
 DDLEE@SEAS.UPENN.EDU

¹University of Pennsylvania

²The Hebrew University

Sensory neurons represent input information via their tuning curves. The optimal tuning curve can be determined by many factors, e.g. optimal criteria, stimulus prior distribution, output noise model and encoding time (Bethge et al. 2002, Ganguli & Simoncelli 2010). The output noise, i.e. the variability in spike-count over time, has been well investigated. Previous theoretical work (Nadal & Parga 1994) has provided some insights into the role input noise can play in optimal encoding, where input noise is considered any noise that affects the input of the neuron. Here we argue that by analyzing the input - output noise pair, we can have a better understanding of the neurons in deeper layers of a neural network since those neurons can only access the noisy output received from upstream neurons. In this work, we assume a Poisson output noise model and generalize our long encoding time limit result (Wang, Stocker & Lee 2012) to the case with the presence of an additive Gaussian input noise. Having the new Fisher information computed, various objective functions, e.g. mutual information and mean asymptotic Lp loss, can be calculated via Cramer-Rao lower bound and Mutual Information Bound (Brunel & Nadal 1998). We analytically solved the optimization problem and derived the optimal condition for the tuning curve. Our work answers the question how neurons should optimally combat with the input noise. Compare to other biophysiological approach (Chance et al. 2002, Lundstrom et al. 2009), our result directly links the phenomena - neuron's adaptation to synaptic noise - to the underlying motivation i.e. optimality criteria. Furthermore, our model can be used to better understand the loss function a neuron is trying to optimize when encoding information. By applying a small input perturbation and observe the change of the tuning curve, the actual optimality criterion can be identified.

III-14. From metacognition to statistics: relating confidence reports across species

Balazs Hangya
Joshua Sanders
Adam Kepecs

BALAZS.CSHL@GMAIL.COM
JOSHUA21@GMAIL.COM
KEPECS@CSHL.EDU

Cold Spring Harbor Laboratory

Confidence judgments, self-assessment about strength of beliefs, are considered a central example of metacognition. While *prima facie* self-reports are the only way to access subjective confidence, it has been proposed that overt behavioral measures can yield measures of confidence in appropriately trained animals. A theoretical basis is needed to relate confidence reports in animals to the human concept of subjective confidence. Here we use the notion of confidence used in statistical hypothesis testing to link explicit self-reported confidence in humans with implicit reports in animal models. We present a confidence-reporting task, in which human and non-human animals make auditory decisions and report their confidence verbally in humans and implicitly in both humans and rats. Subjects respond by making a directional movement indicating their choice, and waiting until they are rewarded following a random delay, or until they give up. For error trials subjects receive feedback after aborting the wait, and we introduce catch trials, for which feedback is omitted. Humans are also prompted to rate their certainty that they had initially made the correct decision on a 1-5 scale. Statistical hypothesis testing was used to formalize properties of confidence measures in two-choice decisions and make predictions that were tested with psychophysical data. We found that both self-reported confidence and the duration subjects were willing to wait for an uncertain outcome followed the predictions of the theory. Moreover, waiting time was strongly correlated with explicit confidence ratings in humans. Next, we tested rats in the same auditory decision paradigm and the duration they were willing to wait for uncertain rewards also followed the theoretically predicted patterns. By relating verbal and implicit confidence reports that can be performed by both rodents and humans, our results provide a basis for studying the neural mechanisms underlying confidence judgments.

III-15. Eyetracking and pupillometry reflect dissociable measures of latent decision processes

James Cavanagh
Thomas Wiecki
Angad Kochar
Michael Frank

JIM.F.CAV@GMAIL.COM
THOMAS_WIECKI@BROWN.EDU
ANGAD_KOCHAR@BROWN.EDU
MICHAEL_FRANK@BROWN.EDU

Brown University

Soon, webcams will be able to track eye movements and pupil sizes to determine what advertisements a user finds appealing, and which are bothersome and ineffective. There is tremendous value in understanding how observable behaviors reflect attitudes and choices, yet these relationships are likely to be latent and unobservable. The Drift Diffusion Model (DDM) provides an algorithmic account of how evidence accumulation and a speed-accuracy tradeoff contribute to decision making through separate latent parameters of drift rate and decision threshold, respectively. While the DDM is often used to describe decision making in noisy sensory environments, it is equally valuable for understanding decision making during value-based choices. In this study, participants completed a reinforcement learning task with a forced choice transfer phase at the end. We used a novel hierarchical Bayesian estimation approach to assess trial-to-trial influence of physiological signals on DDM parameters. To probe the influence of different types of valuation, we investigated if choices between conflicting appetitive and aversive options differed from each other. Posterior estimates indicated that increased dwell time predicted increased drift rates toward the fixated option for both appetitive and aversive choices, with no difference between these conditions. Given the contributions of the salience network (anterior cingulate, insula, locus coeruleus) to pupil dilation, we hypothesized that this measure may index adjustments in decision threshold. Posterior estimates indicated that increased pupil dilation predicted a lower decision threshold specifically during choices between two aversive

options, where larger pupil size predicted faster rejection of an undesirable option, consistent with pupillary marker of an urgency signal when neither choice option is desirable. These findings suggest that combined eyetracking and pupillometry measures can provide an indicator of decision quality that may be useful to an observer.

III-16. Deviations from the matching law reflect reward integration over multiple timescales

Kiyohito Iigaya¹
Leo P. Sugrue²
Greg S. Corrado²
Yonatan Loewenstein³
William T. Newsome²
Stefano Fusi¹

KI2151@COLUMBIA.EDU
LEO@MONKEYBIZ.STANFORD.EDU
GREG.CORRADO@GMAIL.COM
YONATAN@HUJI.AC.IL
BNEWSOME@STANFORD.EDU
SF2237@COLUMBIA.EDU

¹Columbia University

²Stanford University

³Hebrew University

Animals can easily adapt to changing environments. In the case of foraging behavior, animals estimate the relative abundance of resources at different locations and then allocate time to each foraging site in proportion to its relative resource abundance (matching behavior). More generally, the matching law states that the fraction of choices made to any option should match the fraction of total reward earned from the option. Deviations from this behavior are often observed in the form of distortions of the curve expressing the fraction of choices vs fractional income. Undermatching occurs when this curve has a slope smaller than 1. We propose that undermatching is not a failure of the mechanisms underlying the matching law, but instead reflects the ability to integrate information over multiple timescales. This new interpretation was suggested by our study of a decision making model similar to the one described in Fusi et al.(2007). The model assumes that synapses integrate information over multiple timescales, as required to solve general memory problems Fusi et al.(2005). Each synapse has a cascade of states with different levels of plasticity, characterized by different learning rates. The distribution of the synaptic states determines an effective learning rate that changes over time. The model predicts that initially learning is dominated by plastic components, but then less plastic states are progressively populated, introducing a bias in the decision that depends on the long-term statistics. This bias tilts the matching curve generating undermatching. We tested this prediction by analyzing the behavioral data collected by Sugrue et al. (2004). We found that 1) the learning rate decreases with experience 2) undermatching is more prominent when learning is slow 3) despite the deviation from the matching law, the performance improves with experience because of the reduction in the fluctuations of reward probability estimates.

III-17. Changing subjects' fallacies by changing their beliefs about the dynamics of their environment

Friederike Schuur
Brian P Tam
Laurence T Maloney
New York University

F.SCHUUR@NYU.EDU
BPT218@NYU.EDU
LTM1@NYU.EDU

In perception, action, and decision-making, subjects exhibit automatic, apparently suboptimal sequential effects (SQE): they respond more rapidly to a stimulus if it reinforces local patterns in stimulus history. By manipulating beliefs about environmental dynamics, we investigate whether SQE are driven by mechanisms critical for adapting to a changing world. Subjects completed a version of 'Whack-a-Mole'. Sesame Street's Elmo appeared either to the right or left of fixation. Participants were instructed to press a spatially congruent button. In the 1st and 3rd

session, repetitions and alternations in stimulus-location were equally likely. In Session 2 (training), $p(\text{rep.})$ was repeatedly sampled from a Beta-distribution with $p(\text{rep.}) \sim B(6, 12)$ for gambler training (alternations more likely) and $p(\text{rep.}) \sim B(6, 12)$ for hot-hand training (repetitions more likely). Resampling occurred at $p(\text{resampl.}) = 0.18$ and was signaled to subjects to allow them to learn this change-rate. Prior to training, participants were naturally born gamblers: they expected more alternations than repetitions. During training and following each change-signal, participants in the hot-hand group responded increasingly faster on repetition and slower on alternation trials. These expectations carried over to Session 3. We found clear tendencies to expect repetitions rather than alternations in hot-hand trained participants while the gamblers continued to expect alternations rather than repetitions. Modeling reaction times revealed that SQE were driven by combining current evidence, down-weighted according to estimated $p(\text{resampl.})$, with beliefs about the state the world is going to be in after change. In particular, estimated change-rates and beliefs about the most likely state of the world after change (hot-hand vs. gambler) produced distinct SQE. Note that in Session 3, SQE were suboptimal. But while suboptimal, SQE can be modified through training. As such, we are the first to show that training can alter SQE.

III-18. Past failures bias human decisions

Arman Abrahamyan¹
Justin Gardner²

ARMAN@BRAIN.RIKEN.JP
JUSTIN@BRAIN.RIKEN.JP

¹RIKEN Brain Science Institute
²Gardner Research Unit

Past decision can bias present decision even if it is not the 'right' thing to do. For example, when deciding whether a visual stimulus was presented on the left or on the right, mice can exhibit a tendency to switch sides (left to right or vice versa) regardless of which side the stimulus was presented on (Busse et al., 2011, J Neurosci). Here we demonstrate that some people, including experienced observers, exhibit similar suboptimal strategies. Investigating these suboptimal strategies is important both because it gives insight into how humans make sensory decisions and because these strategies bias psychophysical measurements and thus suggest that many reports of psychophysical sensitivity may have been widely underestimated. Similar to the task performed by mice, in our experiment, humans reported if a weak or strong visual stimulus was presented on the left or right side of fixation. Each response was followed by auditory feedback to indicate correct and incorrect decisions. Using a probabilistic choice model, we found that four out of seven subjects demonstrated a tendency to switch sides (i. e., from left to right or vice versa) when they failed on the previous trial. Successful responses on the previous trial had little or no effect on subjects' decisions. We conclude that even experienced subjects employed a suboptimal strategy in which past failures influence current choices. While classical psychophysics has developed methodologies that were thought to be immune to subjects' biases (Green & Swets, 1966), our results suggest that accurate measurements of psychophysical thresholds require new approaches to modeling subjects' decisions. These suboptimal strategies may show systematic differences between species. For mice, this suboptimal strategy of switching could reflect a more global strategy of systematic foraging for randomly located food, while humans may prefer to change answers when the previous answer was wrong.

III-19. The effects of prior, rewards and sensory evidence in perceptual decision making

Yanping Huang
Rajesh Rao
University of Washington

HUANGYP@CS.WASHINGTON.EDU
RAO@CS.WASHINGTON.EDU

Optimal decision-making under uncertainty involves sophisticated combinations of noisy sensory input, prior knowledge, and reward information. The integration-to-bound model has successfully explained many behav-

ioral and electro-physiological experiments that examine how sensory evidence is accumulated to form perceptual decisions. However, two contradictory ways of incorporating prior knowledge have been proposed in the integration-to-bound framework. The first posits an additive static offset to a decision variable, while the second model includes a time-increasing bias signal in the decision variable. Each model only accounts for experimental data in some but not all versions of motion discrimination tasks involving biased choices. Here we use the framework of partially observable Markov decision processes (POMDPs) to explore how sensory input, prior probability and rewards are combined for decision-making. We show how the optimal choice can be selected based on the expected value of the sum of total future rewards associated with each choice. Since the reward of different choices is uncertain in a stochastic environment, we show how prior knowledge and sensory evidence are combined in a Bayesian way to compute the posterior distribution of the unknown environmental state, which is in turn used to compute the expected reward of each choice. We show that predictions from our proposed model require very few free parameters to fit a large set of behavioral data. Model predictions of effects of rewards on the speed and accuracy trade-off are consistent with experimental data reported by Palmer et al and Hanks et al. We also demonstrate how asymmetric rewards can induce biased decisions, as observed in Rorie et al's experiments. While the integration-to-bound model assumes different effects of prior probability to explain data in the fixed time and reaction time versions of motion discrimination tasks, our model predicts behavioral data from both tasks within a single unified framework.

III-20. The effect of novelty-based exploration on reinforcement learning in humans.

Audrey Houillon^{1,2}
 Robert Lorenz³
 Andreas Heinz³
 Jürgen Gallinat³
 Klaus Obermayer⁴

AUDREY.HOULLON@TU-BERLIN.DE
 ROBERT.LORENZ@CHARITE.DE
 ANDREAS.HEINZ@CHARITE.DE
 JUERGEN.GALLINAT@CHARITE.DE
 KLAUS.OBERMAYER@MAILBOX.TU-BERLIN.DE

¹BCCN

²TU Berlin

³Charite

⁴Technische Universität Berlin

Recent research suggests that novelty has an influence on reward-related decision-making. Here, we showed that novel stimuli presented from a pre-familiarized category could accelerate or decelerate learning of the most rewarding category, depending on whether novel stimuli were presented in the best or worst rewarding category. The extent of this influence depended on the individual trait of novelty seeking. Subjects' choices were quantified in reinforcement learning models. We introduced a bias parameter to model exploration toward novel stimuli and characterize individual variation in novelty response. The theoretical framework further allowed us to test different assumptions, concerning the motivational value of novelty. On the one hand SN/VTA activation by novelty has raised the possibility that novelty might have intrinsic rewarding properties. On the other hand specific SN/VTA activations to novelty alone suggest a second, reward-independent mechanism, that favors exploration towards the novel cue. Based on these findings, we proposed that novelty per se can act as an explorative bias, but can also act as a bonus for rewards when these are explicitly attended. The best-fitting model combined both novelty components. The model's bias parameter also showed a significant correlation with the independent novelty-seeking trait. Preliminary fMRI analysis showed the strongest signal change in the condition where reward and novelty were presented together, but only in low action probability trials. These results are in line with previous research, which showed that striatal reward processing could be boosted by novelty, but only in uncertain states during explorative behavior. This effect was furthermore observed in the midbrain, cingulate gyrus and PFC. Altogether, we have not only shown that novelty by itself enhances behavioral and neural responses underlying reward processing, but also that novelty has a direct influence on reward-dependent learning processes.

III-21. How does the brain compute value? A rational model for preference formation.

Paul Schrater
Nisheeth Srivastava
University of Minnesota

SCHRATER@GMAIL.COM
NISHEETHS@GMAIL.COM

At their root, decision theories follow the dictum: when given a choice, choose what maximizes your preferences. The most difficult part of this theory is that preferences must exist before decisions can be made. The standard response to the basic question “Where do preferences come from?” is to suggest that they are learned. While much research has been invested in the idea that animals and people estimate values from experience by estimating utilities from experienced rewards, a broad array of empirical demonstrations of failures of utility theory suggests the idea of representing preferences using expected values and state-specific utility is fundamentally flawed (Vlaev, 2011). We propose that human preferences emerge from a process of inference, wherein subjects learn the desirability of options by rationally compiling noisy and incomplete internal feedback about the relative quality of options within a decision context. Rather than having fixed utilities for items, preferences are constructed anew in each decision by aggregating past experiences across similar contexts using Bayesian inference. This foundational change allows us to build a Bayesian theory of context-sensitive value inference that demonstrates many advantages over existing static representations of value. Specifically, it endogenously explains the emergence of all known context effects in the literature (attraction, similarity, compromise and reference set effects), endogenously induces novelty-seeking biases and is generalizable to unseen options in a value inference framework that is inductively rational, and specializes to the standard economic representation of value when feedback is preference consistent across contexts.

III-22. Phasic locus coeruleus activity changes with practice: a pupillometry study

Alexander Petrov
Taylor R. Hayes
Ohio State University

APETROV@ALEXPETROV.COM
HAYES.335@OSU.EDU

The locus coeruleus (LC) is an important neuromodulatory nucleus that is the principal source of norepinephrine for the forebrain and has been implicated in a range of cognitive functions. Neurophysiological recordings with monkeys show a tight correlation between LC activity and pupil diameter (under constant ambient illumination). We recorded pupil diameter as a non-invasive reporter variable for LC activity in humans and investigated how it changed with practice on a perceptual discrimination task. Such learning data are hard to obtain with animals because they need extensive training just to perform the task properly. After a verbal instruction, brief demo, and pretest, 18 human participants practiced visual motion-direction discrimination for four sessions of 896 trials each. Task difficulty was manipulated between subjects. The stimuli were filtered-noise textures whose direction of motion differed by 8 deg in the Easy group and 4 deg in the Hard group. The behavioral results showed significant improvements in accuracy and response times, with significantly greater learning in the Hard group. Baseline (tonic) pupil diameter was measured during the fixation period on each trial. Task-evoked (phasic) deviations from this baseline were tracked with millisecond resolution throughout the trial. The peak magnitude of the phasic pupillary response decreased systematically across sessions in the Hard group, but not in the Easy group. We interpret this within the framework of the Adaptive Gain Theory (Aston-Jones & Cohen, 2005). On this interpretation, LC acts as a temporal filter, with phasic bursts signaling the onset of task-relevant stimuli and modulating the rate of information accumulation in the decision units. This temporal filter becomes better calibrated and more efficient with practice, but only in the Hard group. The easy discrimination does not need - and hence does not foster - great temporal precision because the high signal-to-noise ratio allows correct classification even under suboptimal timing.

III-23. Dynamic integration of sensory evidence and diminishing reward in perceptual decision making

Alireza Soltani^{1,2}
Ya-Hsuan Liu³
Shih-Wei Wu³

ASOLTANI@STANFORD.EDU
G30001023@YM.EDU.TW
SWWU@YM.EDU.TW

¹Stanford University

²HHMI

³National Yang-Ming University

Most natural decisions require integration of information over time and are made under time constraints, commanding a speed-accuracy-tradeoff (SAT). SAT has been widely used to study the dynamics of information processing in many cognitive functions. However, because in most studies the subjects were instructed about time constraints, and/or fixed payoffs were associated with correct and incorrect responses, it is unclear how SAT is learned and shaped by reward feedback, or what happens when the payoff is directly determined by the reaction time (RT). Here we used a novel experimental paradigm and an accompanying computational model to investigate these questions. The subjects performed two versions of a motion discrimination task in which they had limited time to report the direction of a patch of moving dots in order to obtain reward. Critically, in our experiments, the information about the remaining time (in control experiment) or the amount of reward (in reward experiment) was presented on the screen, and a reward feedback was provided at the end of each trial. Firstly, we found that during the control experiment subjects adjusted their decision processes to time constraints and exhibited SAT. Secondly, during the reward experiment, the subjects learned to combine sensory information and diminishing reward nearly optimally. Thirdly, whereas we found improvements in performance over the course of both experiments, only in the reward experiment there was a decrease in the RT for correct trials and the RT was ~100 msec shorter following rewarded trials than unrewarded trials. Finally, we used a plausible decision network model to explore possible changes that can explain our experimental observations. Overall, our experimental and modeling results provide evidence for emergence of SAT due reward feedback through trial-to-trial adjustments of the decision threshold, and for slow changes in decision processes via adjustments of decision network's connectivity.

III-24. Neuronal signatures of strategic decisions in posterior cingulate cortex

David Barack
Michael Platt
Duke University

DLB28@DUKE.EDU
PLATT@NEURO.DUKE.EDU

Strategic decisions guide the flexible allocation of time between tasks or behaviors. We probed the ability of animals to switch from foraging at one depleting resource to forage at a new one, known as the patch-leaving problem. Previously, we established that neurons in the anterior cingulate cortex (ACC) implement this strategic decision via a rise-to-threshold mechanism. ACC neurons increase their firing rate over the course of a patch, thresholding on the trial that the animals decides to leave. Here, we extend our investigation of the neural mechanisms of strategic choice to the posterior cingulate cortex (CGp). Half of our CGp neurons (44/96) significantly encode leave trials vs. stay trials in a binary manner. Furthermore, this state switch in neuronal activation state occurs in the last one or two trials in a patch. We hypothesize that this state change reflects the strategic decision by the animal to exit the current patch and engage a new one.

III-25. Reconciling decisions from description with decisions from experience

Komal Kapoor
Nisheeth Srivastava
Paul Schrater
University of Minnesota

KOMAL.BITSGOA@GMAIL.COM
NISHEETHS@GMAIL.COM
SCHRATER@GMAIL.COM

Scientists have postulated multiple explanations for a systematic decision-experience gap (Erev & Roth, 2005) in subjects' understanding of the probability of rare events, wherein subjects that are explicitly told the probabilities of success in gambles systematically overweight the likelihood of occurrence of rare outcomes, while subjects that implicitly infer these probabilities are likely to underweight rare outcomes. In this work, we present results from a simulation study where we show a transition from underweighting rare outcomes to overweighting them within the same decision model as a function of experience length. This finding provides theoretical support for the proposal made by (Hertwig, 2011) that increased access to experience is the primary causal factor for the decision-experience gap. We also show accurate prediction results on three different certainty equivalence task settings using a single choice model without having to resort to statistical parameter fitting. Our results are quantitatively competitive with the best results obtained by different algorithms tested against each other in the Technion Prediction Tournament (Erev, 2010), all of which used between 3-6 parameters fit to a training set of human data in order to make predictions. This work builds upon a theory of belief dynamics we have previously developed (Srivastava & Schrater, 2011) which postulates that belief formation in humans emerges as a solution to the computational problem of recalling past beliefs into working memory from long-term memory in as cognitively efficient a manner as possible. By retrieving human-like performance across different experimental conditions without statistical parameter fitting, we obtain robust evidence supporting cognitively efficiency as a theory of belief formation. Finally, the systematic emergence of a bias in probability with increased access to experience cannot be explained by the learning based computational accounts of behavior that have been popularized by the Bayesian modeling paradigm in computational cognitive science.

III-26. Hierarchy of intrinsic timescales across primate cortex

John Murray¹
Alberto Bernacchia¹
Tatiana Pasternak²
Camillo Padoa-Schioppa³
Daeyeol Lee¹
Xiao-Jing Wang⁴

JOHN.MURRAY@YALE.EDU
ALBERTO.BERNACCHIA@YALE.EDU
TANIA@CVS.ROCHESTER.EDU
CAMILLO@WUSTL.EDU
DAEYEOL.LEE@YALE.EDU
XJWANG@NYU.EDU

¹Yale University

²University of Rochester

³Washington University in St. Louis

⁴New York University

Primate cortex is hierarchically organized, and different cortical areas appear specialized for diverse computations. However, the neural circuit basis underlying these areal specializations remains an open question. One hypothesis is that differences in local circuit properties across areas may be responsible for this specialization. We hypothesize that, at the physiological level, these differences can be detected in terms of differential timescales of neural dynamics. To test this hypothesis, we studied temporal autocorrelation of single-neuron spike trains in multiple areas, across sensory, parietal, and prefrontal cortex, recorded in monkeys performing cognitive tasks. We focused on activity during the task's fixation period to isolate internal dynamics and facilitate comparison across different datasets. In a given cortical area, decay of autocorrelation is well described by a characteristic timescale, which reflects intrinsic firing rate fluctuations within single trials. Across areas, timescales follow a hierarchical ordering, with sensory and prefrontal areas exhibiting short and long timescales, respectively, spanning an order of magnitude. The hierarchy of intrinsic timescales correlates with hierarchies derived from long-range anatomi-

cal projections, linking physiological and anatomical measures. The autocorrelation decay also exhibits an offset, which may reflect rate fluctuations on slower timescales. In particular, the offset could reflect the strength of across-trial memory encoded in firing activity. To test this possibility, we used a decision-making task that demands across-trial memory, in which neurons exhibit a characteristic timescale for memory of past rewards. In line with this interpretation, we find that autocorrelation offset correlates with the timescale of reward memory. To explore potential mechanisms, we studied a spiking neural circuit model. We find that strengthening recurrent structure in the network increases the intrinsic timescale, within the range observed experimentally. We propose the hierarchical gradient of intrinsic timescales across cortical areas reflects specialization of local circuit properties for diverse computations.

III-27. Sparse coding and dynamic suppression of variability in balanced cortical networks

Farzad Farkhooi^{1,2}
Martin Nawrot³

FARZAD.FARKHOOI@FU-BERLIN.DE
NAWROT@FU-BERLIN.DE

¹FU Berlin

²BCCN Berlin

³FU and BCCN Berlin

Many lines of evidence suggest that only few spikes carry the relevant stimulus information at later stages of sensory processing. Yet the mechanisms for the emergence of sparse sensory representations remain elusive. Here, we introduce a novel idea in which a temporal sparse and reliable stimulus representation develops naturally in spiking networks[1]. It combines principles of signal propagation with the commonly observed mechanism of neuronal firing rate adaptation. Using a stringent mathematical approach we show how a dense rate code at the periphery translates into a temporal sparse representation in the balanced cortical network. At the same time it dynamically suppresses the trial-by-trial variability, matching the experimental observation in sensory cortices[2]. Computational modelling and experimental measurements suggest that the same principle underlies the prominent example of temporal sparse coding in the insect mushroom body. Our results reveal a computational principle that relates neuronal firing rate adaptation to temporal sparse coding and variability suppression in nervous systems. [1] F. Farkhooi, et al.: arXiv:1210.7165 (2012) [2] M. M. Churchland, et al.: Nat Neurosci 13 (2010) 369

III-28. Combining feed-forward processing and sampling for neurally plausible encoding models

Jorg Lucke¹
Jacquelyn A. Shelton²
Jorg Bornschein²
Philip Sterne²
Pietro Berkes³
Abdul-Saboor Sheikh²

LUECKE@FIAS.UNI-FRANKFURT.DE
SHELTON@FIAS.UNI-FRANKFURT.DE
BORNSCHIEIN@FIAS.UNI-FRANKFURT.DE
STERNE@FIAS.UNI-FRANKFURT.DE
PIETRO.BERKES@GOOGLEMAIL.COM
SHEIKH@FIAS.UNI-FRANKFURT.DE

¹Goethe-University of Frankfurt

²Goethe-University Frankfurt

³Brandeis University

Novel experimental studies support the hypothesis that neural activity states can be interpreted as samples approximating posterior distributions. In the absence of a stimulus, spontaneous activity is accordingly interpreted as samples from a prior. The sampling hypothesis implies, for instance, that averages over the evoked activity given natural stimuli should become similar to activity distributions in the absence of stimuli as an animal ma-

tures. Such increasing similarities are indeed observed in in vivo recordings in ferrets. However, this and related predictions are to large extents independent of a particular model for stimulus encoding, and more detailed predictions could be made based on specific encoding models. Here we study such concrete models that are (a) consistent with the sampling hypothesis and (b) capture the high efficiency of neural inference and learning also for high dimensional tasks. In the studied models sampling based posterior approximations are combined with parametric representations of a variational approximation of the posterior. Neurally, the combined approximation method can be interpreted as a fast feed-forward preselection of the relevant state space, followed by a neural dynamics implementation of Gibbs sampling to approximate the posterior over the relevant states. The used models are variants of sparse coding and describe the encoding of visual stimuli. First, we study the efficiency and scaling behavior for high dimensional stimulus and hidden spaces using an elementary linear model. Numerical experiments verify the efficiency of the approach and demonstrate robustness of learning. Second, we study extensions of this model through the application of nonlinearities and flexible learnable priors. While sampling based inference is retained, flexible priors allow for modeling changes in spontaneous activity and allow for comparisons between learned prior vs. averaged posterior distributions. If trained on natural images, such models furthermore connect the sampling hypothesis to studies on V1 receptive fields.

III-29. Visual target signals are computed via a dynamic ‘and-like’ mechanism in IT and perirhinal cortex

Marino Pagan
Nicole Rust

MPAGAN@MAIL.MED.UPENN.EDU
NRUST@PSYCH.UPENN.EDU

University of Pennsylvania

Finding specific target objects requires the brain to compute nonlinear conjunctions, or ‘and-like’ computations, between visual and target-specific signals to determine when a target is currently in view. To investigate how the brain implements these computations, we recorded the responses of neural populations in inferotemporal cortex (IT) and perirhinal cortex (PRH) as monkeys performed a delayed-match-to-sample sequential target search task that required finding different targets in different blocks of trials. Consistent with combinations of visual and target-specific information that happen within or before IT, we found that the two populations contained similar amounts of total information for this task, as assessed by an ideal observer analysis. However, information about whether a target was currently in view was more accessible to a linear read-out in PRH, consistent with additional conjunction computations in PRH that act on the inputs from IT. To investigate how the brain implements these conjunctions, we focused on the dynamics of the neural signals. In both areas, we found that the latency with which the conjunction information appeared was delayed relative to the latency of visual information alone. Both areas included subpopulations of individual neurons whose initial responses reflected visual information, followed by an evolution of the responses to reflect conjunctions of the visual stimulus and the target. Target signals thus acted via a gating mechanism in these neurons to implement an ‘and-like’ operation with a delay. These results are reminiscent of the delayed ‘and-like’ responses reported for motion processing (Pack and Born, 2001; Smith et al., 2005), shape processing, (Brincat et al., 2006) and other visual search tasks (Chelazzi et al., 1993), and taken together, these results suggest that the ‘and-like’ computations required for visual perception and cognition may be implemented via a canonical mechanism.

III-30. On the information capacity of spike trains and detectability of rate fluctuations

Shinsuke Koyama
The Institute of Statistical Mathematics

SKOYAMA@ISM.AC.JP

Neural activity in cortex exhibits variability, the characteristics of which are specific to each neuron and are strongly

correlated with the function of the cortical area [1]. However, its neural coding advantages, or limitations posed by this variability are not completely understood. In this study, we address a question of how the information about fluctuating firing rates, carried by spike trains, is related quantitatively to the variability of neuronal firing. For this purpose, we introduce the relative entropy of rate-modulated renewal spike trains as a measure of information on fluctuating firing rates. We then give an explicit interpretation of this information in terms of detectability of rate fluctuation: the lower bound of detectable rate fluctuation, below which the temporal variation of firing rate is undetectable with a Bayesian decoder, is entirely determined by this information [2]. We show that the information depends not only of the fluctuating firing rates (i.e., signals), but also significantly on the dispersion properties of firing described by the shape of the renewal density (i.e., noise properties). Interestingly, the gamma density gives the theoretical lower bound of the information among all renewal densities when the coefficient of variation of interspike intervals (ISIs) is given. We provide a procedure to estimate the information from spike trains, and report results of the real data analysis.

III-31. Response properties of sensory neurons artificially evolved to maximize information

Brian Monson
Yaniv Morgenstern
Dale Purves
Duke University

BRIAN.MONSON@DUKE-NUS.EDU.SG
YANIV.MORGENSTERN@GMAIL.COM
PURVES@NEURO.DUKE.EDU

Successful behavior entails efficient processing of sensory inputs from biological sensors. One information theoretic principle proposed to guide such sensory processing is for neural responses to maximize information transfer by reflecting the cumulative probability distribution (CDF) of sensory inputs (Laughlin, 1981). Although there have been models implementing this strategy (e.g., Bell & Sejnowski, 1995), the biological circuitry required to accomplish this is unclear. Here we employ biologically feasible neural networks to examine response properties of early auditory and visual sensory neurons that use a processing strategy similar to this maximum information transfer approach. Modeling each input neuron with a two-layer sum-of-sigmoids network (Poirazi et al., 2003), we optimize the network responses to represent conditional CDFs of (1) natural acoustic intensities distributed across a one-dimensional filter bank sensor array (the basilar membrane), and (2) natural luminance intensities distributed across a two-dimensional spatial sensor array (the retina). Neurons evolve response characteristics similar to those of early-level auditory and visual sensory neurons. The evolved response characteristics include center-surround receptive fields; two-tone suppression (as seen in auditory neurons); and adaptation to ambient stimulus intensity. Evolved characteristics of both auditory and visual neurons are independent of physical location on the sensor array. These results suggest that response properties of biological sensory neurons observed in physiology are the consequence of a strategy that maximizes information transfer based on the frequency of occurrence of natural stimulus patterns.

III-32. A novel method for fMRI analysis: Inferring neural mechanisms from voxel tuning

Rosemary Cowell
David Huber
John Serences
University of California, San Diego

RCOWELL@UCSD.EDU
DHUBER@UCSD.EDU
JSERENCES@UCSD.EDU

Recent methods for analyzing fMRI data produce voxel tuning functions (VTFs) that relate the value of a stimulus feature (e.g., orientation) to the intensity of the BOLD signal. The characteristics of VTFs have been interpreted as reflecting characteristics of the underlying neural tuning functions (NTFs) that contribute to them. However,

knowing how the shape of a voxel response profile is modified by a change in brain state (e.g., viewing stimuli at low versus high contrast) does not tell us how the response profiles of neurons contributing to that voxel are modified. Mapping a VTF back to NTFs is an ill-posed inverse problem: there are two unknown distributions (the shape of underlying NTFs and the response magnitude for each NTF) but only one observed distribution (the BOLD signal across values of the stimulus feature). We tackled this inverse problem by using two BOLD response profiles from two brain states (across which VTF shape is modulated) and solving for modulations in the distributions. We collected BOLD data from V1 in subjects viewing oriented sinusoidal gratings at low and high stimulus contrast. Taking orientation-selective voxel responses at low versus high contrast, we fitted multiple alternative models of the modulation of NTFs (additive shift, multiplicative gain, bandwidth narrowing) assumed to drive the modulation in the VTF. We used parametric bootstrapping to penalize overly flexible models. Although the VTF underwent additive shift from low to high contrast, the best-fitting models of NTF modulation accounting for this shift involved primarily multiplicative gain (in line with electrophysiological evidence). This demonstrates that the method can recover 'ground truth' by making use of the constraints imposed by many voxels across two conditions. The method links monkey neurophysiological data concerning NTFs to human fMRI data on VTFs and should be applicable in other (non-visual) sensory cortices.

III-33. Successful prediction of a physiological circuit with known connectivity from spiking activity

Felipe Gerhard¹
 Tilman Kispersky²
 Gabrielle J Gutierrez²
 Eve Marder²
 Mark Kramer³
 Uri Eden³

FELIPE.GERHARD@EPFL.CH
 TILMAN@BRANDEIS.EDU
 GG99@BRANDEIS.EDU
 MARDER@BRANDEIS.EDU
 MAK@MATH.BU.EDU
 TZVI@MATH.BU.EDU

¹Brain Mind Institute

²Brandeis University

³Boston University

Identifying the structure and dynamics of synaptic interactions between neurons is the first step to understanding neural network dynamics. The presence of synaptic connections is traditionally inferred through the use of targeted stimulation and paired recordings or by post-hoc histology. More recently, causal network inference algorithms have been proposed to deduce connectivity directly from electrophysiological signals, such as extracellularly recorded spiking activity. These algorithms have not been validated on a neurophysiological data set for which the actual circuitry is known. Recent work has shown that traditional network inference algorithms based on linear models typically fail to identify the correct coupling of even a basic three-neuron circuit like the crab stomatogastric nervous system. In this work, we show that point process models that incorporate the spike train nature of the data can correctly infer the physiological connectivity of a three-neuron circuit. We elucidate the necessary steps to derive faithful connectivity estimates from spike train observations alone. We then apply the model to measure changes in the effective connectivity pattern in response to pharmacological interventions, which affect both intrinsic neural dynamics and synaptic transmission. Our results provide the first successful application of a network inference algorithm to a circuit for which the actual physiological synapses between neurons are known. The point process methodology presented here generalizes well to larger networks and can describe the statistics of neural populations. In general we show that advanced statistical models allow for the characterization of effective network structure, deciphering underlying network dynamics and estimating information-processing capabilities.

III-34. Evidence for presynaptic inhibition in shaping retinal processing

Yuwei Cui
 Yanbin Wang
 Jonathan Demb
 Daniel A Butts
 University of Maryland

YWCUI@UMD.EDU
 YANBIN.WANG@YALE.EDU
 JONATHAN.DEMB@YALE.EDU
 DAB@UMD.EDU

Retinal circuitry has been extensively characterized, but it is less clear how different circuit elements contribute to visual processing. For example, it is difficult to untangle the multiple synaptic pathways — both direct and indirect — that contribute to a retinal ganglion cell's response. Furthermore, the models of stimulus processing that can be constrained by data have tended to be overly simplistic. Here, we address both problems by isolating the excitatory synaptic input of specific ganglion cell types, ON- and OFF-Alpha cells, using a new nonlinear modeling approach. Excitatory inputs to ganglion cells driven by temporal noise stimuli showed fine-time-scale features not explainable by typical linear analyses. We hypothesized that presynaptic inhibition might explain these nonlinear response properties, and accordingly we tested a cascade model that combined excitatory and suppressive filters coupled with a two-dimensional nonlinear mapping function. The model could be efficiently fitted to the recordings of synaptic currents and predicted nearly 100% of the explainable variance for both ON- and OFF-Alpha cells, representing a dramatic improvement in predictive power over standard approaches. Moreover, the model's components had a direct physiological interpretation: the identified two-dimensional nonlinearity could be explained by a multiplicative interaction between excitatory and suppressive components, with suppression delayed relative to excitation. By using a stimulus with separately modulated center and surround noise components, we found that such multiplicative suppression arises largely from the surround. The different spatial profiles of excitation and suppression, as well as an apparent ON-OFF tuning of suppression, suggested that inhibition is mediated by amacrine cell synapses onto bipolar cell terminals. By targeting a nonlinear model structured to reveal specific effects of retinal circuitry, we identified a novel computational property of retinal processing.

III-35. Memory maintenance in calcium-based plastic synapses in the presence of background activity

David Higgins^{1,2}
 Michael Graupner³
 Nicolas Brunel²

DAVE@UIGINN.COM
 MICHAEL.GRAUPNER@NYU.EDU
 NBRUNEL@GALTON.UCHICAGO.EDU

¹École Normale Supérieure
²University of Chicago
³New York University

How are synaptic modifications, elicited by specific patterns of activity, stable in the face of ongoing background activity? We investigate the life-time of synaptic changes using a calcium-based synaptic plasticity model (Graupner and Brunel, 2012), which fits spike-timing dependent plasticity results from cortical slices (Sjostrom et al, 2001). The model allows us to compute analytically how the average synaptic strength evolves in time, in the presence of random pre- and post-synaptic Poisson firing. We furthermore investigate the effect of the extracellular calcium concentration on the time scale of synaptic changes by considering two conditions : (i) in vitro experiments for which 2.5 mM extracellular calcium has been used (Sjostrom et al, 2001), and (ii) in vivo-like levels of 1.5 mM calcium (Silver and Erecinska, 1990). We find that the memory of the initial synaptic state decays exponentially, with a time constant that is inversely proportional to the background firing rate if the synapse has a continuum of stable states at rest. The time constant decays according to a power law with an integer exponent that depends on the sizes of the calcium transients triggered by pre- and post-synaptic spikes. Importantly, the predicted memory timescale for the in vivo-like extracellular calcium level is several orders of magnitude longer than that for the in vitro level (hours vs minutes, for a 1/s background rate). In case the synapse has two stable states at rest, we find that in vivo synapses can retain their memories for long periods (~months) under low frequency background

(<1 Hz). These results are qualitatively similar in large-scale simulations of a network of excitatory and inhibitory neurons, operating at low rates. Our findings emphasise the role of synaptic bistability and of the extracellular calcium level for memory retention in cortical circuits in the presence of realistic background activity.

III-36. Dendritic subunits: the crucial role of input statistics and a lack of two-layer behavior

DJ Strouse¹
Balazs B Ujfalussy²
Mate Lengyel²

DANIELJSTROUSE@GMAIL.COM
BBU20@CAM.AC.UK
M.LENGYEL@ENG.CAM.AC.UK

¹Princeton University

²University of Cambridge

Accumulating evidence suggests that dendritic trees play a crucial role in single-neuron information processing, yet there exists no simple, canonical formalization of dendritic computation. At one extreme, multi-compartmental models retain as much biophysical detail as possible, enabling them to exhibit the spatial and temporal dendritic nonlinearities observed in experiments, but sacrificing ease of fitting, mathematical tractability, and computational interpretability. At the opposite extreme, heuristic ‘two-layer network’ models [Poirazi et al., *Neuron* 2003, Polsky et al., *Nat. Neuro.* 2004], which assume that the somatic membrane potential is produced by passing the instantaneous synaptic inputs through a two-layer linear-nonlinear cascade, are easy to analyze mathematically and interpret computationally but were designed to work only with static inputs and outputs, restricting their experimental application to artificial stimulus protocols involving brief, intense stimulation, rather than extended spike trains with realistic statistical properties. Moreover, because of this restriction, the associated metrics for judging nonlinear dendritic behavior were based only on either instantaneous firing rates or peaks/means of somatic membrane potentials, rather than predictiveness of dynamically changing firing rates or full membrane potential traces.

III-37. Physiology and impact of horizontal connections in rat neocortex

Philipp Schnepel^{1,2}
Ad Aertsen^{1,2}
Arvind Kumar^{1,2}
Clemens Boucsein^{1,2}

PHILIPP.SCHNEPEL@BIOLOGIE.UNI-FREIBURG.DE
AERTSEN@BIOLOGIE.UNI-FREIBURG.DE
ARVIND.KUMAR@BIOLOGIE.UNI-FREIBURG.DE
CLEMENS.BOUCSEIN@BIOLOGIE.UNI-FREIBURG.DE

¹Bernstein Center Freiburg

²University of Freiburg

Cortical information processing at the cellular level has predominantly been studied in networks with strong local connectivity, resulting in models displaying high noise correlations. However, recent studies suggest that the bulk of axons targeting pyramidal neurons most likely originate from outside the local volume. This opens interesting new possibilities to influence the input statistics of the neurons within the cortical network. For example, horizontal projections have been implicated to reduce noise correlations and improve the signal-to-noise ratio in reaction to external inputs. Unfortunately, no detailed data about physiology, numbers and spatial extent of horizontal connections is available to date. We, therefore, mapped the horizontal connectivity of L5B pyramidal neurons with photostimulation, identifying intact projections up to a lateral distance of 2mm. Our estimates of the spatial distribution of cells presynaptic to L5B pyramids support the idea that their majority is located outside the local volume. In addition, the synaptic physiology of distant horizontal connections does not differ markedly from that of local connections, while the layer and cell-type dependent pattern of innervation does. Implementing our data in a reduced model of a neocortical network shows that, indeed, the identified horizontal connections can promote robust asynchronous on-going activity states and reduce noise correlations in stimulus-induced activity and may,

thus, be a means for the neocortex to improve signal detection. More specifically, a new role for layer 6A emerges, since it provides a strong source of horizontal connections to L5B pyramids. In addition to its feedback projections to thalamus and its modulatory influence on the principal input layer of cortex (L4), L6A also seems to exert a strong influence on the principal output stage of cortical processing.

III-38. Gain-control via shunting-inhibition in a spiking-model of leech local-bend.

Edward Frady
William Kristan

EFRADY@UCSD.EDU
WKRISTAN@UCSD.EDU

University of California, San Diego

It was originally theorized that shunting-inhibition could be playing a key role in gain-control by computing division. However, several studies have suggested that spiking dynamics interfere with the shunting mechanism and prevent divisive scaling. New hypotheses have emerged that suggest that division could be implemented via shunting-inhibition, but require synaptic noise. Other hypotheses suggest that non-shunting circuit mechanisms must be computing division. We describe implementations of both shunting and circuit-based neuronal gain-control in a model of leech local-bend. Much of the neural circuitry responsible for the local-bend behavior has been well worked-out, giving us a strong biological framework to test our hypotheses. The key experimental insight of the local-bend behavior is that the information necessary for the reflex is encoded by a population-code in the neural activity. Experimental work further shows that GABAergic inhibition has a divisive effect on neuronal responses. We describe the challenges of implementing division with a circuit mechanism, requiring overly complex and highly specific wiring to maintain the population-code. We derive a new implementation of shunting-inhibition as a mechanism for division in a multi-compartmental spiking neuron model. This mechanism enables us to implement gain-control without the need for noise, and build a complete spiking model of the local-bend behavior based on a population-code. We further explore the differences of feed-forward and feed-back inhibitory gain-control. This local-bend model has a direct correspondence to contrast-invariant orientation-selectivity as seen in V1 simple-cells and implicates soma-targeting inhibitory neurons as possibly controlling gain through this mechanism. This model is biologically plausible, computationally efficient, and, importantly, mathematically tractable. The model will allow for rate-coded theories of neural computation — especially those that depend on population codes or divisive normalization, to be translated into theories and simulations that include spiking dynamics.

III-39. Detecting and quantifying topographic order in neural maps

Stuart Yarrow¹
Khaleel Razak²
Aaron Seitz²
Peggy Series¹

S.YARROW@ED.AC.UK
KHALEEL@UCR.EDU
AARON.SEITZ@UCR.EDU
PSERIES@INF.ED.AC.UK

¹University of Edinburgh

²University of California, Riverside

Topographic maps are an often-encountered feature in the brains of many species. The degree and spatial scale of smooth topographic organisation in neural maps vary greatly, as do the sampling density and coverage of techniques used to measure maps. An objective method for quantifying topographic order would be valuable for evaluating differences between, e.g. experimental and control conditions, developmental stages, hemispheres, individuals or species; to date, no such method has been applied to experimentally-characterised maps. Neural maps are typically identified and described subjectively, but in cases where the scale of the map is close to the resolution limit of the measurement technique, just identifying the presence of a map can be a challenging sub-

jective task. In such cases, an objective map detection test would be advantageous. To address these issues, we assessed seven measures (Pearson distance correlation, Spearman distance correlation, Zreihen measure, topographic product, topological correlation, wiring length and path length) by quantifying topographic order in three classes of cortical map model: linear gradient, orientation-like, and randomly scattered homogeneous clusters. We found that the first five of these measures were sensitive to weakly-ordered maps and effective at detecting statistically significant topographic order, based on noisy simulated measurements of neuronal selectivity and sparse spatial sampling of the maps. We demonstrated the practical applicability of these measures by using them to examine the arrangement of spatial cue selectivity in pallid bat primary auditory cortex. This analysis shows for the first time that statistically significant systematic representations of inter-aural intensity difference and source azimuth exist at the scale of individual binaural clusters. An analysis based on these measures could be applied in any situation where it is useful to demonstrate the presence of a neural map, or to quantify the degree of order in a map.

III-40. Efficient and optimal Little-Hopfield auto-associative memory storage using minimum probability flow

Christopher Hillar¹
 Jascha Sohl-Dickstein²
 Kilian Koepsell²

CHILLAR@MSRI.ORG
 JASCHA.SOHL-DICKSTEIN@GMAIL.COM
 KILIAN@BERKELEY.EDU

¹Redwood Center for Theoretical Neuroscience

²University of California, Berkeley

We present an algorithm to store binary memories in a Little-Hopfield neural network using minimum probability flow, a recent technique to fit parameters in energy-based probabilistic models. In the case of memories without noise, our algorithm provably achieves optimal pattern storage (which we show is at least one pattern per neuron) and outperforms classical methods both in speed and memory recovery. Moreover, when trained on noisy or corrupted versions of a fixed set of binary patterns, our algorithm finds networks which correctly store the originals. We also demonstrate this finding visually with the unsupervised storage and clean-up of large binary fingerprint images from significantly corrupted samples.

III-41. Signal processing in neural networks that generate or receive noise

Isao Nishikawa^{1,2}
 Kazuyuki Aihara¹
 Taro Toyoizumi³

ISAO-N@SAT.T.U-TOKYO.AC.JP
 AIHARA@SAT.T.U-TOKYO.AC.JP
 TARO.TOYOIZUMI@BRAIN.RIKEN.JP

¹University of Tokyo

²RIKEN

³RIKEN Brain Science Institute

Cortical neurons exhibit irregular activity patterns. However, the source of the cortical variability is unknown. Here, we study two different types of randomly connected networks of quadratic integrate-and-fire neurons that produce irregular spontaneous activity patterns: (a) a network model that has strong synaptic interactions and actively generates variability by chaotic nonlinear dynamics and (b) a network model that has weak synaptic interactions and receives noisy input, for example, by stochastic vesicle releases. Despite the difference in their sources of variability, these two models can behave almost identically in their baseline activity. Indeed, when parameters are set appropriately, the two models are indistinguishable based on their spontaneous activity patterns unless majority of neurons in the network are simultaneously recorded. In spite of the close similarity in their spontaneous activity patterns, the two models can exhibit remarkably different sensitivity to external input. External input to the former network can reverberate within the network and be successfully read out over long time because of the

strong synaptic interactions between neurons. On the other hand, input to the latter network rapidly decays because of the small synaptic interactions. We describe implications of this difference on population coding and activity dependent plasticity.

III-42. A spiking model of superior colliculus for bottom-up saliency

Jayram Moorkanikara Nageswaran¹

Micah Richert¹

Csaba Petre¹

Filip Piekniewski¹

Sach Sokol²

Botond Szatmary¹

Eugene Izhikevich¹

¹Brain Corporation

²Johns Hopkins University

NAGESWARAN@BRAINCORPORATION.COM

RICHERT@BRAINCORPORATION.COM

CSABA.PETRE@BRAINCORPORATION.COM

PIEKNIWSKI@BRAINCORPORATION.COM

SSOKOL3@JHU.EDU

SZATMARY@BRAINCORPORATION.COM

EUGENE.IZHIKEVICH@BRAINCORPORATION.COM

Active exploration of visual space is controlled by a saccadic system using a combination of bottom-up and top-down signals. Rate based models, (Itti and Koch 2001) and simple V1 models (eg. Li 2002), have been proposed to explain bottom-up or 'pop-out' attentional mechanisms. We present a biologically detailed spiking model of superior colliculus to generate saccadic eye movements using a competitive mechanism based on spike arrival times. Superior colliculus is modeled as a two layer network with feature-specific superficial layers driving a common deep layer. The superficial layers of superior colliculus are driven by retinal parasols cells, V1 layer 2/3 orientation-selective cells, V1 layer 2/3 color-selective cells, and MT cells. These enable luminance, orientation, color and motion popout respectively. In this model, the earliest spikes are considered the most salient. A saccade is triggered when there is a consistent winner in the superficial layers that is in the same spatial region over a period of approximately 100ms. The superficial layers implement a temporal winner-take-all mechanism; specifically, there exists a separate inhibitory subnetwork that allows each superficial layer to choose a winner. The deep layer integrates these winners to select a single spatial region. It implements a bump attractor network mediated by short-range excitation (AMPA and NMDA) and long range inhibition that generate a bump of activity before and during a saccade. The reaction time of the model was tested on simple and conjunctive pop-out tasks. We observe a only a mild shift in reaction time based on the number of distractors. Higher reaction time was observed as the complexity of the search task was increased. Also, reactions times qualitatively match latencies observed in the visual system. Finally, the model captures known behavioral and neurophysiological properties of the superior colliculus such as saccadic suppression and saccade build up.

III-43. Using optogenetics to probe neural circuit mechanisms in the alert, behaving non-human primate

Jonathan Nassi¹

Ali Cetin¹

Anna Roe²

John Reynolds¹

¹Salk Institute for Biological Studies

²Vanderbilt University

NASSI@SALK.EDU

ACETIN@SALK.EDU

ANNA.ROE@VANDERBILT.EDU

REYNOLDS@SALK.EDU

Optogenetics has quickly proven to be a powerful new tool for understanding the function of specific cell-types and circuits within the central nervous system, yet its application and widespread use in non-human primates has been slow to develop. One particular challenge has been the transdural delivery of viruses and light to the brain, which has so far required the use of injectrodes and optrodes that are damaging to the intricate cortical

circuitry being studied. Here, we report on a new approach to optogenetic techniques in the alert and behaving monkey which makes use of a silicone-based, optically-clear artificial dura. Replacing the native dura with an artificial dura allows for the use of fine glass micropipettes to inject virus, thin electrodes to record neural activity and optical fibers that need not enter the brain. These measures greatly reduce the potential to cause damage while enabling light-based activation with a high level of spatial precision. We injected several viruses into distinct locations within macaque primary visual cortex (V1), all of which were engineered to preferentially express the depolarizing opsin C1V1 in local excitatory neurons. Several weeks after virus injections, light delivered to V1 of an alert, fixating monkey through an optical fiber positioned above the artificial dura modulated the spontaneous and visually-evoked activity of recorded neurons. We observed both facilitation and suppression of spontaneous activity and visually-evoked contrast response functions - consistent with the interpretation that light stimulation leads to both direct excitation and indirect excitation and inhibition within the cortical circuit. This new approach promises to greatly assist in testing models of normalization and dissecting the cortical circuits underlying visual perception, cognition and behavior.

III-44. Being balanced: the role of the indirect pathway of the basal ganglia in threshold detection

Wei Wei¹
Jonathan Rubin²
Xiao-Jing Wang^{1,3}

WEI.WEI@YALE.EDU
JONRUBIN@PITT.EDU
X.JWANG@YALE.EDU

¹Yale University
²University of Pittsburgh
³New York University

The basal ganglia (BG) are believed to play an important role in voluntary motor control. For perceptual decision-making, the direct pathway of the BG has been suggested to provide a mechanism for effectively modulating the threshold for evidence accumulation in cortex [1]. Here we examined the contribution of the indirect pathway of the BG, through the external segment of the globus pallidus (GPe) and the subthalamic nucleus (STN), in this process. We noted that the experimentally observed ramping activity in caudate nucleus (CD), the input nucleus of the BG, during the evidence accumulation period [2], could have curtailed the efficiency of the direct pathway in adjusting the threshold as in [1]. We found that with the balancing between the direct and indirect pathways of the BG, the activity in the output nucleus of the BG, the substantia nigra pars reticulata (SNr), was no longer influenced by the ramping activity in cortex and CD. After the GPe activity was suppressed completely by CD ramping, the input from CD could sharply suppress the SNr activity. The resulting disinhibition of the downstream superior colliculus (SC), together with the input directly from cortex, elicited SC bursting, which then sent feedback to cortex and terminated the evidence accumulation process there, thereby indicating a threshold crossing for the decision process. We also investigated the impact of a parkinsonian state, characterized by enhanced bursting and synchrony in the beta band in GPe and STN, on threshold detection and performance in reaction time tasks.

III-45. The amygdala is critical for reward encoding in the orbital, but not medial, prefrontal cortex

Peter Rudebeck¹
Andrew Mitz²
Joshua Ripple²
Ravi Chacko²
Elisabeth Murray²

RUDEBECKP@MAIL.NIH.GOV
ARM@NIH.GOV
JOSHUA.RIPPLE@NIH.GOV
RAVI.CHAKO@GMAIL.COM
MURRAYE@MAIL.NIH.GOV

¹National Institute of Mental Health
²NIMH/NIH

It has been hypothesized that input from the amygdala is critical for reward-related neural responses in the orbital (PFo) and medial prefrontal cortex (PFm). Causal evidence for this hypothesis is scarce. We recorded the activity of neurons in the PFo and PFm of macaques while they made choices between sequentially presented visual stimuli that were associated with rewards, both before and after lesions of the amygdala. This is the first study of its kind in macaques. Preoperatively, we found that neurons in the PFo and PFm encoded stimulus-reward associations, but that neurons in PFo better represent the relationship between stimuli and rewards. Additional analyses revealed that neurons in the PFo encoded the reward amount associated with the temporally separated stimuli independently. Removing input from the amygdala abolished the difference between stimulus-reward encoding in the PFo and PFm, as fewer neurons in the PFo represented stimulus-reward associations. Preoperatively, neurons in both the PFo and PFm also encoded the amount of reward expected and then received after a choice, but again encoding was stronger in the PFo. Removing input from the amygdala dramatically reduced encoding of received rewards, but only in PFo. The difference was particularly apparent immediately following the onset of reward delivery. These data suggest that the amygdala is important for signaling the outcome of a choice to neurons in the PFo, but not PFm, and that these signals are important for encoding of stimulus-reward associations. Given the role of the PFo in contingent learning, we tested whether a lack of such a reward-outcome signal would alter monkey's ability to learn new stimulus-reward associations. Confirming this, monkeys with amygdala lesions were slower to learn stimulus-reward associations. These data have implications for reinforcement learning models and our understanding of frontal-temporal interactions during reward processing.

III-46. Spatial representation in the ventral hippocampus

Alexander Keinath¹
Joshua Dudman^{2,3}
Isabel Muzzio¹

AKEINATH@SAS.UPENN.EDU
DUDMANJ@JANELIA.HHMI.ORG
IMUZZIO@SAS.UPENN.EDU

¹University of Pennsylvania

²Janelia Farm Research Campus

³HHMI

The ways in which information is stored along the longitudinal hippocampal axis are not fully understood. The dorsal hippocampus receives heavy visual and spatial inputs, whereas the ventral hippocampus is connected with brain regions that control emotion and anxiety. These neuroanatomical differences are paralleled by an increase in the scaling of place fields from the dorsal to the ventral pole along a continuum. Typical analysis of neuronal activity in the dorsal hippocampus is based on the sparse tuning of dorsal representations and assumes that the activity of a given neuron represents the probability that the animal is in a particular location. However, for broadly tuned ventral cells this approach provides a poor description of position. We hypothesized that if activity in the ventral hippocampus reflects spatial experience, then one should observe neural trajectories that correlate with the trajectory of the animal through physical space. Furthermore, emotional inputs to the ventral region suggest that network activity in this area may be modulated by contextual valence. To investigate these possibilities, we conducted chronic single unit recordings from area CA1 in different regions along the longitudinal hippocampal axis while mice were exposed to a variety of visual and olfactory stimuli with different degrees of emotional valence. Using principal component analysis to obtain a low-dimensional representation of the population activity for all units recorded during a single session, we found that position information is represented in the ventral hippocampus. Bayesian and template matching methods were used to extract this information. Mean reconstruction error (MRE) was calculated to quantify the accuracy of this information. We found that the MRE was similar in the dorsal and ventral hippocampus, and was not significantly affected by the valence of the context. In summary, our results demonstrate that population activity in the ventral hippocampus provides accurate spatial information.

III-47. Data assimilation of individual HVC neurons using regularized variational optimization.

Mark Kostuk¹
 Daniel Meliza²
 Hao Huang²
 Alian Nogaret³
 Daniel Margoliash²
 Henry Abarbanel¹

MKOSTUK@PHYSICS.UCSD.EDU
 DMELIZA@UCHICAGO.EDU
 HUANGHAO198604@UCHICAGO.EDU
 PYSARN@BATH.AC.UK
 DAN@BIGBIRD.UCHICAGO.EDU
 HABARBANEL@UCSD.EDU

¹University of California, San Diego

²University of Chicago

³University of Bath

Accurately modeling single neuron behavior is imperative to understanding neuronal systems, both when a cell is regarded as a node in a larger network or when tracking long term changes in its underlying biological machinery (e.g. due to environmental modifications or learning induced plasticity). We estimate the parameters and unknown state variables of a single compartment neuron model based upon short segments (1.5 sec) of standard electrophysiology data taken from neurons in the pre-motor nucleus HVC of the birdsong system. The resultant models are capable of predicting the real neuron responses to novel stimuli as well as providing biophysically plausible justification for membrane voltage using the unmeasured state variables, which correspond to Hodgkin-Huxley style gating particles of the constituent ion channel currents. This approach allows for the efficient characterization of potentially large sample of neurons, on which work has begun in HVC. The assessment of model parameters over a population of neurons is crucial to the creation of representative, distribution-based network models. The estimates are obtained individually for multiple neurons using a variational optimization algorithm with a modification adapted from control theory to regularize nonlinear instabilities in the high-dimensional search surface. The optimization returns the set of unknowns with which the model dynamics produce the maximum likelihood trajectory of the distribution defined by the data-assimilation path integral of the model dynamics conditioned upon the experimental voltage measurements. The principles of the method are applicable to any deterministic dynamical system (including networks) provided sufficient measurements are available; furthermore the number and specific state variables that must be measured is an answerable question once a model is defined.

III-48. Decoding sound source location using neural population activity patterns

Mitchell Day^{1,2}
 Bertrand Delgutte¹

DAY@MEEI.HARVARD.EDU
 BERTRAND_DELGUTTE@MEEI.HARVARD.EDU

¹Harvard Medical School

²Massachusetts Eye and Ear Infirmary

The acuity of sound localization in the horizontal plane is sharpest for sounds directly in front of a listener (the midline). However, neurons along the auditory pathway from the brainstem to the cortex are, in general, broadly tuned to contralateral source locations, placing the slopes of azimuth tuning functions over the midline. It has been proposed that sound azimuth in a given hemifield is estimated from the summed firing rates across the population of neurons on the side of the brain contralateral to the sound source, presumably placing the steepest slope of the population-rate tuning function over the midline. We tested the localization performance of this 'population-rate' decoder using azimuth tuning functions measured from single units in the inferior colliculus (IC) of awake rabbits. The population-rate decoder performed very poorly at localizing sounds in the contralateral hemifield, inconsistent with unilateral lesion studies that have demonstrated normal localization of sounds contralateral to the intact side. The population-rate decoder ignores the substantial heterogeneity of azimuth tuning functions observed across the IC. This heterogeneity suggests the pattern of activity across IC neurons may be unique for each azimuth and serve as a basis for estimating source location. We created a 'population-pattern' decoder where a multivariate

Gaussian probability density of spike count was estimated for each azimuth, with each dimension being the spike count response of each neuron in the population. Each test trial of population activity was then classified to the azimuth with maximum likelihood. The population-pattern decoder accurately localized sources throughout the frontal hemifield, demonstrating that the pattern of firing rates in the IC is sufficient to estimate source azimuth.

III-49. Feedforward inhibition controls tuning in auditory cortex by restricting initial spike counts

Ashlan Reid
Tomas Hromadka
Anthony Zador
Cold Spring Harbor Laboratory

REIDA@CSHL.EDU
TOMAS.HROMADKA@GMAIL.COM
ZADOR@CSHL.EDU

Cortical neurons respond selectively to features of the sensory world based on the interplay of their incoming excitatory and inhibitory synaptic inputs. The relative role of inhibition in shaping selective responses, however, remains controversial in many sensory modalities. Here, we examined the role of a particular source of cortical inhibition, specifically parvalbumin (PV) expressing interneurons, in establishing tone frequency selectivity in the auditory cortex. We used temporally controlled optogenetic silencing of PV interneurons to isolate their contribution to the tone frequency selectivity of their postsynaptic targets. We performed whole cell and cell-attached recordings from cortical neurons in vivo in order to measure both spiking responses and the excitatory and inhibitory synaptic inputs underlying them. Surprisingly, we found a dual phase role of PV interneurons: in the early phase of the response, silencing PV inhibition increased the number of spikes by directly reducing tone-evoked inhibition, without changing the shape of the frequency tuning curve; in the late phase of the response, however, these extra spikes were correlated with aberrant excitation and broadening of the tuning curve. Our results suggest that PV interneurons enforce sharp selectivity in the auditory cortex by restraining excitation rather than by directly imposing lateral inhibition. Thus, the early gain modulation imposed by PV interneurons can have more far reaching effects than simply imposing a linear shift in cortical firing rates. By quenching the response in time, early feedforward inhibition from PV interneurons prevents late accumulation of excitation at non-optimal frequencies. This tight temporal control might also enhance the response fidelity to more complex, temporally modulated sounds. Together, our results demonstrate the importance of the precise control over spike number in establishing receptive fields in auditory cortex and show that transient alterations in activity can have long lasting consequences.

III-50. Toward a mechanistic description of shape-processing in area V4

Anirvan Nandy¹
Tatyana O. Sharpee^{1,2}
John Reynolds¹
Jude Mitchell¹

NANDY@SALK.EDU
SHARPEE@SALK.EDU
REYNOLDS@SALK.EDU
JUDE@SALK.EDU

¹Salk Institute for Biological Studies
²University of California, San Diego

Area V4 receives inputs from early visual areas (V1/V2) and provides input to inferotemporal cortical areas (TEO & TE). We have a poor understanding of V4's role in transforming the simple features encoded in early visual areas into complex object selectivity found in inferotemporal cortical neurons. We sought to gain insight into this by examining the fine structure of shape tuning in area V4. Using a set of parametrically controlled artificial stimuli, we find that V4 neurons exhibit considerable diversity in their shape tuning properties. At one end of the spectrum, neurons exhibit homogenous fine-scale orientation tuning maps (Fig 1). At a coarse scale, these neurons respond best to straight contours and exhibit robust position invariance. At the other end of the spectrum

are neurons that have very heterogeneous fine-scale orientation maps (Fig 2). These neurons respond best to curved contours, but exhibit very limited position invariance. Curvature preference for these neurons varies considerably across the receptive field. A simple local pooling model derived from a neuron's fine-scale response map predicts neuron's selectivity for curved contours. This holds true across the gamut of shape selective cells. We also simultaneously characterized a sub-set of neurons using natural scene stimuli. Preliminary analysis by projecting the scene information into a V1-like orientation space followed by gain-control (divisive normalization, [1]) and correlating with spikes, reveals spatio-temporal kernels that were very similar to those obtained with the artificial stimuli (Fig 3). Thus our characterization appears to be robust across diverse stimuli categories. The combination of pooling of fine-scale orientation selectivity and normalization allow us to construct a parsimonious and mechanistic description of shape processing in V4. We are now examining how these computations are mediated by the laminar cortical circuit, by simultaneously recording from neurons within a cortical column using linear array electrodes.

III-51. Object selectivity and tolerance to variation in object appearance trade off across rat visual corti

Sina Tafazoli
Houman Safaai
Matilde Fiorini
Gioia De Franceschi
Davide Zoccolan

TAFAZOLISINA@GMAIL.COM
HOUMAN1359@YAHOO.COM
FIORINI.MATILDE@GMAIL.COM
GIOIA.DFR@GMAIL.COM
ZOCCOLAN@SISSA.IT

School for Advanced Studies (SISSA)

The key computational problem of object recognition is attaining both selectivity among different visual objects and tolerance to variation in each object's appearance (e.g., as a result of position and size changes). The primate visual system appears to have solved this problem by gradually building increasingly selective, yet tolerant, representations of visual objects across the hierarchy of visual areas known as the ventral visual stream. It remains unknown whether, in rodents, a similar processing stream exists, thus accounting for the recently demonstrated invariant object recognition abilities of rats (Tafazoli et al, 2012). We are currently investigating this topic, by recordings the neuronal responses to the presentation of a battery of visual objects in two different cortical areas of the rat brain: primary visual cortex (V1) and temporal association cortex (TeA), with the latter being the candidate highest visual area of a putative rat ventral visual stream. Our object set consists of 10 different objects, each transformed across a variety of axes (position, size, in-depth azimuth rotation and in-plane rotation), to assess both object selectivity and transformation-tolerance of neuronal responses. Visual objects are presented in rapid sequence (for 250 ms, followed by a blank of 250 ms) in anesthetized rats. Our preliminary results show that object selectivity increases from V1 to TeA, while the opposite happens for tolerance to size and azimuth changes. Moreover, selectivity and tolerance appear to trade off both within each area and across V1 and TeA, with pairs of selectivity-tolerance values spreading along the same negative-slope line in a selectivity vs. tolerance plot. Thus, although very preliminary, these data suggest that the rat visual system may not have achieved the concurrent increase of selectivity and tolerance that is typical of primates. This raises the question of what mechanisms can support rat invariant visual object recognition.

III-52. Invariant population representations of objects are enhanced in IT during target search

Noam Roth
Margot Wohl
Nicole Rust

University of Pennsylvania

NOAMROTH@MAIL.MED.UPENN.EDU
WOHLMP@SAS.UPENN.EDU
NRUST@PSYCH.UPENN.EDU

Finding a specific visual target requires the brain to solve two computational challenges. First, the brain must determine the identities of the objects in a currently viewed scene across variation in details such as the objects' position, size and background context. Second, the brain must combine visual information with target-specific signals to determine whether a currently viewed scene contains a target. The neural mechanisms responsible for addressing both of these challenges are thought to be implemented, at least in part, in inferotemporal cortex (IT), but how target-specific signals combine with invariant object representations remains little-understood. To investigate, we recorded the simultaneous activity of moderately-sized neural populations in IT while a monkey performed a delayed-match-to-sample task that involved reporting when one of four objects was in view, across five different types of identity-preserving transformations. We then assessed how well a linear read-out could determine object identity, invariant of changes in position, size and background context, based on the spike count responses in a window that preceded the monkey's reaction times. We found higher population performance for each object on trials when it was a target as compared to trials when it was a distractor. Higher performance could be accounted for by an increase in the average firing rate response to targets relative to distractors and did not depend on correlated trial-by-trial variability within the population. Furthermore, just before a behavioral error, average firing rate responses were lower and similar to distractors when targets were not detected ('misses'), and responses were higher and similar to targets when distractors were misclassified ('false alarms'). These results establish that target-specific signals combine with invariant object representations in a manner that enhances the IT population representation of a target object, and these results suggest that IT population representations are likely used downstream to guide behavior.

III-53. Distinct neuronal responses in the human substantia nigra during reinforcement learning

Ashwin Ramayya¹
Kareem Zaghloul²
Bradley Lega¹
Christoph T Weidemann³
Michael Kahana¹
Gordon Baltuch¹

ASHWINRAMAYYA@GMAIL.COM
KAREEM.ZAGHLOUL@NIH.GOV
BRADLEGA@GMAIL.COM
CTW@COGSCI.INFO
KAHANA@PSYCH.UPENN.EDU
BALTUCHG@MAIL.MED.UPENN.EDU

¹University of Pennsylvania

²National Institutes of Health

³Swansea University

Midbrain dopaminergic (DA) neurons are thought to drive learning by reinforcing actions that lead to unexpected rewards. However, at least some DA neurons in the non-human primate midbrain respond to both positive and negative outcomes and thus may be involved in identifying salient events, rather than positive reinforcement. Because the primary feedback (e.g., airpuffs) used in these studies may activate DA neurons via low-level sensory representations, it is unclear whether this subset of DA neurons specifically encode sensory salience or more general motivational salience. Here, we show that some putative DA neurons in the human Substantia Nigra (SN) respond to positive and negative outcomes even when participants are learning from higher-order, abstract feedback. Taken together with recent work describing a population of human DA neurons that respond to unexpected virtual rewards, our findings suggest two motivational signals among putative DA neurons in the human SN: a positive valence signal that may drive reinforcement learning, and a salience signal that may detect behaviorally-

relevant events, regardless of their value. Furthermore, we describe a population of putative GABA-ergic neurons which respond shortly after DA neurons; these neurons may represent a local inhibitory mechanism to maintain a DA baseline during reinforcement learning.

III-54. Single neurons vs. population dynamics: What is changed through learning and extinction?

Anan Moran
Donald B Katz
Brandeis University

AMORAN@BRANDEIS.EDU
DBKATZ@BRANDEIS.EDU

Ample evidence suggests that although extinction of learning returns behavior to ‘baseline,’ this return represents not forgetting but rather learning anew. This evidence is ubiquitously indirect, however, as little is known about how extinction is reflected in neural responses. We set out to test whether single neurons return to pre-learning responses following extinction, and if not, how extinction-related changes might support a return of baseline behavior. We hypothesized that the answer might lie in the distinction between single-unit and population response dynamics. To test this hypothesis, we gave rats conditioned taste aversions (CTA), whereby they learned to dislike a sucrose solution that has been paired with nausea, and then extinguished that learning by presenting the sucrose in the absence of nausea, while recording small ensembles of single neurons from the gustatory cortex (GC), an area known to code taste information. As expected, the vast majority of single neurons held across learning and extinction (the entire procedure took 48 hours) did not return to pre-CTA responses. Using hidden Markov models (HMM, a technique that has successfully described taste coding in this area, see Jones et. al 2007), however, we observed several parameters of the population dynamics that reflected the return of ‘naïve-like’ behavior. Specifically, the span of states changed with learning and changed back with extinction, as did the percentage of single trials that could be described with single-state sequences (this phenomenon was 6-fold more likely after training than before learning or after extinction). Moreover, we found that as sucrose became aversive with training, sucrose trials were frequently misclassified as aversive acid, but that this tendency disappeared following extinction. Together these results show that although single neurons in cortex that change with CTA learning change (often more) with extinction, ‘taste processing’ as reflected in coherent population activity continues to subservise behavior.

III-55. Associative learning as an emergent property of spatially extended spiking neural circuits with STDP

Pulin Gong
John Palmer
University of Sydney

GONG@PHYSICS.USYD.EDU.AU
JPAL8929@PHYSICS.USYD.EDU.AU

Association of sequential events happening at different time moments is of fundamental importance for perceptual and cognitive functions. One of the important paradigmatic forms of such association is classical conditioning, in which the pairing of two subsequent stimuli is learned such that the presentation of the first stimulus (conditioned stimulus) is taken as a predictor of the second one (unconditioned stimulus). Most of the theoretical models proposed to account for classical conditioning have focused on individual neurons or synapses by assuming the presence of slowly decaying firing activity of neurons, which along with some synaptic plasticity such as spike timing dependent plasticity (STDP), enables associative learning between temporally separated events. However, the experimental evidence of such slowly decaying firing activity for associative learning is still inconclusive. Here, we present a novel, alternative account for the association based on the emergent properties of spiking neural circuits instead of individual neurons. Our proposal relies on two basic, known neurophysiological features of neuronal circuits: (1) lateral inhibitory coupling for neural circuits, and (2) spike timing dependent plasticity. In

a two-dimensional, spatially extended spiking neural circuit with these features, we find that each event can be represented by an emergent spiking sequence in terms of a propagating pattern in the network, and associative learning between the two stimuli can happen even when the timescale of their separation is significantly larger than that of individual neurons and synapses. Particularly, our network model can reproduce and then account for the contiguity of classical conditioning as found in behavioral studies, which states that the successful association rate is a non-monotonic function of time separation of conditioned and unconditioned stimuli. We find this emergent associative learning is quite robust to noise added to the network.

III-56. Self-tuning spike-timing dependent plasticity curves to simplify models and improve learning

Micah Richert
Botond Szatmary
Eugene Izhikevich
Brain Corporation

RICHERT@BRAINCORPORATION.COM
SZATMARY@BRAINCORPORATION.COM
EUGENE.IZHIKEVICH@BRAINCORPORATION.COM

Plasticity for feed-forward excitation ought to optimally assign credit to which synapses cause a postsynaptic cell to spike. It is common to use a double-exponential fit of the LTP and LTD curves (Bi and Poo, 1998); however, exponential curves are not always optimal and are prone to some pathologies. For example, if there are repeated patterns in the input spikes, learning will degenerate to only the selection of the earliest spikes (Masquelier et al., 2008). Often the parameters for STDP are hand tuned for particular problems and networks. Measuring of the cross-correlogram offline can provide useful insight as to what the optimal STDP curve should be. We propose an adaptive STDP curve that is derived online from the cross-correlogram, and will discuss its relationship to biology. This dynamic plasticity automatically incorporates an estimate of the dendritic and neuronal integration/processing time in order for a presynaptic input to cause a postsynaptic spike. This plasticity results in faster learning and greater diversity in a model of V1 simple cells since it more optimally accounts for which input spikes cause a post spike. For learning temporal patterns this plasticity does not shift to the earliest spikes of a repeated pattern. This enables a simple system to learn the whole temporal pattern instead of just the beginning. Further, for different neural models and input statistics, different STDP curves will be learned and yet still result in good V1 receptive fields. Because the STDP curve is adaptive to the statistics for each cell, it can be different for each cell in the same population. The model requires only a few meta parameters and these parameters are intuitive and learning is stable over a large range. Most importantly, instead of having to fiddle with parameters, this synapse model is self-tuning.

III-57. Formation and regulation of dynamic patterns in two-dimensional spiking neural circuits with STDP

John Palmer
Pulin Gong
University of Sydney

JPAL8929@PHYSICS.USYD.EDU.AU
GONG@PHYSICS.USYD.EDU.AU

Spike-timing-dependent plasticity (STDP) is an important synaptic dynamics that is capable of shaping the complex spatiotemporal activity of neural circuits. We examine the effects of STDP on the spatiotemporal patterns of a spatially extended, two-dimensional spiking neural circuit. We show that STDP can promote the formation of multiple, localized spiking wave patterns or multiple spike timing sequences in a broad parameter space of the neural circuit. Furthermore, we illustrate that the formation of these dynamic patterns is due to the interaction between the dynamics of ongoing patterns in the neural circuit and STDP. In particular, the subdiffusive motion of wave packets interacts with STDP to cause symmetry breaking which results in propagating wave fronts. This process is due to spontaneous symmetry breaking, which occurs in a fundamentally self-organizing manner, without fine-

tuning of the system parameters. Moreover, we find that STDP provides a synaptic mechanism to learn the paths taken by spiking waves and their behaviour during interactions, enabling them to be regulated. This regulation mechanism is error-correcting; in particular, it is able to correct for shifts in stimulus locations. If a stimulus, which creates a propagating wave is repeatedly applied at a fixed location and this location is then shifted, the new wave will quickly converge to the previously learned path. Our results, therefore, highlight the important roles played by STDP in facilitating the formation and regulation of spiking wave patterns that may have crucial functional roles in brain information processing.

III-58. Spinal presynaptic inhibition promotes smooth limb trajectories during reaching

Andrew Fink¹
Eiman Azim^{1,2}
Katherine Croce^{1,2}
Z. Josh Huang³
Larry Abbott¹
Thomas M. Jessell^{1,2}

AF2243@COLUMBIA.EDU
EA2471@COLUMBIA.EDU
KCROCE92@GMAIL.COM
HUANGJ@CSHL.EDU
LFA2103@COLUMBIA.EDU
TMJ1@COLUMBIA.EDU

¹Columbia University

²HHMI

³Cold Spring Harbor Laboratory

Sensory input to spinal motor neurons is subject to regulation via presynaptic inhibition, but little is understood about the role played by this form of inhibition during motor behavior. A specialized set of inhibitory interneurons, GABApre neurons, contact sensory afferent terminals and are hypothesized to mediate presynaptic inhibition. Here, we isolate GABApre neurons via their selective expression of Gad2 and ask (1) Does activation of the GABApre population evoke presynaptic inhibition at sensory-motor synapses? and (2) What are the consequences of acute ablation of GABApre neurons during motor behavior? We targeted GABApre neurons using a viral strategy, injecting Cre-dependent AAV into the spinal cord of Gad2::Cre mice. After injection of AAV encoding channelrhodopsin-YFP (ChR2-YFP) we detected ChR2-YFP expression in >70% of inhibitory synaptic contacts onto proprioceptive sensory terminals and found that ChR2-mediated photoactivation of GABApre neurons reduced neurotransmitter release probability at sensory-motor synapses. To examine the role of GABApre neurons during motor behavior we injected AAV encoding the diphtheria toxin receptor into the cervical spinal cord of Gad2::Cre mice. Interperoneal administration of diphtheria toxin following viral injection resulted in a loss of >80% of inhibitory contacts onto proprioceptive terminals. In mice trained to perform a forelimb reaching task, GABApre neuron ablation induced a marked deterioration in reach accuracy accompanied by the onset of pronounced reverberations in forepaw trajectory. To clarify the influence of presynaptic inhibition on limb extension we simulated a simplified joint with excitatory extensor drive, stretch-sensitive feedback, and a presynaptic inhibitory gate. Elimination of the inhibitory gate caused joint oscillation due to sensory-driven activation of antagonist muscles during joint extension. These results establish GABApre neurons as mediators of presynaptic inhibition at sensory-motor synapses and suggest that presynaptic inhibition serves to promote smooth trajectories during limb extension.

III-59. Integrative properties of motor cortex pyramidal cells during quiet wakefulness and movement

Paolo Puggioni
 Miha Pelko
 Mark van Rossum
 Ian Duguid
 University of Edinburgh

P.PUGGIONI@SMS.ED.AC.UK
 MIHA.PELKO@ED.AC.UK
 MVANROSS@INF.ED.AC.UK
 IAN.DUGUID@ED.AC.UK

The primary motor cortex (M1) plays a prominent role in the initiation and control of voluntary movements. Due to its direct link with behaviour, M1 is an ideal platform to study how brain state and behaviour are related to single neuron dynamics. We perform patch-clamp recordings and somatic current injections in the M1 of awake mice to characterise the intracellular activity and integrative properties of excitatory neurons in supercial (L2/3) and deep (L5B) layers during quiet wakefulness and movement. We find that during quiet wakefulness, L2/3 neurons display sparse spiking activity (0.5+0.7 Hz) while L5B cells display sustained firing (5.6+3.5 Hz) and that the membrane potential (Vm) in both cortical layers is characterized by slow fluctuations in the delta-band range (2-4 Hz). We identified two subpopulations of pyramidal cells in L5B -the main output layer of M1- that either suppressed (L5Bsupp) or enhanced (L5Benh) their firing rates during movement. In L5Bsupp neurons, movement decreased slow Vm oscillations and variance with no change in mean Vm, resulting in divisive gain modulation and reduced spike rates. In L5Benh neurons, movement also reduced slow Vm oscillations but this effect was counterbalanced by a net depolarization and increased Vm fluctuations in the high frequency band (12-50 Hz), resulting in increased firing rates. Based on integrate-and-fire simulations, we estimate that during movement L5Benh neurons preferentially receive an increase in excitatory inputs (%) with more substantial correlations on a fine time-scale. Together, these changes have a linear multiplicative effect on the input-output gain of L5Benh neurons. Our data demonstrate a remarkable diversity among cortical layers, a strong modulation of integrative properties depending on brain state and suggest that the cortex exploits behavior-dependent modes of operation.

III-60. Neural dynamics following optogenetic disruption of motor preparation

Daniel O'Shea¹
 Werapong Goo¹
 Paul Kalanithi¹
 Ilka Diester²
 Charu Ramakrishnan¹
 Karl Deisseroth¹
 Krishna Shenoy¹

DJOSHEA@STANFORD.EDU
 WGGO@STANFORD.EDU
 PKALANITHI@STANFORD.EDU
 ILKA.DIESTER@ESI-FRANKFURT.DE
 CHARUR@STANFORD.EDU
 DEISSERO@STANFORD.EDU
 SHENOY@STANFORD.EDU

¹Stanford University

²ESI, Max Planck, Frankfurt

The structure of neural activity in primate primary motor (M1) and dorsal premotor (PMd) cortex suggests that movement preparation drives cortical activity to a beneficial state for initiating movement, leading to quicker reaction times (RTs). Subthreshold electrical microstimulation in PMd largely erases this RT benefit (Churchland 2007), hypothesized to result from a time-consuming 'replanning' process. However, recent evidence demonstrates that movement may proceed without passing through this preparatory state (Ames 2012). Here, we asked whether disrupting preparatory activity instead slows RTs because post-stimulation neural activity follows a distinct peri-movement trajectory relative to non-stimulated trials. We employed optogenetic perturbation to probe the neural dynamics which underlie this behavioral disruption. We targeted the excitatory opsin C1V1TT to putative excitatory neurons in the arm region of PMd of two rhesus macaques trained on an instructed-delay reaching task. Optical stimulation (200ms continuous-pulse) in PMd delivered during the delay period within 100ms of the movement cue (late-stimulation) slowed RTs relative to non-stimulated trials, demonstrating that optical stimulation can disrupt movement preparation. To our knowledge, this represents the first optogenetically-mediated

behavioral effect in the primate arm-motor system. Stimulation delivered $>200\text{ms}$ before the movement cue (early-stimulation) did not affect RTs, suggesting that PMd can recover from optogenetic disruption on a rapid timescale. We computed the difference in PMd/M1 population firing rates on stimulated versus non-stimulated trials. Both early- and late-stimulation pushed neural state significantly away from non-stimulated trajectories. By movement onset, neural differences largely (though not completely) decayed for early stimulation trials, whereas these differences persisted into the movement time for late-stimulation trials. Therefore, targeted optogenetic perturbation of premotor cortical activity can disrupt motor preparation by diverting neural activity onto a distinct, presumably less-beneficial peri-movement trajectory.

III-61. Characterization of dynamical activity in motor cortex

Gamaleldin Elsayed¹
 Matthew Kaufman²
 Stephen Ryu³
 Krishna Shenoy³
 Mark M Churchland⁴
 John Cunningham¹

GAMALELDIN@WUSTL.EDU
 MKAUFMAN@CSHL.EDU
 SEOULMANMD@GMAIL.COM
 SHENOY@STANFORD.EDU
 MC3502@COLUMBIA.EDU
 CUNNINGHAM@WUSTL.EDU

¹Washington University in Saint Louis

²Cold Spring Harbor Laboratory

³Stanford University

⁴Columbia University

There has been increasing interest recently in understanding the role that internal dynamics play in the response of neural populations, in both sensory [1] and motor systems [2]. In particular, [2] has recently shown evidence of consistent, internally-driven dynamical activity in populations of motor cortical neurons by focusing on rotations in the neural state space. However, that work stopped short of exploring other dynamical features or characterizing their dominance. Here we characterize the structure of neural population dynamics by studying canonical features including: i) expansive vs. rotational structure (the two building blocks of simple linear systems); ii) time invariance vs. variance (to examine the temporal complexity of data); and iii) normality vs. nonnormality (a feature of theoretical importance due to its connection to Hebbian vs. balanced amplification [3]). Here we present algorithms that enable analysis of all these dynamical features. While time invariance vs. variance can be studied with simple least squares, the other features require novel contributions. Expansive systems are fit using new extensions to the method in [2], and fitting normal systems requires extensions to and combinations of classic results from differential geometry and linear algebra, which we derive below. We use these novel methods to analyze data from motor cortex. Despite the simplicity of the linear time invariant model, we show that activity across many different experimental conditions has consistently strong dynamics fitting this model (48.5% of the data variance is explained), and furthermore that the linear time invariant component is highly normal and rotational (93.2% and 91.3% of the time invariant system, respectively). In contrast, the time invariant system has trivial expansive component (8.7%). In all, this work deepens the characterization of dynamics in motor cortex and introduces analyses that can be used similarly across other cortical areas.

III-62. Decoding arm movements from hybrid spike-field potentials in human motor cortex

Janos Perge
 Brown University

JANOS_PERGE@BROWN.EDU

Action potentials recorded in motor cortex offer an information-rich input to motor prosthetic applications, yet recording stable spiking activity from the same neuronal population across months to years is an unresolved

technical challenge. Local Field Potentials (LFPs) — which can be recorded simultaneously with spikes — also contain information about intended movement while LFP recording stability and the inferred movement kinematics are assumed to be more robust. The degree to which LFP can encode movement information in humans and the stability of decoded signals remains to be rigorously tested. To evaluate kinematic information of LFPs in humans, we analyzed multiunit spiking activity and five frequency bands of the LFP signals (low =0.3-8Hz, beta=10-40Hz, gamma=45-65Hz, high1 =65-200Hz, high2 = 200-400Hz) in one person with tetraplegia using silicon multielectrode arrays chronically implanted in the arm area of M1. The participant, enrolled in the BrainGate2 (IDE) pilot clinical trial, was asked to try to move his dominant arm in four cardinal directions. In multiple sessions recorded over more than a year, the accuracy of offline target classification (using a naïve Bayesian classifier) in the highest frequency band was comparable to spikes, while other bands performed below spikes. Although the stability of classification results using LFPs did not exceed that obtained with spikes, LFPs contained complementary information, as 50% of the channels with no significant modulation in spikes were modulated in one of the five LFP bands (one-way ANOVA, $p < 0.01$, $n = 572$ channels over 12 sessions) or modulated in spikes but not in any of the LFP bands (4%, $n = 47$ channels). Due to this complementary information, classifying movement direction using both spikes and LFPs improved motor prosthetic performance and decreased day-to-day variance. Thus the reliability of future intracortical neuroprosthetic applications may be improved by using hybrid spike-LFP signals.

III-63. A model of hierarchical motor control and learning: from action sequence to muscular control

Ekaterina Abramova

Luke Dickens

Daniel Kuhn

Aldo Faisal

Imperial College London

E.ABRAMOVA10@IMPERIAL.AC.UK

LUKE.DICKENS@IMPERIAL.AC.UK

D.KUHN@IMPERIAL.AC.UK

ALDO.FAISAL@IMPERIAL.AC.UK

Fundamental question in neuroscience is how the brain translates and learns across levels of representation with ease (and optimality): e.g. at the behavioural level we often plan and specify tasks symbolically (e.g. "grab cup", "pour coffee"), while at the cellular level the brain continuously controls spike patterns for our 600 muscles. The neural computations that allow our brain to bridge across the divide between symbolic action selection and low-level actuation control, let alone learning at these two levels, are unclear. We propose a neural architecture which enables us to learn optimal control of a dynamical system from experience alone. Our hierarchical system can learn optimal control of non-linear dynamics - a challenging open problem in control theory. Our model — Reinforcement Learning Optimal Control (RLOC) - uses a top-level reinforcement learner (putatively located in the basal ganglia), which selects symbolic actions. Each action corresponds to low-level locally optimal linear feedback controllers (putatively implemented across M1, pMC and spinal cord). Our model uses low-level motor experience to learn the local system dynamics for local optimal linear control (as encountered in human reaching movements). The learning loop is closed by the high-level reward signal being driven by the low-level optimal control costs, which enables the system to learn the optimal sequence of local linear optimal controllers. Our model can learn, starting from unknown task dynamics, the non-linear optimal control of planar arm reaching movements and the pendulum-on-a-cart swing up and balance problem - the system learns quickly and finds solution rivaling or beating the performance of existing state-of-the-art non-linear control algorithms (which have full knowledge of the dynamics), while outperforming monolithic reinforcement learning approaches. Our model demonstrates how experimentally established findings in symbolic reinforcement learning and linear optimal control of movements can be combined to learn non-linear controls.

III-64. Gamma band activity in the human parahippocampal gyrus predicts performance in a sequence task

Radhika Madhavan^{1,2}
 Hanlin Tang³
 Daniel Millman¹
 Nathan Crone⁴
 Joseph Madsen²
 William Stan Anderson⁵
 Gabriel Kreiman¹

RADHIKA.MADHAVAN@CHILDRENS.HARVARD.EDU
 HTANG@FAS.HARVARD.EDU
 MILLMAN.DAN@GMAIL.COM
 NCRONE@JHMI.EDU
 JOSEPH.MADSEN@CHILDRENS.HARVARD.EDU
 WANDERS5@JHMI.EDU
 GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU

¹Harvard Medical School

²Childrens Hospital Boston

³Harvard University

⁴Johns Hopkins University

⁵JHH

A hallmark of episodic memory formation is the encoding and association of spatiotemporal sequences of events. En route towards a systematic characterization of the neural mechanisms underlying formation of complex episodic memories, here we investigated the neurophysiological responses underlying learning of temporal sequences of four images. After each sequence presentation, subjects were instructed to arrange the images in the order in which they appeared ('recall'). Subjects' performance increased over trials reaching >60% accuracy within 20 trials. We characterized neural responses from 549 electrodes during sequence learning at multiple temporal scales and different locations in the human brain using intracranial field potential recordings from 7 epileptic patients undergoing invasive monitoring for clinical purposes. We quantified oscillatory power at frequencies from 0.1-100Hz. Thirty-five out of 163 electrodes in the medial temporal lobe showed elevated levels of low gamma power (30-55Hz) during the recall period compared to pre-trial baseline. For these electrodes, we sought to quantify the relationship between their responses and the time course of learning over trials evaluated at the behavioral level. Strikingly, during the course of learning, peak power in the low gamma band decreased over trials. This decrease in gamma strength was directly correlated with the increase in performance accuracy. Learning-related modulation of gamma power was most prominent in electrodes over the parahippocampal gyrus (10/19 electrodes). In addition, in sessions wherein subjects did not learn (accuracy < 40%) gamma amplitudes remained unchanged over trials, further supporting the observed correlation between this physiological signature and learning. Several authors have suggested that oscillations in the gamma frequency band may help synchronize inputs with millisecond precision to facilitate learning. Based on those studies, we speculate that the decrease in gamma strength may reflect the need for plasticity during learning and relatively reduced levels of gamma to maintain learnt sequence traces.

III-65. Temporal windowing of stimulus processing in V1 by saccade-driven alpha oscillations

James McFarland¹
 Adrian Bondy²
 Bruce Cumming²
 Daniel A Butts¹

JMMCFARL@UMD.EDU
 BONDYAL@MAIL.NIH.GOV
 BGC@LSR.NEI.NIH.GOV
 DAB@UMD.EDU

¹University of Maryland

²National Institutes of Health

During free viewing, the visual input is temporally structured by eye movements, comprised of large jumps in eye position known as saccades that typically occur several times a second, interspersed with intervals where the gaze location is relatively constant. To determine how this temporal structuring of the visual input influences stimulus processing in primary visual cortex (V1), we recorded spiking activity and local field potentials (LFPs)

from a macaque during free viewing of natural images. To measure the neurons' tuning to features of the stimulus in this context, we used Gabor-filter 'energy' models, modified such that the degree of stimulus tuning ('response gain') was a function of time since fixation onset. We also fit similar models to the LFP power in different frequency bands to describe the stimulus-evoked network activity. We found that, in addition to evoking a large, transient response in the spiking activity and LFP, saccades entrained ~10 Hz alpha oscillations in the LFP that persisted throughout the subsequent fixation. By transforming to time coordinates of alpha cycles, we found that these alpha oscillations modulated the response gain of V1 neurons, resulting in a temporal windowing of the stimulus processing following saccades. The stimulus tuning of gamma (35-60 Hz) power was similarly modulated by the alpha rhythm, and the gamma oscillations provided further temporal structure to the spiking activity through spike-phase coupling. These results show that during free viewing, alpha oscillations following saccades create a temporal windowing of V1 activity during fixations, and more generally suggest that stimulus-driven network dynamics may play an important role in shaping feedforward stimulus processing.

III-66. Theta and gamma activity during human episodic memory retrieval.

John Burke
Michael Kahana
University of Pennsylvania

JOHN.FRED.BURKE@GMAIL.COM
KAHANA@PSYCH.UPENN.EDU

The ability to successfully retrieve stored memories is a fundamental process in episodic memory, but the neural substrate mediating memory retrieval is largely unknown. To investigate this issue, we conducted an experiment in which 60 neurosurgical patients with implanted subdural electrodes participated in a delayed free recall task. The location and timing of neural processes that support memory retrieval were examined by calculating instantaneous power fluctuations in theta and gamma frequency bands. We found that memory retrieval is initially marked by an increase in theta activity in the lateral temporal cortex and medial temporal lobe, which lateralized to the right versus the left hemisphere. Increases in gamma power followed this theta activation and were especially prominent in the left hippocampus. Together, these data suggest that both theta and gamma activity play a central role in episodic memory retrieval and can be used to track memories with high-temporal precision leading up to the point in which they are spontaneously vocalized.

III-67. A recurrent neural network that produces EMG from rhythmic dynamics

David Sussillo¹
Mark M Churchland²
Matthew Kaufman³
Krishna Shenoy¹

SUSSILLO@STANFORD.EDU
MC3502@COLUMBIA.EDU
MKAUFMAN@CSHL.EDU
SHENOY@STANFORD.EDU

¹Stanford University

²Columbia University

³Cold Spring Harbor Laboratory

It remains an open question how the firing rates of neurons in motor cortex (M1) lead to the EMG activity that ultimately drives movement. Recently, (Churchland et al., 2012)[1] reported that neural responses in monkey M1 exhibit a prominent quasi-rhythmic pattern during reaching, even though the reaches themselves are not rhythmic. They argued that M1 could be understood as 'an engine of movement that uses lawful dynamics', i.e., that M1 could be viewed as a dynamical system. A major question posed by their work is finding a concise set of equations for a dynamical system that uses rhythmic patterns to drive EMG. We approached this problem by training a nonlinear recurrent neural network (RNN) (Sussillo and Abbott, 2009) to generate the recorded EMG during the same reach tasks used in [1]. We trained the RNN to simultaneously generate the EMG activity recorded from three muscles for 27 'conditions' (reach types). The network was provided with condition-specific

static inputs as an initial condition, derived from the actual preparatory activity of recorded neurons. The RNN architecture consisted of a simulated M1 circuit, which provided input to three separate spinal cord circuits, one for each muscle. The model makes two main points. First, it is possible to produce realistic EMG activity using a network of this structure with the inputs provided. In particular, the input received during planning (derived from real neural data) provided a sufficiently detailed set of initial states to allow successful production of EMG for all 27 conditions. Second, the network naturally employs a solution that seems not unlike that used by the brain. This is true both on the surface level – simulated neurons appear very realistic, and on a mechanistic level, a large untuned component carries the neural state into a region of dynamics that produces strong oscillations.

III-68. Natural firing patterns reduce sensitivity of synaptic plasticity to spike-timing

Michael Graupner¹
Srdjan Ostojic^{2,3}

MICHAEL.GRAUPNER@NYU.EDU
SRDJAN.OSTOJIC@GMAIL.COM

¹New York University

²Group for Neural Theory, ENS, Paris

³École Normale Supérieure

Synaptic plasticity is sensitive to both the rate and the timing of pre- and postsynaptic spikes. In experimental protocols used to induce plasticity, the imposed spike trains are regular and the relative timing between every pre- and postsynaptic spike is fixed. This is at odds with natural firing patterns observed in the cortex of intact animals, where cells fire irregularly and the timing between pre- and post-synaptic spikes varies. To investigate synaptic changes elicited by in vivo-like irregularly firing neurons at different rates and realistic correlations between pre- and post-synaptic spikes, we use numerical simulations and mathematical analysis of synaptic plasticity models. We concentrate on a calcium-based model (Graupner and Brunel 2012), and further consider a voltage-based model (Clopath et al. 2010) and a spike-timing based model (Pfister and Gerstner 2006). To allow for comparison, all models are fitted to plasticity results obtained in vitro (Sjoestroem et al. 2001). We show that standard stimulation protocols overestimate the influence of spike-timing on synaptic plasticity. Using a simple modification of regular spike-pair protocols, we allow for neurons to fire irregularly. Such irregular spike-pairs reduce the amplitude of potentiation and depression obtained by varying the time difference between pre- and postsynaptic spikes. This protocol allows us to quantify the relative effects of firing rate and timing in natural firing patterns, and to predict changes induced by an arbitrary correlation function between pre- and post-synaptic spikes. We show that spike correlations change synaptic plasticity at low firing rates in all models; whereas their influence becomes negligible at high firing rates for the calcium-based model but remains significant for the other two models. Our findings yield predictions for novel experiments and help bridge the gap between existing results on synaptic plasticity and plasticity occurring under natural conditions.

III-69. Local edge statistics provides significant information regarding occlusions in natural scenes

Kedarnath Vilankar¹
James Golden¹
Damon Chandler²
David Field¹

KPV9@CORNELL.EDU
JRG265@CORNELL.EDU
DAMON.CHANDLER@OKSTATE.EDU
DJF3@CORNELL.EDU

¹Cornell University

²Oklahoma State University

This study investigated the statistical properties of edges found in natural scenes. We focused on the relative proportions and distribution of occlusions edges: edges due to the occlusion of one region of the image from another.

Since occlusion edges play an important role in object recognition, we investigated whether local information can be used to identify the probable edge type. Thirty-eight high-resolution natural scenes from the McGill Color Image Database were used in this analysis. In the first study, 1,000 edge locations were randomly selected from the image set using Canny edge detection. Eight subjects classified these edges into three types: 1) Occlusion edge, 2) Non-occlusion edge, and 3) Indeterminate. We used these labels to estimate the proportion of edges that were due to occlusions and the statistical differences between these edge types. We found 46% of the edges were occlusion edges, 48% were non-occlusion edges and 6% could not be classified. Further analyses determined that the different edge types differed based on local properties. Patches (41x81 pixels) were extracted at each location and then aligned with the local orientation of the edge. For patches that had 75% agreement in the classification task, we selected 356 occlusion and 347 non-occlusion edges. The most striking results were found with the distribution of Michelson contrasts. We found that 75% of the occluding edges had contrast values more than 0.55 and 75% of non-occluding edges had contrast values less than 0.2. An 80%-20% cross-validation revealed that nearly 91% of test edges can be classified correctly as occluding or non-occluding edges. We present further analyses using hand-labeled occlusion edges. Because local information can be used to provide significantly accurate edge classification, we argue that the early visual system has significant information regarding object boundaries in natural scenes.

III-70. Singular dimensions in spike triggered ensembles of correlated stimuli

Johnatan Aljadeff^{1,2}

Ronen Segev³

Michael J. Berry II⁴

Tatyana O. Sharpee^{2,1}

ALJADJEFF@UCSD.EDU

RONENSGV@EXCHANGE.BGU.AC.IL

BERRY@PRINCETON.EDU

SHARPEE@SALK.EDU

¹University of California, San Diego

²Salk Institute for Biological Studies

³Ben Gurion

⁴Princeton University

Many biological systems perform computations on inputs that have very large dimensionality. Characterizing what input combinations underlie the computation can go a long way towards understanding its function. Often, the relevant input dimensions are found as those along which the variance changes as a result of the computation, a method known as spike-triggered covariance (STC). This method has been highly successful in characterizing relevant input dimensions for neurons in a variety of sensory systems. So far, most studies used STC method with weakly correlated Gaussian inputs. Here we show that to be used with strongly correlated stimuli, such as 1/f signals typical of the natural environment, the STC method needs to be modified. The reason is that a random sample from ensemble of strongly correlated inputs has a covariance matrix with one (or more) outstanding eigenvalues. The presence of these "outstanding" modes interferes with analysis of statistical significance of candidate relevant input dimensions. In some cases, the outstanding mode can appear as the strongest mode with the largest change in variance, although this mode may have very little to do with the neural computation. Most commonly, the outstanding mode obscures many of the significant dimensions. We show that these issues are not resolved by removing correlations prior to doing the STC analysis because of the variability associated with estimated variance along different stimulus dimensions. However, evaluating significance of dimensions produced by the STC method in the subspace orthogonal to the outstanding mode does help to avoid these artifacts. Analyzing the responses of retinal ganglion cells probed with 1/f noise, we find that taking into account outstanding modes is crucial for recovering relevant input dimensions for these neurons.

III-71. The adaptive trade-off between discriminability and detectability in the vibrissa system

Douglas Ollerenshaw^{1,2}

He Zheng¹

Qi Wang¹

Garrett Stanley¹

D.OLLERENSHAW@GATECH.EDU

NEURON@GATECH.EDU

QI.WANG@BME.GATECH.EDU

GARRETT.STANLEY@BME.GATECH.EDU

¹Georgia Institute of Technology

²Emory University

The nervous system has long been known to adapt to the ongoing presence of a sensory stimulus, leading to important changes in the operating characteristics of the circuitry. Recent findings from our laboratory (Wang et al., 2010) demonstrated that from the perspective of an ideal observer of activity in the vibrissa cortex of the anesthetized rat, ongoing adaptive stimuli enhance discriminability of sensory inputs at the expense of detectability, which was shown to be mediated by thalamic synchrony. Here, we show that this generalizes to spatial cortical representations through voltage sensitive dye (VSD) imaging in the anesthetized rat and psychometric performance in the behaving rat. VSD measures of cortical activation show that both the magnitude and spatial spread of cortical activation decrease with adaptation, resulting in decreased detectability but increased discriminability in the adapted state. To test this in the behaving animal, we trained head-fixed rats in a single vibrissa detection task, followed by a two-whisker spatial discrimination task, both in the presence or absence of a preceding adapting stimulus. Animals in the detection task displayed a higher threshold for detection in the adapted state. In the discrimination task, animals showed an improved ability to discriminate between two whiskers with adaptation. In both tasks, an increased period of time between the adapting and test stimuli leads to a weaker effect, pointing to recovery from adaptation on the timescale of a few seconds. More recently, animals performing the detection task have been implanted with extracellular recording electrodes targeting the ventral-posteriomedial (VPM) nucleus of the thalamus, and preliminary results point to an adaptive reduction of both thalamic firing rate and timing precision, consistent with the process observed in the anesthetized animal. Taken together, results from this set of studies suggest a general principle of adaptive shaping of feature selectivity in sensory processing.

III-72. How humans and rats accumulate information in a tactile task, and a putative neuronal code

Arash Fassihi¹

Athena Akrami^{1,2}

Vahid Esmaeili¹

Mathew E Diamond¹

FASSIHI@SISSA.IT

ATHENA.AKRAMI@GMAIL.COM

VESMAEIL@SISSA.IT

DIAMOND@SISSA.IT

¹School for Advanced Studies (SISSA)

²Princeton University

We have devised a delayed comparison task, appropriate for human and rats, in which subjects discriminate between pairs of vibration delivered either to their whiskers, in rats, or fingertips, in human. To learn how signals are integrated over time, we varied the duration of the second stimulus. In rats, the performance was progressively improved when the comparison stimulus duration increased from 200 to 400 and then to 600 ms but in humans this increment led to a perceptual bias. Psychophysical reverse correlation revealed the perpetual strategies used by rats and humans in this task. Both rats and humans relied principally on features related to velocity, speed, and acceleration. For rats, the single feature that best predicted choice was velocity standard deviation. The difference is that rats judged the stimulus after normalizing by stimulus duration, whereas human subjects tended to be influenced by the summated values of those features over time. This explains why humans did not improve performance when comparison stimulus duration increased, but instead overestimated the value of the longer stimulus. Neuronal activity recorded from rat primary somatosensory cortex during the behavior revealed that 35% of neurons coded velocity standard deviation, the principal feature that affected choice. But the performance

predicted from individual neurons was far inferior to the actual performance of the subject. We employed a generalized linear model (GLM)-based analysis of neuronal population activity. The performance supported by the GLM was much closer to the animal's true performance, suggesting that the population code accounts better for this behavior.

III-73. Interplay of confidence and value in the transformation of olfaction-to-action

Gil Costa^{1,2}

Zachary F. Mainen¹

GIL.COSTA@NEURO.FCHAMPALIMAUD.ORG

ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

¹Champalimaud Neuroscience Programme

²PDBEB

Difficult decisions can occur because stimuli are hard to perceive or because the rules of what should be done given a certain stimulus are uncertain to the decision maker. We would like to understand how this second form of uncertainty is represented by the brain and may be assessed and used for adaptive behavior. In addition, we are interested in the relationship between uncertainty and value in such decisions. To study these issues we developed a perceptual decision task in which rats (n=6) were trained to perform a binary categorization of an odor mixture and report their decision confidence by reward waiting time. In order to dissociate neuronal representations of confidence and value we manipulated reward magnitude: in alternating blocks of trials one of the two choice sides was rewarded with 3 times more water than the opposite side. As expected, this reward manipulation biased the choice function of the animals towards the more rewarded side. In contrast, the confidence report was unaltered by the value manipulation: leaving times for the high and low reward ports were not different. This suggests the processes contributing to choice bias and confidence report in this task are not identical. We next performed simultaneous multi-electrode recordings in animals performing this task (n=6). We recorded from the olfactory tubercle (OT), considered to be the 'olfactory striatum'. OT single units (n=64) showed diverse firing properties correlated with variables including stimulus identity or contrast and action value, before the decision, and expected value of chosen actions during reward anticipation. We hypothesize that OT may be a locus where odor representations are linked to action values and therefore a site at which uncertainty in olfactory categorization may also originate.

III-74. Transformation of a temporal sequence to a spatial pattern of activity in piriform cortex.

Honi Sanders¹

Brian Kolterman²

Dmitry Rinberg³

Alexei Koulakov²

John Lisman¹

HONI@BRANDEIS.EDU

KOLTERMA@CSHL.EDU

DMITRY.RINBERG@NYUMC.ORG

AKULA@CSHL.EDU

LISMAN@BRANDEIS.EDU

¹Brandeis University

²Cold Spring Harbor Laboratory

³New York University

Mitral cells of the rodent olfactory bulb (OB) respond to stimuli with high frequency bursts of spikes (sharp events) at very precise phases of the sniff cycle (Shusterman, et al., 2011). Here, we show that sharp events occur at a preferred phase of the gamma oscillation measured in the field potential. Temporal precision and gamma phase preference of sharp events mean that the output of the OB is a discretized temporal sequence. Classification of the sharp event sequence is an example of the general problem of a network decoding a sequence of inputs lasting orders of magnitude longer than the time constants of its constituent neurons. Solution to this problem may have

applicability to a classic problem in perception, the ‘chunking’ processes that convert a temporal sequence to a spatial pattern. Here we propose a class of computational models that solve this problem by a brute-force method. The network is comprised of modules, each of which is dedicated to maintaining the OB activity of a particular gamma cycle. At the end of the sniff, each item in the sequence is represented by the activity of one module. We show how this can be achieved through two processes. 1) Activation of one module primes the next module so that it can be activated by the input that arrives in the next gamma cycle. 2) Once a module is maintaining activity, it must be impervious to further OB input. We successfully implement a version of this model using dendritic bistability as the working memory mechanism. We also show that our model network performs better with increasing gamma modulation of OB sharp events, demonstrating a possible function of gamma rhythms in information processing. Finally, we discuss possible network and biophysical mechanisms can underlie the needed computations.

III-75. Seeing turbulence - real-time visualization of turbulent odor plumes

Venkatesh Gopal¹
Robert Morton¹
Alexander Grabenhofer²

VGOPAL@ELMHURST.EDU
MORTONR@NET.ELMHURST.EDU
GRABENHOFERA825@NET.ELMHURST.EDU

¹Elmhurst College
²Northern Illinois University

We describe a novel schlieren flow visualization system developed to study the chemotactic behavior of *Drosophila Melanogaster* larvae that are exposed to a turbulent odor plume. Our system can image a quasi-two dimensional plume in real-time without the need to seed the flow with a visual marker such as smoke or Helium-filled soap bubbles. To our knowledge, this is the first behavioral assay in which the animal can be presented with a turbulent odor plume, and both the flow and the behavior of the animal tracking the flow can be imaged simultaneously, and in real-time. The temporal resolution of the system is limited solely by the camera frame-rate, thus making it possible to image behavioral responses to temporally patterned odor plumes at sub-millisecond timescales. Chemotactic behavior - orienting to, or tracking an odor gradient - is a vitally important component of the behavioral repertoire, as it allows animals to locate food and mates or orient away from the position of a predator. Chemotaxis and odor-following are complicated by the fact that the source of the odor is usually not a diffusing point source that can be found by ‘climbing up’ a smooth odor gradient. Instead, odors are most often transported by turbulent air or water currents. Turbulent odor plumes are amongst the most complex stimuli in nature. Spatially, turbulent mixing breaks up a smooth odor plume into ‘patchy’ discrete filaments. The odor concentration within these filaments has large, and essentially unpredictable, fluctuations. Thus, an animal navigating through a turbulent flow field experiences a rapidly fluctuating odorant concentration.

III-76. Representation of natural images in V4 using feature invariances

Yuval Benjamini¹
Julien Mairal²
Ben Willmore
Michael Oliver
Jack Gallant¹
Bin Yu¹

YUVALB@STAT.BERKELEY.EDU
JULIEN.MAIRAL@INRIA.FR
BENJAMIN.WILLMORE@DPAG.OX.AC.UK
MICHAELOLIVER@BERKELEY.EDU
GALLANT@BERKELEY.EDU
BINYU@BERKELEY.EDU

¹University of California, Berkeley
²INRIA Grenoble

Visual area V4 is believed to play an important role in the recognition of shapes and objects and in visual attention, but its function in shape representation is poorly understood. In particular, no known neuronal model of V4

provides good predictions of responses to natural images (Roe et al., Neuron 2012). In this work, we develop a model based on invariance and sparse coding principles that predicts well for V4 neurons and provides a rich description of their shape selectivity. We recorded electrophysiological data from 71 neurons in area V4 of two rhesus macaques, responding to sequences of natural images. The images are encoded according to a dictionary of features tuned on natural images; the dictionary elements represent complex shape properties such as corners, bars, curves and texture. We developed predictive models for each neuron, which were then validated on a separate set of images. Although V4 are more complex and diverse than their V1 counterparts, we achieve prediction accuracy levels comparable to previous studies of V1. The resulting models describe a diverse population of neurons, with many distinct selectivity profiles. Among these are neurons excited by corners, by thin bars, by specific orientations, or by texture; all with high prediction accuracy. We analyze the population of models using sparse principal component analysis, and discover two main groups of neurons: those selective to texture versus those selective to figures. This supports the hypothesis that one primary role of V4 is to extract image characteristics discriminating objects from background.

III-77. Tuning to low-level visual features is conserved during mental imagery

Thomas Naselaris¹
Cheryl Olman²
Dustin Stansbury³
Jack Gallant³
Kamil Ugurbil²

TNASELAR@MUSC.EDU
CAOLMAN@UMN.EDU
STAN_S_BURY@BERKELEY.EDU
GALLANT@BERKELEY.EDU
KAMIL@CMRR.UMN.EDU

¹Medical University of South Carolina

²University of Minnesota

³University of California, Berkeley

The relationship between mental and external images has been intensely debated for a millennium. Since the advent of fMRI much of the debate has focused on whether tuning to visual features of external images is conserved during mental imagery. For example, when we observe a unicorn are the cortical locations activated by the orientation of its horn also activated when we generate a mental image of the unicorn from memory? Here, we use a receptive field model of tuning to low-level visual features (i.e., retinotopic location, orientation, and spatial frequency) to identify complex mental images from ultrahigh-field fMRI measurements of human brain activity. This result indicates that tuning to low-level visual features of external images is at least partially conserved during mental imagery. Our finding paves the way for new technologies driven by visual mental imagery; we provide a proof-of-principle demonstration of brain-aided internet image search.

III-78. Optimal de-noising and predictive coding account for spatiotemporal receptive field of LGN neurons

Ziqiang Wei^{1,2}
Tao Hu^{3,4}
Dmitri Chklovskii^{3,4}

WEIZ@JANELIA.HHMI.ORG
HUT@JANELIA.HHMI.ORG
MITYA@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus, HHMI

²Johns Hopkins University

³Janelia Farm Research Campus

⁴HHMI

The efficient coding hypothesis suggests that the statistics of natural stimuli can account for spatio-temporal receptive fields (STRFs) of visual neurons (Attneave, 1954; Barlow, 1961). Yet, previous models (Atick & Redlich, 1990, 1992; van Hateren, 1992; Ruderman, 1994; Dong & Atick, 1995) do not yield a unique STRF prediction

without additional assumptions (Graham et al., 2006) because the analysis is performed in the Fourier domain and the transformation back to real space is ambiguous without knowing the phases of the Fourier components. Here, we propose an efficient coding theory in real space, in which a unique visual STRF is derived for a given natural scene statistics. In our model, the first processing step implements optimal linear de-noising of incoming stimuli. The second step performs predictive coding by subtracting an optimal linear prediction from the signal. With only one adjustable parameter (signal-to-noise ratio of the incoming stimuli), and using the same natural scene movie (Kayser et al. 2003), our model reproduces all the salient features of the STRFs measured in LGN neurons (Lesica et al., 2007): (1) the STRF center has expected biphasic shape in time; (2) the surround is weak and delayed relative to the center; (3) the shape of STRF changes with stimulus contrast as observed in experiment. Interestingly, our consistent implementation of predictive coding achieves de-correlation within and between spatial channels but only at different times. Thus, our results take the efficient coding hypothesis to its logical conclusion and resolve the STRF non-uniqueness found previously. Our STRF predictions can be tested experimentally by presenting stimuli with varying spatiotemporal correlations to the visual system.

III-79. Compensation of heading tuning for rotation in area VIP: Retinal and extra-retinal contributions

Adhira Sunkara^{1,2}
 Greg DeAngelis³
 Dora Angelaki²

SUNKARA@BCM.EDU
 GDEANGELIS@CVS.ROCHESTER.EDU
 ANGELAKI@CABERNET.CNS.BCM.EDU

¹Washington University in St. Louis

²Baylor College of Medicine

³University of Rochester

Moving through a 3D environment results in a projected vector field on the retina - optic flow. For linear translations, this pattern is radially symmetric around a 'Focus of Expansion' (FOE) corresponding to the heading direction. Eye movements introduce rotations in the optic flow which disrupt the FOE, obscuring the true heading. Psychophysical studies have shown that humans are capable of accurately perceiving their heading during pursuit eye movements. There are two strategies the brain can use to recover heading in the presence of rotation: (a) Subtract extra-retinal velocity signals from optic flow or (b) Use differences in optic flow properties of translation and rotation to separate the two components using retinal cues. To evaluate the contributions of each strategy, we introduce rotations to translational flow fields through eye pursuit (real pursuit, RP) or add rotation to the stimulus, simulating eye pursuit while the monkey fixates (simulated pursuit, SP). RP and SP produce the same retinal stimulation, but RP has additional extra-retinal signals. To evaluate extra-retinal and retinal contributions to pursuit compensation, we recorded extracellularly from ventral intraparietal area (VIP), which has heading direction tuning and eye pursuit responses. We compared the full tuning curve for translations in the horizontal plane (using a 3D dot cloud) to the tuning during either RP or SP. This comparison reveals slightly larger shifts in tuning for SP than RP, but much smaller than a purely extra-retinal strategy suggests, highlighting the importance of retinal cues in pursuit compensation. We also observe larger gain fields during RP compared to SP, suggesting the use of gain fields to signal the presence of eye rotations. The gain fields can potentially be useful for identifying whether the rotation is a result of eye movements or another source such as head movements, or moving on a curved path.

III-80. Emergence of bottom-up saliency in a spiking model of V1

Botond Szatmary¹

Micah Richert¹

Jayram Moorkanikara Nageswaran¹

Csaba Petre¹

Filip Piekniewski¹

Sach Sokol²

Eugene Izhikevich¹

¹Brain Corporation

²Johns Hopkins University

SZATMARY@BRAINCORPORATION.COM

RICHERT@BRAINCORPORATION.COM

NAGESWARAN@BRAINCORPORATION.COM

CSABA.PETRE@BRAINCORPORATION.COM

PIEKNIIEWSKI@BRAINCORPORATION.COM

SSOKOL3@JHU.EDU

EUGENE.IZHKEVICH@BRAINCORPORATION.COM

We present anatomically detailed spiking model of the parvo and magno pathways of the retina, primary visual cortex (V1), and superior colliculus (SC) to enable active saccades. Due to STDP and visual experience, the model shows the emergence of a saliency map, resulting in the perceptual behavior of bottom-up (pop-out) attention. In contrast to previous models proposed to explain pop-out based attention for simple features (e.g., saliency map hypothesis of Itti and Koch, 2001), where feature selectivity and inhibitory mechanisms between similar features are pre-wired, connectivity in our spiking model is neither pre-wired nor are neurons pre-labeled, but feature selectivity still emerges. Projections between cell types in the V1 model (L4 and L2/3) are in agreement with anatomical data. Both excitatory and inhibitory synapses are subject to different forms of STDP. These plasticity mechanisms and exposure to rich natural visual stimuli lead to (i) neuronal responses similar to those recorded in vivo, (ii - parvo) formation in color selective cells, and (iii - magno) formation of simple and complex cells covering a broad range of orientations and spatial frequencies. Pop-out mechanism is mediated by modifying the activity in layer 2/3 with long-range effective inhibition using a narrow form of STDP, which selectively picks up short temporal correlations between neurons responding to similar features but depresses ignores neurons responding to different features. Stronger within-feature long-range inhibition dampens the population response to features that are abundant in the input, but allows strong response to salient input features. The activity of V1 drives the SC, resulting in pop-out saccades. (The SC model is presented in a separate submission.) The model connects electrophysiology (spiking activity) and perception, and it explains animal behavior in a variety of standard pop-out tasks.

III-81. Gating of retinal information transmission by saccadic eye movements

Pablo Jadzinsky

Stephen Baccus

Stanford University

JADZ@STANFORD.EDU

BACCUS@STANFORD.EDU

The efficiency of information transmission in neurons has been a long standing focus of study. Classic results (Laughlin, 1981) indicate that a neuron could maximize information transmission by positioning its response curve to match the range of sensory input. However, most sensory neurons have a firing rate much lower than expected by virtue of a high threshold that reflects a different constraint, namely reducing the energetic cost of neural activity (Pitkow & Meister, 2012). Both considerations, however, assume stationary stimulus statistics, which are not true of natural sensory input. We examined information transmission in the salamander retina surrounding visual stimuli that simulated saccadic eye movements. Between 50-100 ms after a saccade, ganglion cells showed on average more than a two-fold average increase in transmitted information, which in many cases approached the theoretical maximal amount of information transmission given the Poisson-like noise of ganglion cell spiking. This increase in information was achieved by increased adaptation to the stimulus contrast. Subsequently, cells then reverted to a lower-rate energy conserving mode where only strong stimuli were encoded. These changes in the neural code were captured by a simple model whereby dynamic changes in peripheral presynaptic inhibition yielded a change in threshold that occurred prior to the stage of adaptation at the bipolar to ganglion cell synapse. We find that alternating between different representations of the same scene achieves two opposing goals. By

using a high information-rate code only very briefly after a saccade, the cell encodes a new visual scene quickly. In contrast, low firing rates are maintained when peripheral stimulus statistics would predict that central stimuli are less likely to change. Rather than choosing a single code given the tradeoff between energy efficiency and information transmission, a dynamic stimulus representation gains advantages and avoids drawbacks inherent to each neural code.

III-82. The role of the pulvinar in visual processing and attention

Ethan Meyers
Robert Schafer
Ying Zhang
Tomaso Poggio
Robert Desimone

Massachusetts Institute of Technology

EMEYERS@MIT.EDU
RJSCHAFFER@GMAIL.COM
YZHANG00@MIT.EDU
TP@AI.MIT.EDU
DESIMONE@MIT.EDU

The neural processing that underlies visual object recognition in primates occurs in the ventral visual pathway, and it is widely believed that this pathway operates in a feed-forward manner where information is passed from one cortical region to the next. It is also known that the pulvinar nucleus of the thalamus is interconnected with visual brain regions, and studies have suggested that all visual cortical regions that have direct anatomical connections also have a connection that passes through the pulvinar. However, it is still unclear what role the pulvinar plays in visual processing, or even how strongly the pulvinar's visual selectivity and attention modulation are relative to the visual response properties seen in cortex. To address these questions we recorded neurons in visual cortical areas V4 and the inferior temporal cortex (IT), and from the lateral pulvinar which is interconnected with these regions, while monkeys engaged in a visual attention task. Our results revealed several unexpected findings including that the population of neurons in the pulvinar was almost as visually selective as populations in V4 and IT, and that the pulvinar contains information that is position invariant. Additionally, we observed that the magnitude of attention effects in the pulvinar were similar to those seen in V4 and IT. From these findings it appears that the pulvinar could be part of a major visual pathway that passes information that has been traditionally associated with the ventral visual pathway between visual cortical regions.

III-83. Motion estimation involves high-order correlations differentially tuned for light and dark edges

Justin Ales
Damon Clark
James Fitzgerald
Daryl Gohl
Marion Sillies
Anthony Norcia
Thomas Clandinin
Stanford University

JUSTIN.ALES@GMAIL.COM
CLARKDA@STANFORD.EDU
JAMES.ELIOT.FITZGERALD@GMAIL.COM
DGOHL@STANFORD.EDU
MSILIES@STANFORD.EDU
AMNORCIA@STANFORD.EDU
TRC@STANFORD.EDU

Visual motion is a critical behavioral cue for many animals. Contemporary neural models estimate motion by computing pairwise space-time correlations in light intensity. Moving natural scenes, however, contain more complex correlational structures. By simulating motion using natural scenes we show that specific third-order correlations resulting from asymmetries in above- and below-mean regions of the visual scene contain useful information about motion. Moreover, motion estimation models that utilize odd-ordered correlations are able to distinguish between light and dark edges, something that 2nd order models cannot. Given that this information exists in moving natural images, we tested whether two very different visual systems, those of flies and humans,

extract third-order correlations to estimate motion in a manner that distinguishes light and dark edges. To isolate sensitivity to high-order correlations, we used 3-point 'glider' stimuli that contain no net 2-point correlations. These stimuli separate the motion information contained in 3rd and higher-order correlations from that specified by 2nd-order correlations. To isolate edge-polarity specific responses, we used novel stimuli that separately manipulated motion direction and edge polarity. Psychophysical measurements were obtained from both flies and humans as well as neural recordings in humans. Flies turn in the direction of motion allowing a behavioral measure of the fly's motion percept. We found that flies, like humans, detect motion in the 3-point 'gliders' and by using genetic manipulations, we show that this sensitivity is correlated with edge-polarity selectivity. Using Steady-State Visual Evoked Potentials we demonstrate that humans exhibit neural adaption that is specific to the combination of edge direction and edge polarity. Moreover, measured psychophysically, this adaptation in humans interacts with the perception of 3-point gliders, indicating that high-order correlations are differentially involved in edge polarity-selective motion processing in humans.

III-84. S cone isolating stimuli evoke express saccades in a chromatic contrast dependent manner

Nathan Hall
Carol Colby
University of Pittsburgh

NJH5@CNBC.CMU.EDU
CCOLBY@CNBC.CMU.EDU

Neurophysiological experiments in behaving macaque monkeys provide a unique opportunity to study sensory-motor transformation by relating sensory driven neural activity to motor output. We studied the superior colliculus (SC) as instance of this transformation because it contains neurons that respond to visual input and drive saccadic eye movement output. The circuit from retina through the SC is thought to enable very rapid 'express saccades' by omitting cortical processing. Our data suggest novel circuitry that can rapidly produce motor output in a stimulus-dependent manner. We probed this circuit by using stimuli that only activate short wavelength sensitive retinal cones (S cones). It has long been thought that SC neurons do not receive input from S cones. Here we show that the SC not only receives visual input from S cones but that it can very rapidly transform S cone input into motor output. We recorded behavioral data and SC neural activity while monkeys performed a visually triggered saccade task. The target was an individually calibrated stimulus defined by S cone contrast alone. We found that S cone stimuli can evoke express saccades. The latency of evoked express saccades varied with S cone contrast: higher contrast targets evoke faster express saccades. Neuronally, we discovered that the visual response of individual SC neurons had greater magnitude and shorter latency to high contrast S cone stimuli. The large visuomotor burst characteristic of express saccade trials also shifted in time as a function of contrast. Likewise, higher contrast targets influenced behavior by increasing the probability of express saccades while decreasing their latency. These results demonstrate a tight connection between visual input and oculomotor output within a single brain structure. However, the relationship between sensory input and behavioral output in the SC is not deterministic but rather affects express saccade generation in a probabilistic manner.

III-85. Motion-induced gain modulation in V1 improves contour detection

Torsten Lüdge¹
Robert Urbanczik²
Walter Senn²

TOERST@GMAIL.COM
URBANCZIK@PYL.UNIBE.CH
SENN@PYL.UNIBE.CH

¹Bern

²University of Bern

The recurrent interaction among orientation selective neurons in the primary visual cortex (V1) is well suited to enhance contours in a noisy visual scene. Motion selective neurons in V1 additionally support contour de- tectio

beyond the cooperative effect of the orientation selective neurons. In fact, slight motion of an object hidden in a still background can cause a strong pop-up effect. We ask how to wire a microcircuitry in V1 of orientation and motion selective neurons to explain the motion-induced pop-up effect exerted on contours. The circuitry has to cope with the observation that motion selective neurons themselves may show only weak orientation selectivity, and that motion may even be equally strong on the background than on the contour itself. This precludes a simple summation of orientation and motion induced evidence on the existence of local contours. We show that best performances of the contour detection network are achieved if the motion selective neurons locally enhance the gain of all orientation selective neurons at the spot of the motion. This local gain modulation makes use of the recurrent connectivity between the orientation selective neurons. Due to the local gain increase, the excitatory feedback loop among co-aligned neurons with the same orientation selectivity locally enhances the response to contours. Both, the locality of the modulation and its multiplicative form are crucial for the contour enhancement: a global gain modulation would unspecifically make the network over-excitable, and a local additive modulation would not fully exploit the power of a local self-excitation. The suggested local gain modulation of orientation selective neurons by motion selective neurons may readily be implemented by known elements of a cortical microcircuitry. Motion input to the apical tree of layer 5 pyramidal neurons may increase their gain to recurrently enhance their response to co-aligned oriented input.

III-86. Stimulus-dependence of membrane potential and spike count variability in V1 of behaving mice

Gergo Orban¹
 Pierre-Olivier Polack²
 Peyman Golshani²
 Mate Lengyel¹

GO223@CAM.AC.UK
 POLACKPO@HOTMAIL.COM
 PGOLSHANI@MEDNET.UCLA.EDU
 M.LENGYEL@ENG.CAM.AC.UK

¹University of Cambridge

²University of California, Los Angeles

Neuronal responses in the cortex show large trial-by-trial variability. This variability has profound consequences for cortical information coding and computations, yet its nature and the factors controlling it are still poorly understood. Importantly, in the primary visual cortex (V1), this variability has been shown to depend on the presence or absence of a stimulus – as quantified by the Fano factor of spike counts in extracellular recordings (Churchland et al, Nat Neurosci 2010). Intracellular recordings in anaesthetised animals indicate that changes in membrane potential fluctuations parallel those in spike count variability, found in extracellular recordings (Churchland et al, Nat Neurosci 2010). Moreover, contrast was also found to systematically affect membrane potential variability (Finn et al, Neuron 2007). However, neural activity under anaesthesia is characterised by large amplitude slow oscillatory patterns that can be measured both intra- and extracellularly and are quite unlike those seen in the awake animal (Steriade et al, J Neurosci, 1993). These correlated oscillations can thus potentially confound measured response variabilities. Therefore the relevance of stimulus-dependent changes in response variability in the awake state remains to be determined. We analyzed intracellular recordings from V1 neurons of awake mice in the absence of a stimulus as well as during stimulation with drifting sinusoidal gratings at different contrast levels. We found that the membrane potential fluctuations of V1 neurons showed a clear dependence on contrast and these changes were paralleled by changes in their firing activity as measured by the Fano factor. Furthermore, spontaneous activity (measured in the absence of a stimulus) is characterized by increased trial-by-trial variability compared to stimulus-evoked activity. These results suggest that models of the primary visual cortex need to address how these systematic changes in membrane potentials and spike counts emerge and what role they have in the computations carried out by the brain.

III-87. Fly photoreceptors are tuned to detect and enhance higher-order phase correlations

Daniel Coca¹
 Uwe Friederich¹
 Roger Hardie²
 Stephen Billings¹
 Mikko Juusola^{1,3}

D.COCA@SHEFFIELD.AC.UK
 U.FRIEDERICH@SHEFFIELD.AC.UK
 RCH14@CAM.AC.UK
 S.BILLINGS@SHEFFIELD.AC.UK
 M.JUUSOLA@SHEFFIELD.AC.UK

¹University of Sheffield

²University of Cambridge

³Beijing Normal University

In his seminal papers, Horace Barlow argued that sensory pathways must possess mechanisms for detecting, enhancing and encoding efficiently the stimuli that are behaviourally relevant for the animal. Here we used a photoreceptor model, which accurately predicts voltage responses of real photoreceptors to naturalistic stimuli under a wide luminance range, to test these hypotheses. In particular, we examined the significance of both local and non-local high-order phase correlations for animal vision. We simulated the photoreceptor model using synthetic stimuli sequences incorporating local phase correlations (edges) and non-local phase correlations (quadratic phase coupling) superimposed with Gaussian white noise. By decomposing voltage output of photoreceptor somata into linear second- and higher-order responses, we could elucidate three key encoding principles: the maximization of sensitivity to behaviourally relevant higher-order statistical features of the temporal stimuli, the selective improvement of signal/noise ratio and the efficient coding of these features. Using higher-order frequency response functions, derived analytically for the estimated model, we explain the nonlinear mechanisms responsible for coding the local and non-local higher-order statistical features in the stimuli as well as improving their signal-to-noise ratio. To validate the results, we carried out electrophysiological experiments using a specially designed stimuli sequence, which allows extracting the nonlinear component of the photoreceptor response directly from data, without a model. We show that the experimentally derived nonlinear components of the photoreceptor responses, elicited by the local and non-local phase correlations, agree very well with the responses predicted by the generic photoreceptor model. Finally, models derived for the somatic voltage responses of blind HDC-JK910 mutant photoreceptors strongly suggest that the nonlinear transformations, underlying the detection of high-order phase correlations in the temporal light patterns, are performed by their phototransduction alone, independently of neighbouring neurons.

III-88. A computational model of the role of eye-movements in object disambiguation

Lena Sherbakov¹
 Gennady Livitz¹
 Aisha Sohail¹
 Anatoli Gorchetchnikov²
 Ennio Mingolla³
 Massimiliano Versace¹

LENA.SHERBAKOV@GMAIL.COM
 GLIVITZ@GMAIL.COM
 AISHAS110@GMAIL.COM
 ANATOLI@BU.EDU
 E.MINGOLLA@NEU.EDU
 MAXVERSACE@GMAIL.COM

¹Boston University

²CompNet

³Northeastern University

The ability to explore and learn about one's surroundings is a critical capacity that facilitates everyday goal-directed behavior. Disambiguating the identity of objects that occupy a scene in the face of conflicting or incomplete evidence is an imperative sub-task of visual investigation. While research to date has focused on supervised or unsupervised classification algorithms for object recognition, little effort has been devoted to studying how small eye movements aid the visual system in object disambiguation. We present CogEye, a biologically-inspired,

active vision system that autonomously makes decisions about where to look, acts on those decisions in the form of saccades, learns view-invariant representations of salient objects, and builds mental spatial maps of its surroundings. A central feature of CogEye is the what-to-where visual stream connection that, in conjunction with bottom-up saliency, influences saccade selection. The model continuously learns prototypical size-invariant 'object-feature' maps and uses them to bias a saccade toward the location of greatest object disambiguation power. In addition to providing a testable hypothesis for the function of eye movements in object recognition, the novelty of our approach resides in the GPU-based computational platform. CogEye is built on Cog Ex Machina (Snider et al. 2011), a high performance computing platform designed for simulating large-scale integrative brain systems. We show that our model agrees with human saccade psychophysics during object disambiguation tasks, and we highlight the idea that within-object saccades can increase the efficiency and accuracy of object recognition.

III-89. How biased are maximum entropy models of neural population activity?

Jakob H Macke^{1,2}
Iain Murray³
Peter Latham⁴

JAKOB.MACKE@GMAIL.COM
I.MURRAY@ED.AC.UK
PEL@GATSBY.UCL.AC.UK

¹Max Planck Institute Tübingen

²Bernstein Center Tübingen

³School for Informatics, Edinburgh University

⁴University College, London

Maximum entropy models have become popular statistical models in neuroscience and other areas of biology, and can be useful for quantifying the coding properties of sensory systems. However, maximum entropy models fit to small data sets can be subject to sampling bias; i.e. the true entropy of the system can be severely underestimated. Here we study the sampling properties of estimates of the entropy obtained from maximum entropy models. We focus on the pairwise binary model, which is used extensively to model neural population activity. We show that if the data is well described by a pairwise model, the bias is equal to the number of parameters divided by twice the number of observations. However, if the higher order correlations in the data deviate from those predicted by the model, the bias can be larger. Using a phenomenological model of neural population recordings, we find that the additional bias due to higher-order correlations is largest for small firing probabilities, high correlations, and large population sizes. However, numerical studies indicate that even in the worst case, it is only about a factor of four larger. We derive guidelines for how much recording time one needs to achieve a bias which is smaller than some specified level of accuracy. Finally, we show how a modified plug-in estimate of the entropy can be used for bias correction.

III-90. Unsupervised learning of latent spiking representations

Ryan P Adams¹
Geoffrey Hinton²
Richard Zemel²

RPA@SEAS.HARVARD.EDU
HINTON@CS.TORONTO.EDU
ZEMEL@CS.TORONTO.EDU

¹Harvard University

²University of Toronto

Connectionist machine learning originates in the idea of designing adaptive computational frameworks which resemble neural systems. One of the most salient features of real neural systems, however, is that they most often communicate via action potentials. It has been difficult to incorporate spiking behavior into connectionist learning frameworks, as our tools for learning from point processes are much more limited than those for learning

from vectorized representations. In this work, we examine how at the limit of the restricted Boltzmann machine (RBM), a popular class of undirected graphical models for unsupervised feature discovery, naturally yields a latent point process representation. We discuss various issues surrounding this representation, and show how it easily allows topological constraints to be enforced in spike-based hidden representations. We also examine how to build learning systems that incorporate neurally-motivated phenomena such as inter-neuron competition. Finally, we demonstrate that this framework can be practically implemented to learn from real data.

III-91. Internally generated transitions in multi-state recurrent neural networks

Sean Escola
Larry Abbott

SEAN@NEUROTHEORY.COLUMBIA.EDU
LFA2103@COLUMBIA.EDU

Columbia University

Neural networks appear to have the ability to operate in multiple states, each with unique dynamics. For example, networks in primate motor cortex during reaching tasks exhibit very different dynamics during the delay period versus after movement onset (Churchland et al., Nature, 2012). Recently, a hidden Markov model technique was developed to infer states defined by state-dependent dynamics. When applied to data, this method uncovers unexpected dynamics such as oscillations that are phase-locked to transition times (Escola and Paninski, Cosyne, 2010; Escola et al., Neural Computation, 2011). To better understand the mechanisms by which neural networks can maintain multiple dynamical regimes and make appropriate transitions between them, we use FORCE learning (Sussillo and Abbott, Neuron, 2009) to train the linear readout unit of a random recurrent neural network to produce one of a set of distinct trajectories depending on which constant input the network is receiving. By visualizing these trajectories with PCA, it is clear that each state-dependent trajectory lies in a distinct region of firing rate space. We are also able to train the network to have near-instantaneous transitions between states when the constant input changes, with transition times approaching the neuronal time constant (10 ms). By adding readout units that detect when the network completes each of the state-dependent trajectories, we can use these outputs to drive an external fixed-point network to jump between attractors. By using this fixed-point network to generate the constant inputs into the primary network, we are able to close the loop and have the network activity itself drive transitions between states, unlike prior models of state-dependent computations (notably, Pascanu and Jaeger, Neural Networks, 2011) that relied on an external signal to effect transitions. Furthermore, by adding noise at the level of the fixed-point network, the system begins to resemble an arbitrary Markov chain.

III-92. Scalability properties of multimodular networks with dynamic gating

Daniel Marti
Omri Barak
Mattia Rigotti
Stefano Fusi

DM2913@COLUMBIA.EDU
OMRI.BARAK@GMAIL.COM
MR2666@COLUMBIA.EDU
SF2237@COLUMBIA.EDU

Columbia University

Brain processes arise from the interaction of a vast number of elements. Despite the enormous number of involved elements, interactions are generally limited by physical constraints. Typically a neuron is connected to thousands other neurons, a far lower number than the hundred billion neurons in the brain. Unfortunately, it is the number of connections per neuron, not the total number of neurons, what often determines the performance of large neural networks (measured, e.g., as memory capacity), a fact that hinders the scalability of such systems. We hypothesize that the scalability problem can be circumvented by using multimodular architectures, in which individual modules composed of local, densely connected recurrent networks interact with one another through sparse connections. We propose a general model of multimodular attractor neural networks in which each module state changes only upon external event and the change depends on the state of a few other modules. To

implement this scheme, every module has to disregard the state of any module not involved in a particular interaction. Because a module can potentially interact with several others, ignoring the states of non-relevant modules would require learning of an exponentially large number of conditions. We solve this problem by adding a group of neurons that dynamically gate the interactions between modules. These neurons receive inputs from the modules and event signals through random sparse connections, and respond to combinations of event-states. This information is then sent back to the modules. Because they implement conjunctive representations, the number of necessary gating neurons grows only polynomially with the number of modules. We hypothesize that gating neurons reside in cortical layer 2/3, and that they mediate the interactions between modules in layer 5/6. The laminar organization of the neocortex could thus be a crucial architectural solution to the scalability problem.

III-93. A mathematical theory of semantic development

Andrew Saxe
James L. McClelland
Surya Ganguli
Stanford University

ASAXE@STANFORD.EDU
MCCLELLAND@STANFORD.EDU
SGANGULI@STANFORD.EDU

Many psychology experiments have revealed remarkable regularities in the developmental time course of semantic cognition in childhood, as well as its progressive disintegration in adult dementia. For example, infants tend to learn to make broad categorical distinctions between concepts (i.e. plant/animal) before they can make finer scale distinctions (i.e. dog/cat), and this process reverses in dementia, where finer scale distinctions are lost before broad ones. What are the theoretical principles underlying the ability of neuronal networks to discover categorical structure from experience? We address this question by developing a phenomenological, mathematical theory of semantic development through an analysis of the learning dynamics of multilayer networks exposed to hierarchically structured data. We find new exact solutions to the nonlinear dynamics of error corrective learning in deep, 3 layer networks. These solutions reveal that networks learn input-output covariation structure on a time scale that is inversely proportional to its statistical strength. We further analyze the covariance structure of data sampled from hierarchical generative models, and show how such models yield a hierarchy of input-output modes of differing statistical strength, leading to a hierarchy of time-scales over which such modes are learned. Our results reveal, quite strikingly, that even the second order statistics of hierarchically structured data contain powerful statistical signals sufficient to drive many complex experimentally observed phenomena in semantic development, including the progressive differentiation of concepts and its reversal in dementia, sudden stage-like transitions in performance, and erroneous illusory correlations in early learning. Moreover, our work reveals how deep network structure is essential for exhibiting these phenomena. Beyond semantic cognition, our analytical results provide an extremely general formalism for understanding how the statistical structure of experience drives learning in deep networks, and could provide insight into the learning dynamics of many different network models in psychology and neuroscience.

III-94. Disentangling serial adaptation using a biophysically constrained model

Bongsoo Suh
Stephen Baccus
Stanford University

BSSUH@STANFORD.EDU
BACCUS@STANFORD.EDU

In the brain, sensory information is transformed across successive layers of spatiotemporal filtering, thresholds and nonlinear adaptive processing. Because neural recordings of typically have no access to transformations occurring at lower levels, the presence of such intermediate processing poses a barrier to understanding both the sensory neural code and its biophysical implementation. Retinal ganglion cells adapt to temporal contrast in at least two stages, with separate components of adaptation in the membrane potential and spiking response

(Zaghoul et al. 2005). Recently we reported that a simple biophysical model of synaptic vesicle cycling captures adaptive processing seen in the salamander ganglion cell membrane potential (Ozuyal and Baccus, 2012). This LNK model consists of a linear temporal filter, a threshold nonlinearity and a first-order kinetic model of the type used to capture chemical reactions. Here we endeavor to capture both stages of adaptation using only the spiking output. We extended the LNK model with a spiking stage containing negative feedback, yielding a Linear-Nonlinear-Kinetic-Spiking (LNKS) model. To validate this approach we fit LNKS model parameters to intracellular recordings of ganglion cells responding to changing temporal contrast where both membrane potential and spikes were recorded. Using only spikes, we optimized the model using a constrained gradient descent method. With a single set of model parameters, the overall correlation coefficient of the firing rate between the data and model was 85%, and 80% for the membrane potential that was never compared while fitting the model. The intrinsic variability between repeats of the same stimulus was 88% for the firing rate and 89 % for the membrane potential. This approach reveals that by imposing biophysically reasonable constraints on different model stages, we can discover parameters for successive stages of adaptive processing without the need for explicit recording at different levels.

III-95. Efficient associative memory storage in cortical circuits of inhibitory and excitatory neurons

Julio Chapeton
Armen Stepanyants
Northeastern University

CHAPETON.J@HUSKY.NEU.EDU
A.STEPANYANTS@NEU.EDU

Many features of synaptic connectivity are ubiquitous among cortical systems. Cortical networks are dominated by excitatory neurons and synapses, are sparsely connected, and function with stereotypically distributed connection weights. We show that these basic structural and functional features of synaptic connectivity arise readily from the requirement of efficient associative memory storage. Our work builds on the study of Brunel and colleagues, which demonstrated that a cerebellar Purkinje cell receiving excitatory inputs from parallel fibers is analogous to a perceptron operating at its maximum associative memory storage capacity. Motivated by this, along with recent observations of inhibitory synaptic plasticity, we developed a new theory on memory storage which can be applied to networks of not only excitatory but also inhibitory neurons. Our theoretical results indicate that efficient memory storage leads to network connectivity with predictable statistics that can be measured experimentally. First, we predict that despite the large number of neuron classes, functional connections between potentially connected cells must be realized with $<50\%$ probability if the presynaptic cell is excitatory and $>50\%$ probability if it is inhibitory. Second, we establish a unique relation between probability of connection and coefficient of variation in connection weights. These predictions do not depend on any parameters and must hold for all circuits involved in efficient associative memory storage. We also show that the addition of a small fraction of inhibitory neurons into and all-excitatory network increases the network's capacity for robust associative memory storage. What is more, our theory explains the shapes of IPSP/EPSP distributions obtained in electrophysiological experiments. These predictions are consistent with a dataset of 74 published experiments reporting connection probabilities and distributions of IPSP/EPSP amplitudes in various cortical systems.

III-96. Probabilistic inference reveals synapse-specific short-term plasticity in neocortical microcircuits

Rui P. Costa¹
P. Jesper Sjöström²
Mark van Rossum¹

RUI.COSTA@ED.AC.UK
JESPER.SJOSTROM@MCGILL.CA
MVANROSS@INF.ED.AC.UK

¹University of Edinburgh
²McGill University

Short-term synaptic plasticity is highly diverse and varies with brain area, cortical layer, cell type, and developmental stage. Since this form of plasticity shapes neural dynamics, its diversity suggests a specific and essential role in neural information processing. Therefore, a correct identification of short-term plasticity is an important step towards understanding and modeling neural systems. Accurate phenomenological models have been developed, but they are usually fitted to experimental data using least-mean square methods. We demonstrate that, for typical synaptic dynamics, such fitting gives unreliable results. Instead, we introduce a Bayesian approach based on a Markov Chain Monte Carlo method, which provides the posterior distribution over the parameters of the model. We test the approach on simulated data. First we show that common protocols to measure short-term plasticity protocols yield broad distributions over some model parameters, i.e. with inaccurate estimates. Using this insight, we find a better experimental protocol for inferring the true synaptic parameters and show that our Bayesian formulation provides robust identification of changes in the model parameters. Next, we infer the model parameters using experimental data from three different neocortical excitatory connection types, revealing novel synapse-specific distributions, while the approach yields more robust clustering results. Our approach to demarcate synapse-specific synaptic dynamics is an important improvement on the state of the art and reveals novel features from existing data as well as guiding future experimental work.

III-97. A spiking input-output relation for general biophysical neuron models explains observed $1/f$ response

Daniel Soudry¹
Ron Meir²

DANIEL.SOUDRY@GMAIL.COM
RMEIR@EE.TECHNION.AC.IL

¹Electrical Engineering, Technion

²Technion

Cortical neurons contain many sub-cellular processes, operating at multiple timescales, which, through non-linear and stochastic interactions, may entail a highly complex Input-Output (I/O) relation. A glimpse at such a relation was observed in recent experiments, in which synaptically isolated individual neurons, from rat cortical culture, were stimulated with periodic extra-cellular current pulses. The neurons exhibited a ' $1/f^\alpha$ spectrum', responding in a complex and irregular manner over the entire range of experimental timescales - from seconds to days. The large number of neuronal processes present within these timescales (mostly unknown) precludes a purely simulation-based investigation which requires a precise specification of all parameters. To circumvent this problem, we derive the neuronal spiking I/O for a broad class of biophysical neuron models under sparse spiking input, relating input spike trains to output spikes based on known biophysical properties, given a few sufficient assumptions. Thus we obtain closed-form expressions for the mean firing rates, all second order statistics and construct optimal linear estimators for the neuronal response and internal state. We numerically verify that these results continue to hold even in cases where our assumptions fail to hold. We then relate to the experimental results and find that, given these assumptions, the neuron must contain processes over a wide range of timescales, where the number of relevant ion channels must also scales with exponent α . The mathematical analysis shows that very slow processes indeed affect the mean and fluctuations of the response, but this does not imply that the neuron retains a very long memory trace of its inputs - and we discuss how this can directly tested.

III-98. Storing structured sparse memories in a large-scale multi-modular cortical network model

Alexis Dubreuil¹
Nicolas Brunel²

ALEXIS.DUBREUIL@ENS-CACHAN.FR
NBRUNEL@GALTON.UCHICAGO.EDU

¹UMR 8118, CNRS & Universite Paris Descartes

²University of Chicago

Different areas of the cortex are linked to each other thanks to long range connections that travel through white matter. It has been proposed that the resulting network, made of interconnected local modules, could be used as a global auto-associative memory device, storing multi-modal memories in the form of dynamical fixed-point attractors [1]. However, it remains unclear which types of large-scale connectivity matrices are well suited to perform such a function. To investigate this question, we studied network models of binary neurons, made of local highly connected modules that communicate with each other through diluted long range connections. Once the global form of the synaptic matrix is chosen, the network is loaded with patterns that are imprinted using an Hebbian-type learning rule. We characterized the ability of different synaptic structures and different pattern distributions to give rise to an auto-associative memory network. This was done by computing quantities such as the maximal capacity of the network, or the relationship between the amount of information stored in the network and the typical size of the basin of attraction associated with each pattern. We found synaptic structures that allow to store global memory patterns in an efficient way, i.e. they store a finite amount of information per synapse, with a finite basin of attraction. For memories to have a finite size basin of attraction, the number of long range connections should be above a certain threshold that we characterized. The effect of the coding level, both at the microscopic and at the macroscopic scale, on memory properties of the network was also investigated. [1] O’Kane D., Treves A. ‘Short-and Long-range Connections in Autoassociative Memory.’ Journal of Physics A: Mathematical and General 25 (1992): 5055.

III-99. Conditional random fields for spiking populations of neurons

Emily Denton
Richard Zemel
University of Toronto

DENTON@CS.TORONTO.EDU
ZEMEL@CS.TORONTO.EDU

With the advances of multi-electrode recording technologies comes an increased need for techniques that can effectively model the relationship between complex sensory stimuli and the multivariate, correlated responses of a neural population. To address this problem, we propose a new framework for modelling spiking responses in correlated populations of neurons which offers several advantages. First, we utilize conditional random fields as a modelling framework, providing a novel approach to this problem. A standard CRF could capture stimulus dependent response properties and interactions between neurons. We extend this framework by incorporating higher order potentials which allow for a natural characterization of global response properties. An additional novel aspect of the model is that the generated spikes correspond to the MAP estimate, offering an alternative view of neural computation as maximizing probabilistic quantities. We apply this framework to model time-varying stimulus data and corresponding spiking population responses, achieving promising preliminary results.

III-100. The influence of network structure on neuronal dynamics

Patrick Campbell¹
Michael Buice²
Duane Nykamp¹

PRCAMP@MATH.UMN.EDU
MABUICE@GMAIL.COM
NYKAMP@MATH.UMN.EDU

¹University of Minnesota

²The Allen Institute for Brain Science

Understanding the influence of network structure on neural dynamics is a fundamental step toward deciphering brain function, yet presents many challenges. We show how networks may be described in terms of the occurrences of certain patterns of edges, and how the frequency of these motifs impacts global dynamics. Through analysis and simulation of neuronal networks, we have found that two edge directed paths (two-chains) have the most dramatic effect on dynamics. Our analytic results are based on equations for mean population activity and correlations that we derive using path integrals and moment hierarchy expansions. These equations reveal

the primary ways in which the network motifs influence dynamics. For example, the equations indicate that the propensity of a network to globally synchronize increases with the prevalence of two-chains, and we verify this result with network simulations. Finally, we present ongoing work investigating when these second-order equations break down, and how they might be corrected by higher order approximations to the network structure, such as the prevalence of three edge chains beyond that predicted by the two-chains.

Author Index

- Abarbanel H., 183
Abbott L., 42, 69, 80, 152, 153, 189, 208
Abrahamyan A., 167
Abramova E., 192
Acuna D., 34
Adams G., 110
Adams R. P., 157, 162, 207
Adibi M., 141
Aertsen A., 177
Agarwal G., 71
Ahmadian Y., 139
Ahrens M. B., 44
Aihara K., 179
Aimone J., 77
Aitchison L., 116
Akam T., 131
Akrami A., 46, 197
Ales J., 203
Ali F., 44
Aljadef J., 196
Alonso J., 117, 143
Alstott J., 83
Alvina K., 42
Andersen R., 33
Anderson W. S., 139, 193
Andrei A., 96
Angelaki D., 36, 56, 140, 201
Angueyra J., 40
Anselmi F., 75
Aoi M., 62
Aonishi T., 127
Arabzadeh E., 141
Archer E., 47, 154
Ardid S., 56
Arnstein D., 153
Averbeck B., 134
Avery M., 46
Avitan L., 151
Avni O., 124
Awatramani G., 96
Azim E., 189
- Ba D., 156
Babadi B., 125
Baccus S., 202, 209
Bacher D., 133
Baden T., 144
Badreldin I., 118
Balasubramanian K., 118
Balasubramanian V., 72, 84, 96, 106, 127
Balcarras M., 56
Ballard D., 82
Baltuch G., 186
Banerjee A., 69
Bansal A., 53
Bao M., 85
Barack D., 170
Barak O., 208
Barbieri F., 61
Barrett D. G., 33, 98
Bathellier B., 74
Battaglia D., 138
Beck J., 49, 108, 142
Beierholm U., 50
Benjamin B., 103
Benjamini Y., 199
Benna M., 102
Bennett D., 147
Berens P., 143, 144
Bergmann U., 121
Berkes P., 55, 76, 172
Bernacchia A., 156, 171
Berniker M., 34
Berry II M. J., 123, 196
Bethge M., 50, 143, 144, 153
Bhardwaj M., 113
Bharioke A., 66
Bialek W., 23
Billings S., 206
Blazeski R., 67
Boahen K., 23, 103, 150
Bonci A., 59
Bondy A., 193
Born R., 91
Bornschein J., 172
Bouchard K., 111, 135
Boucsein C., 177
Bourdoukan R., 33
Bozza T., 32
Brainard M., 111
Brea J., 99
Breton Y., 160
Briguglio J., 84

- Brody C., 24, 110
Brown E., 137
Brunel N., 61, 176, 211
Bruno R., 27, 65, 67, 120
Brunton B., 110
Buesing L., 151, 162
Buibas M., 145
Buice M., 212
Bullmore E., 83
Burge J., 159
Burke J., 194
Burkitt A. N., 97
Burns S., 139
Butts D. A., 176, 193
Buzsaki G., 71
Böhmer W., 114
Büchel C., 114
- Cadiou C., 36
Caetano M., 53
Cai M., 87
Calabrese A., 51
Calhoun A., 78
Campbell P., 212
Carcea I., 58
Cardin J., 119
Carnevale F., 32
Carroll S., 106
Carver F., 83
Cash S., 119, 133
Cavanagh J., 53, 165
Cayco-Gajic N. A., 107
Cetin A., 180
Chacko R., 181
Chagas A. M., 153
Chaisanguanthum K., 135
Chalasanani S., 78
Chalk M., 115
Chandler D., 195
Chang E., 135
Chapeton J., 210
Chaudhuri R., 156
Chechik G., 80
Chemali J., 137
Chiu C., 76
Chklovskii D., 49, 66, 92, 200
Christopoulos V., 54
Churchland A., 112, 115
Churchland M. M., 153, 191, 194
Clandinin T., 30, 203
Clark D., 30, 203
Clifford C. W., 141
Clopath C., 120
Coca D., 206
Coen-Cagli R., 39
- Colby C., 204
Collins A., 58
Colonnese M., 64
Cone J., 132
Conover K., 160
Constantinople C., 27
Conte M. M., 84
Cools R., 133
Coppola D. M., 125, 148
Coppola R., 83
Cormack L., 52, 91
Corrado G. S., 166
Costa G., 198
Costa R., 131
Costa R. P., 210
Costa V., 134
Costello G., 38
Cotton R. J., 93, 143
Cowell R., 174
Crevecoeur F., 101
Croce K., 189
Crone N., 193
Cui Y., 176
Cumming B., 193
Cunningham J., 191
Curto C., 99, 155
Czuba T. B., 91
- Dadarlat M., 88
Das D., 50
Daw N., 28
Day M., 183
Dayan P., 51, 131, 160
De Franceschi G., 185
de la Rocha J., 41
de Lafuente V., 32
de Ruyter van Steveninck R., 60, 140
De Schutter E., 82
DeAngelis G., 56, 201
Dehghani N., 119
Deisseroth K., 190
Delgutte B., 183
Demb J., 176
Deneve S., 33, 98, 115, 117
Denton E., 212
DePasquale B., 153
Desimone R., 203
Destexhe A., 119
DeWitt E., 57
Diamond M. E., 46, 197
DiCarlo J., 36, 76
Dickens L., 192
Diester I., 190
Doiron B., 39
Donoghue J., 133

- Doya K., 134
Dragoi V., 96
Drew P., 142
Drugowitsch J., 29
Duan C. A., 110
Dubreuil A., 211
Dudman J., 182
Duguid I., 190
Dutt N., 46
- Eagleman D., 87
Ebihara A., 75
Ecker A. S., 143
Eden U., 62, 175
Egnor R., 124
El-Shamayleh Y., 130
Eliasmith C., 63
Elsayed G., 191
Engel S., 85
Engel T., 150
Engelken R., 102
Engert F., 44
Erlich J., 110
Ernst U., 124, 158
Escola S., 208
Eskandar E., 133
Esmaeili V., 46, 197
Euler T., 144
Everling S., 56
- F. Mainen Z., 57, 198
Fagg A., 118
Faisal A., 100, 104, 192
Farkhooi F., 172
Fassihi A., 46, 197
Felsen G., 59
Fernandes H., 34
Ferrera V., 55
Fetter R., 66
Field D., 195
Field G., 89
Fiete I., 24, 68, 71, 116
Fink A., 189
Fiorillo C., 65
Fiorini M., 185
Fiser J., 55, 76, 149
Fisher D., 145
Fitzgerald J., 123, 203
Florez-Weidinger J., 148
Frady E., 178
Frank M., 53, 58, 77, 165
Franklin N., 77
Freeman J., 89
Fregnac Y., 27
Freifeld L., 30
- Freiwald W., 75
Fried I., 80, 138
Friederich U., 206
Froemke R. C., 58
Frohlich F., 147
Froudarakis E., 93, 143
Fukai T., 64
Funamizu A., 134
Fusi S., 102, 166, 208
- Gallant J., 105, 199, 200
Gallinat J., 168
Ganguli S., 28, 111, 129, 131, 163, 209
Gao P., 163
Gao Y., 142
Gardner J., 111, 167
Gardner T., 62, 73, 128
Geisel T., 138
Geisler W., 100, 159
Genkin A., 49
Gerhard F., 175
Gerich F., 146
Gerkin R., 70
Gerstner W., 67, 81, 108
Gilson M., 64, 97
Giret N., 28
Glimcher P., 34
Gohl D., 203
Golby A., 53
Gold J., 159
Golden J., 195
Golshani P., 205
Gomez C., 91
Goncalves P., 68
Gong P., 187, 188
Goo W., 190
Goodhill G., 151
Gopal V., 199
Gorchetchnikov A., 206
Gordon D., 25
Goris R., 149
Grabenhofer A., 199
Grabska-Barwinska A., 142
Graf A., 33
Graupner M., 176, 195
Gray C., 48, 89
Grayden D. B., 97
Gregor K., 92
Greschner M., 89
Grinvald A., 139
Grosenick L., 37
Grothe I., 158
Gu Y., 56
Guetig R., 41
Guitchounts G., 73

- Gunning D., 89
Guo R., 114
Gutierrez G. J., 175
Gutkin B., 115, 117
- Haefner R., 55, 149
Hagens O., 67
Hahnloser R., 28, 129
Halgren E., 119
Hall N., 204
Hangya B., 121, 165
Hanks T., 110
Hanuschkin A., 129
Hardie R., 206
Harris K., 41
Haslinger R., 156
Hass C., 40
Hatsopoulos N., 118
Hawken M., 90
Hayes T. R., 169
Heinz A., 168
Hennequin G., 81
Henry C., 90
Henson R., 83
Hermundstad A., 84
Higgins D., 176
Hillar C., 89, 94, 179
Hinton G., 207
Histed M. H., 37
Hochberg L., 133
Hollender L., 41
Holroyd T., 83
Hong H., 76
Hong S., 65, 82
Horowitz M., 30
Horwitz G., 40, 105
Houillon A., 168
Hromadka T., 184
Hrovat C., 74
Hu M., 96
Hu T., 49, 200
Hu Y., 98
Huang H., 183
Huang K., 44
Huang X., 146
Huang Y., 167
Huang Z. J., 121, 189
Huber D., 174
Hudson T., 61
Huh D., 137
Huk A., 52, 91
Hunsberger E., 63
Huo B., 142
- Igarashi Y., 109
- Iigaya K., 166
Isik L., 129
Issa E., 36
Ito M., 134
Itskov V., 99, 155
Izhikevich E., 122, 145, 180, 188, 202
- Jadzinsky P., 202
Jansen M., 143
Jepson L., 89
Jessell T. M., 136, 189
Jin J., 117, 143
Jogan M., 35
Johnson K., 135
Jones S., 38
Josic K., 92, 98, 106, 113
Juusola M., 206
- Kaardal J., 123
Kahana M., 186, 194
Kaifosh P., 42
Kalanithi P., 190
Kanitscheider I., 49, 108
Kanzaki R., 134
Kaping D., 56
Kapoor K., 171
Kaschube M., 125, 148
Katz D. B., 187
Katz Y., 123
Kaufman M., 80, 112, 115, 191, 194
Keil W., 125, 148
Keinath A., 182
Kelly C., 72
Kelly S., 117
Kempster R., 121
Kennedy A., 42
Kepecs A., 57, 121, 165
Kerr R., 97
Khaligh-Razavi S., 130
Khaw M. W., 28, 34
Kim H., 65
Kinkhabwala A., 71
Kispersky T., 52, 175
Kleberg F., 64
Knoblich U., 119
Knudson K., 154
Kochar A., 165
Koechlin E., 29
Koepsell K., 179
Kohn A., 80, 91
Kolterman B., 198
Komban S. J., 143
Kording K., 34
Kornfeld J., 28
Kornprobst P., 95

- Koster U., 48, 89, 94
Kostuk M., 183
Kouh M., 95
Koulakov A., 198
Koyama S., 173
Koyluoglu O. O., 116
Kramer M., 175
Kreiman G., 53, 91, 193
Kreiter A., 158
Kremkow J., 117, 143
Krichmar J., 46
Kriegeskorte N., 130
Krishnamurthy K., 159
Kristan W., 178
Kuhn D., 192
Kumar A., 177
Kvitsiani D., 121
- Lahiri S., 131
Lajoie G., 42
Lamp I., 123
Landy M., 61
Laquitaine S., 111
Lashgari R., 143
Latham P., 108, 116, 142, 207
Laubach M., 53
Laughlin S., 93
Lee A. M., 59
Lee D., 164, 171
Lee S. W., 113
Lega B., 186
Leibo J. Z., 75, 129
Lengyel M., 30, 51, 76, 149, 177, 205
Li J., 44
Li P., 89
Li X., 143
Lieberman M., 137
Lieder F., 112
Lim Y., 62, 128
Lin K. K., 42
Lindbloom-Brown Z., 40
Linderman S., 162
Lisberger S., 80
Lisman J., 198
Litke A., 89
Litwin-Kumar A., 39
Liu C., 31
Liu S., 56
Liu Y., 170
Livitz G., 206
Loewenstein Y., 166
Logothetis N. K., 61
Lopour B., 138
Lorenz R., 168
Louie K., 34
- Lu Z., 66
Lucke J., 172
Luedtke N., 50
Luz Y., 79
Löwel S., 125, 148
Lüdge T., 204
- Ma W. J., 45, 85, 113, 140
Machado T., 37, 136
Machens C., 33, 98
Macke J. H., 151, 207
Madhavan R., 193
Madsen J., 53, 193
Mairal J., 199
Maloney L. T., 166
Mandon S., 158
Marder E., 25, 52, 175
Margoliash D., 183
Markowitz D., 163
Markowitz J., 73
Marti D., 208
Mason C., 67
Masset P., 115
Masson G. S., 95
Mathieson K., 89
Maunsell J., 37
Mazzoni A., 61
McClelland J. L., 209
McFarland J., 193
Meier P., 122, 145
Meinertzhagen I., 66
Meir R., 211
Meliza D., 183
Mendonca A. G., 57
Mensi S., 67, 108
Merel J., 65
Meso A., 95
Mesquita L., 140
Meyers E., 129, 203
Michelson C., 109
Miconi T., 103
Milekovic T., 133
Mill R. W., 73
Miller K., 65, 139
Millman D., 193
Mingolla E., 206
Mitchell J., 184
Mitz A., 181
Miyata R., 127
Mizuseki K., 71
Mlynarski W., 128
Mochol G., 41
Monson B., 174
Monteforte M., 87, 102
Moore C., 38

- Moore T., 150
Moorkanikara Nageswaran J., 122, 145, 180, 202
Moran A., 187
Morgenstern Y., 174
Morton R., 199
Movshon A., 25, 80, 149
Mudigonda M., 89
Murray E., 181
Murray I., 207
Murray J., 171
Mutch J., 75
Muzzio I., 182
Mély D. A., 148
- Nandy A., 184
Narayanan N., 53
Naselaris T., 200
Nassar M., 159
Nassi J., 180
Naud R., 108
Nawrot M., 172
Negrello M., 82
Neishabouri A., 100
Neitzel S., 158
Newsome W. T., 166
Nishikawa I., 179
Niven J. E., 93
Niyogi R., 160
Nogaret A., 183
Norcia A., 203
Nykamp D., 212
- O'Doherty J., 88
O'Shea D., 190
Obermayer K., 114, 168
ODoherty J., 113
Ohayon S., 124
Okada M., 109
Oliver M., 199
Ollerenshaw D., 86, 197
Olman C., 200
Olshausen B., 48, 89, 94
Olveczky B., 44, 125
Op de Beeck H. P., 146
Orban G., 205
Orger M., 44
Osborne L., 63
Ostojic S., 195
Ota K., 127
Otchy T. M., 44
Oweiss K., 118
- Pachitariu M., 162
Padoa-Schioppa C., 171
Pagan M., 173
- Pakman A., 161
Palmer J., 187, 188
Palmer S. E., 63
Palmigiano A., 138
Pamplona D., 145
Paninski L., 37, 48, 65, 136, 152, 161
Panzeri S., 61
Parga N., 32
Park I. M., 47, 52, 60, 154
Park M., 105
Parra L. C., 83
Pastalkova E., 126
Pasternak T., 171
Pasupathy A., 130
Pawelzik K. R., 158
Paz R., 78, 80
Pearson J., 110
Pehlevan C., 44
Pelko M., 190
Perge J., 191
Perona P., 124
Perry J., 100
Pesaran B., 163
Petre C., 145, 180, 202
Petrov A., 169
Peyrache A., 119
Pfau D., 152
Pfister J., 99, 101
Piekiewicz F., 145, 180, 202
Pillow J. W., 47, 52, 60, 105, 109, 154
Pipa G., 156
Piray P., 133
Pitkow X. S., 56, 108
Platt M., 110, 170
Plaza S., 66
Plenz D., 83
Pnevmatikakis E. A., 37, 65, 89, 152
Poggio T., 75, 129, 203
Poirazi P., 118
Polack P., 205
Poole B., 37
Popovic M., 149
Portugues R., 44
Pouget A., 49, 56, 108, 142
Powell N., 155
Pozzorini C., 67, 108
Prentice J., 72, 106, 127
Priebe N., 80, 154
Puelma Touzel M., 87
Puggioni P., 190
Pujuc Z., 151
Purves D., 174
- Rabinowitz N., 89
Rahnama Rad K., 48

- Ramakrishnan C., 190
Ramaswamy V., 69
Ramayya A., 186
Ramirez A., 65
Ranade S., 121
Rankin J., 95
Rao R., 167
Raposo D., 112, 115
Rasumov N., 93
Razak K., 178
Reato D., 83
Reid A., 184
Remme M., 121
Renart A., 41
Resulaj A., 32
Reynolds J., 95, 180, 184
Richert M., 122, 145, 180, 188, 202
Rieke F., 40
Rigotti M., 208
Rinberg D., 32, 86, 198
Ringach D., 138
Ripple J., 181
Robles-Llana D., 68, 71
Robson D., 44
Roe A., 180
Roitman M., 132
Romani S., 139
Romo R., 32
Rosenberg A., 36
Rotermund D., 158
Roth N., 186
Rothkopf C. A., 82, 145
Rouault H., 92
Roy S., 60
Rozell C. J., 48
Rubin J., 181
Rudebeck P., 181
Rumpel S., 74
Rust N., 173, 186
Ryu S., 80, 163, 191
- Sabes P. N., 45, 88, 135
Safaai H., 185
Saggau P., 93, 143
Sahani M., 151, 162
Sakata S., 41
Salinas E., 38
Sanders H., 198
Sanders J., 165
Santaniello S., 139
Santhanam G., 163
Sarma S., 139
Sarode S., 159
Savin C., 51, 76
Savitskaya J., 70
- Sawtell N. B., 42
Saxe A., 209
Schafer R., 203
Scheffer L., 66
Schier A., 44
Schiess M., 31
Schilling V., 67
Schmid A., 104
Schnabel M., 125, 148
Schnepel P., 177
Schoonover C., 67
Schottdorf M., 125
Schrater P., 26, 54, 155, 169, 171
Schreiber S., 121
Schuur F., 166
Schwab D., 96, 106
Schwartz O., 39
Schwarz C., 153
Scott E., 151
Scott M., 63
Scott S. H., 101
Seely J., 80
Segev R., 196
Seidemann E., 109
Seitz A., 178
Sejnowski T., 27
Sellers K., 147
Senn W., 31, 99, 204
Serences J., 174
Series P., 178
Serre T., 148
Shahidi N., 96
Shamir M., 79
Shanechi M., 137
Shankar K., 161
Sharpee T. O., 78, 95, 123, 184, 196
Shea-Brown E., 42, 98, 107
Sheikh A., 172
Shelton J. A., 172
Shen H., 135
Shen S., 85
Shenoy K., 80, 163, 190, 191, 194
Sherbakov L., 206
Shimojo S., 113
Shin H., 45
Shinn-Cunningham B., 26, 128
Shizgal P., 160
Shriki O., 83
Shusterman R., 86
Sidiropoulou K., 118
Sillies M., 203
Simmons K., 106
Simoncelli E., 149
Sinha S., 60, 140
Sirotin Y., 86

- Sjöström P. J., 210
Smear M., 32
Smith C., 48, 161
Smith M., 80, 83
Smolyanskaya A., 91
Snoek J., 157
Sober S., 72
Sohail A., 206
Sohl-Dickstein J., 179
Sokol S., 145, 180, 202
Solomon E., 76
Solomon R., 160
Solt K., 137
Soltani A., 170
Sommer F., 71
Sommer T., 114
Sompolinsky H., 69, 164
Soudry D., 211
Southerland J., 118
Sprekeler H., 120, 121
Srivastava N., 169, 171
Stanford University T., 38
Stanley G., 86, 117, 197
Stansbury D., 105, 200
Steinmetz N., 103, 150
Stepanyants A., 210
Stern M., 69
Sterne P., 172
Stevenson I., 48, 71, 89
Stocker A., 35, 160, 164
Strouse D., 177
Sugrue L. P., 166
Suh B., 209
Sumner C., 73
Sunkara A., 201
Surace S. C., 101
Sussillo D., 194
Szatmary B., 145, 180, 188, 202
- Tafazoli S., 185
Tai L., 59
Takahashi H., 134
Takemura S., 66
Tam B. P., 166
Tang H., 193
Tank D., 71
Tankus A., 80
Tapia J., 67
Taub A., 123
Taubman H., 80
Taylor A., 124
Tee S. P., 74
Teichert T., 55
Theis L., 50, 153
Thomas D. A., 97
- Thompson J., 59
Ticchi A., 104
Tkacik G., 84
Tobia M., 114
Toemen N., 124
Tolias A., 92, 93, 143
Toni I., 133
Toyoizumi T., 179
Trautmann E., 163
Trenholm S., 96
Triesch J., 145
Tripathy S. J., 70
Trousdale J., 98, 106
Tsodyks M., 139
- Ugurbil K., 200
Ujfalussy B. B., 30, 177
Urban N., 70
Urbanczik R., 31, 204
- Vaadia E., 80
Vaidya M., 118
Van den Berg R., 45, 85, 113
Van den Bergh G., 146
van der Schaaf M., 133
van Rossum M., 190, 210
Veliz-Cuba A., 155
Vermaercke B., 146
Versace M., 206
Victor J., 84, 104
Vilankar K., 195
Vineyard C., 77
Vitaladevuni S., 66
Vogels T. P., 81, 120
Vogelstein J., 37
- Wang Q., 86, 117, 197
Wang X., 75, 156, 171, 181
Wang Y., 117, 126, 143, 176
Wang Z., 164
Wayne G., 42, 152
Wei W., 181
Wei X., 127, 160
Wei Z., 200
Weidemann C. T., 186
Weller J. P., 105
White L. E., 125, 148
Wiecki T., 165
Wilbrecht L., 59
Williams Z., 156, 157
Willmore B., 199
Wimmer V., 67
Witt A., 138
Wohl M., 186
Wolf F., 87, 102, 125, 148

Womelsdorf T., 56
Woolley S. M. N., 51
Wu S., 170
Wyart V., 29

Xiong Q., 43

Yamins D., 76
Yang H., 30
Yarrow S., 178
Yates J., 52
Yates T., 50
Yatsenko D., 92
Yoon K., 71
Youngs N., 155
Yu B., 163, 199
Yu D., 55
Yu Y., 104
Yun S., 65
Yvert B., 133

Zador A., 43, 184
Zaghloul K., 186
Zaidel A., 140
Zaidi Q., 143
Zaika N., 58
Zeldenrust F., 117
Zemel R., 157, 207, 212
Zhang J., 32
Zhang W., 120
Zhang Y., 203
Zheng H., 86, 197
Zhou Y., 75
Zhu D., 38
Zhu M., 48
Znamenskiy P., 43
Zoccolan D., 185
Zylberberg J., 107