

# COSYNE 2014

[www.cosyne.org](http://www.cosyne.org)

**MAIN MEETING**

**Salt Lake City, UT**

**Feb 27 - Mar 2**

---



---

# Program Summary

## Thursday, 27 February

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

## Friday, 28 February

- 7:30 am **Breakfast**
- 8:30 am **Session 2: Circuits I: From wiring to function**  
Invited speaker: Thomas Mrsic-Flogel; 3 accepted talks
- 10:30 am **Session 3: Circuits II: Population recording**  
Invited speaker: Elad Schneidman; 3 accepted talks
- 12:00 pm **Lunch break**
- 2:00 pm **Session 4: Circuits III: Network models**  
5 accepted talks
- 3:45 pm **Session 5: Navigation: From phenomenon to mechanism**  
Invited speakers: Nachum Ulanovsky, Jeffrey Magee; 1 accepted talk
- 5:30 pm **Dinner break**
- 7:30 pm **Poster Session II**

## Saturday, 1 March

- 7:30 am **Breakfast**
- 8:30 am **Session 6: Behavior I: Dissecting innate movement**  
Invited speaker: Hopi Hoekstra; 3 accepted talks
- 10:30 am **Session 7: Behavior II: Motor learning**  
Invited speaker: Rui Costa; 2 accepted talks
- 11:45 am **Lunch break**
- 2:00 pm **Session 8: Behavior III: Motor performance**  
Invited speaker: John Krakauer; 2 accepted talks
- 3:45 pm **Session 9: Reward: Learning and prediction**  
Invited speaker: Yael Niv; 3 accepted talks
- 5:15 pm **Dinner break**
- 7:30 pm **Poster Session III**

---

## Sunday, 2 March

7:30 am **Breakfast**

8:30 am **Session 10: Decision making**

Invited speakers: Anne Churchland, Joshua Gold; 2 accepted talks

11:00 am **Session 11: Modulation: (Dis)Engaging sensory processing**

4 accepted talks

12:00 pm **Lunch break**

2:00 pm **Session 12: Perception**

Invited speaker: Doris Tsao; 2 accepted talks

---

# Poster Session Topics

<b>Topic Area</b>	<b>Session I</b> <i>Thursday</i>	<b>Session II</b> <i>Friday</i>	<b>Session III</b> <i>Saturday</i>
Optogenetics	1–5		
Behavior	6–12		
Neural encoding, decoding	13	1–31	
Cognition	14–55		1–3
Computational, modeling	56–100	32–35	4–62
Other representations		36–39	
Sensory systems		40–88	63–64
Techniques		89–100	
Learning, plasticity			65–84
Motor systems			85–100

Brain Corporation leverages neuroscience and machine learning towards a platform enabling robots who learn. Our software and algorithms add intelligence to robotics platforms allowing them to become teachable, much like animals. Brain Corporation partners with Qualcomm, leader in mobile technology.

## Robots who learn.

### Seeking Experts in Sensorimotor Control and Software Development

A number of full-time positions at all levels are available immediately in theoretical and computational modeling of sensorimotor control at Brain Corporation, San Diego, CA.

Submit your CV/resume, relevant papers, and source code samples to [jobs@braincorporation.com](mailto:jobs@braincorporation.com). In your cover letter, please address at least two of the requirements below.

#### Requirements:

We are seeking candidates who are exceptional in one or more of the following areas:

- **Feed-forward and recurrent neural networks, spike-timing plasticity and dynamics.**
- **Machine learning techniques, including convolutional networks and deep learning.**
- **Sensorimotor transformations, sensorimotor processing and reinforcement learning.**

Plus have hacker level ability in regards to programming:

- **Programming in two or more languages, including Python and C/C++/Objective-C.**

#### Competitive Compensation Packages

The employee compensation package includes a stock option grant, matching 401k retirement contributions, medical, dental and vision insurance coverage, and annual performance bonuses. Employees have access to facilities on the Qualcomm campus and the office is located in proximity to UCSD and the Salk Institute in San Diego, California.

Please visit [www.braincorporation.com](http://www.braincorporation.com) for more information.



brain  
CORPORATION



# 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/research](http://www.qualcomm.com/research)

**QUALCOMM**<sup>®</sup>





## Record. Process. Stimulate. Store.


### Thousands of Channels Without Compromise.

TDT delivers the industry's most dependable and flexible research instrumentation.

Our new High-Density Cortical Arrays offer precision signal capture, perfect for multichannel recording and stimulation. Achieve maximum real-time control, when pairing with TDT's RZ2 Signal Processor.

Optimize your research by building your arrays with TDT's Online Cortical Array Designer. With configuration ranges from 16 to 400 channels, define each electrode length to customize to your exact needs.


Pair your arrays with TDT's PZ5 NeuroDigitizer. With the best signal-to-noise ratio and dynamic range of any amplifier available today, the PZ5 is an industry favorite for multichannel recordings. Acquire evoked potentials, EEG, EMG and analog signals - all with a single device.



[sales@tdt.com](mailto:sales@tdt.com) | 386.462.9622



 The MIT Press



The MIT Press [mitpress.mit.edu](http://mitpress.mit.edu)

### **THE NEW VISUAL NEUROSCIENCES**

edited by John S. Werner and  
Leo M. Chalupa

A comprehensive review of contemporary research in the vision sciences, reflecting the rapid advances of recent years.

584 pp., 575 b&w illus., 281 color plates, \$250 cloth

### **NEUROSCIENCE**

A Historical Introduction

Mitchell Glickstein

An introduction to the structure and function of the nervous system that emphasizes the history of experiments and observations that led to modern neuroscientific knowledge.

376 pp., 52 color illus., 119 b&w illus., \$50 cloth

### **BRAIN STRUCTURE AND ITS ORIGINS**

Function, Evolution, Development

Gerald E. Schneider

An introduction to the brain's anatomical organization and functions with explanations in terms of evolutionary adaptations and development.

656 pp., 243 color illus., 127 b&w illus., \$75 cloth

### **HUMAN ROBOTICS**

Neuromechanics and Motor Control

Etienne Burdet, David W. Franklin,  
and Theodore E. Milner

A synthesis of biomechanics and neural control that draws on recent advances in robotics to address control problems solved by the human sensorimotor system.

288 pp., 104 illus., \$45 cloth

### **THE COGNITIVE-EMOTIONAL BRAIN**

From Interactions to Integration

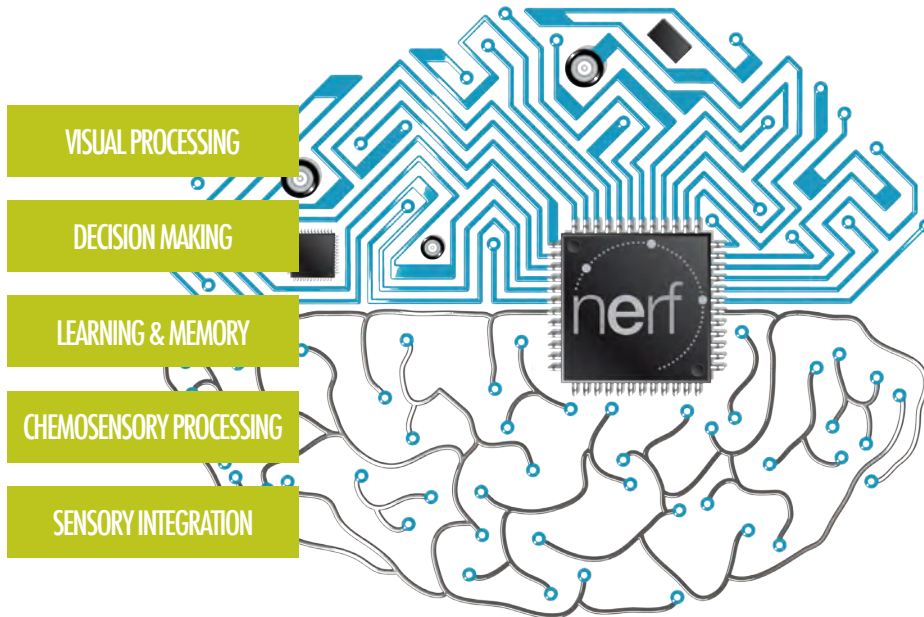
Luiz Pessoa

A study that goes beyond the debate over functional specialization to describe the ways that emotion and cognition interact and are integrated in the brain.

336 pp., 70 b&w illus.,  
14 color plates,  
\$40 cloth

Visit the  
MIT PRESS  
BOOTH  
for a 30%  
DISCOUNT

# NEURO-ELECTRONICS RESEARCH FLANDERS



Also visit us during our annual **NERF Neurotechnology Symposium** on **May 5th and 6th, 2014**. The symposium focuses on novel technologies for studying neural circuits, brain machine interfacing and neuromorphics engineering.

Confirmed speakers:

- ▶ Elisabetta Chicca, Universität Bielefeld, Germany
- ▶ Valentina Emiliani, INSERM, France
- ▶ Michele Giugliano, University of Antwerp, Belgium
- ▶ Andreas Hierlemann, ETH, Switzerland
- ▶ Adi Mizrahi, The Hebrew University of Jerusalem, Israel
- ▶ Arto Nurmikko, Brown University, USA
- ▶ Bernardo Sabatini, HHMI, USA
- ▶ Andreas Schaefer, NIMR, UK
- ▶ Ian Wickersham, MIT, USA
- ▶ Mehmet Fatih Yanik, MIT, USA

We look forward to welcoming you at this exciting event !



KU LEUVEN



**NERF** is a joint basic research initiative, set up by imec, VIB and KU Leuven to unravel the functioning of the brain. At the NERF labs, located in the heart of Europe in Leuven Belgium, world-class neuroscientists look into fundamental neuroscientific questions through collaborative, interdisciplinary research, combining nanoelectronics with neurobiology, physics, engineering and computer science.

[www.nerf.be](http://www.nerf.be)



# 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






Being the **leader** has many advantages...



# AlphaLab SnR™

What makes the AlphaLab SnR the most advanced microelectrode recording system on the market...

- ✓ Unmatched Stimulation and Recording Capabilities
- ✓ Exceptional Versatility and Scalability
- ✓ Complete Experimental Control
- ✓ Online Data Access for Closed-Loop Applications



**Stop by our booth for free sample electrodes!**  
 Visit our new website at [www.alphaomega-eng.com](http://www.alphaomega-eng.com)




Booth #2024

**North America**  
 Alpha Omega Co. USA Inc.  
 5755 North Point Parkway  
 Unit 229, Alpharetta, GA 30022  
 Toll-Free: 1-877-919-6288  
 Phone: 770-521-2049  
 Fax: 1-877-471-2055


**Europe, Middle East & Africa**  
 Alpha Omega GmbH  
 Ubstadter Str. 28  
 76698 Ubstadt-Weiher  
 Germany  
 Toll Free: 00-800-2574-2111  
 Fax: +49 (0) 721-2391034

**Asia Pacific & Latin America**  
 Alpha Omega Engineering  
 PO Box 810  
 Nazareth Illit 17105  
 Israel  
 Tel 972-4-656-3327  
 Fax 972-4-657-4075


[www.alphaomega-eng.com](http://www.alphaomega-eng.com)    [info@alphaomega-eng.com](mailto:info@alphaomega-eng.com)



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



- ◆ Extensive online spike sorting
- ◆ 16 - 256 channels with DigiAmp™
- ◆ Real-time low latency performance
- ◆ Ultrafast file loading with PL2™



- ◆ Industry leading LED intensities
- ◆ LED/Laser driver with pattern generation
- ◆ Ultra high performance cables & implants
- ◆ Dual LED commutator with electrical data lines




- ◆ Stand-alone CinePlex Lite, or
- ◆ Full integration with Neural Recording
- ◆ Real-time & offline tracking
- ◆ Full range of cameras up to 200fps

Registration  
Now  
Open

5th Plexon Neurophysiology & Behavior Workshop  
March 17 - 20, 2014, Dallas, TX USA

[www.plexon.com](http://www.plexon.com)





Focus

 **BLACKROCK**<sup>®</sup>  
MICROSYSTEMS

[sales@blackrockmicro.com](mailto:sales@blackrockmicro.com)

*Electrophysiology, Simplified*

## 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 2014 Leadership

#### Organizing Committee:

- *General Chairs*: Marlene Cohen (University of Pittsburgh), Peter Latham (The Gatsby Computational Unit, UCL)
- *Program Chairs*: Stephanie Palmer (University of Chicago), Michael Long (New York University)
- *Workshop Chairs*: Tatyana Sharpee (Salk Institute), Robert Froemke (New York University)
- *Communications Chair*: Eugenia Chiappe (Champalimaud Neuroscience Programme)

#### 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:

- Michael Long (New York University), co-chair
- Stephanie Palmer (University of Chicago), co-chair
- Misha Ahrens (Janelia Farm)
- Asohan Amarasingham (The City University of New York)
- Bruno Averbeck (National Institutes of Health)
- Jeffrey Beck (Rochester University)
- Claudia Clopath (Imperial College)
- Brent Doiron (University of Pittsburgh)
- Robert Froemke (New York University)
- Vivek Jayaraman (Janelia Farm)
- Alla Karpova (Janelia Farm)
- Adam Kohn (Albert Einstein College of Medicine, Yeshiva University)
- Sean MacEvoy (Boston College)
- Katherine Nagel (Harvard University)
- Leslie Osborne (University of Chicago)
- Anitha Pasupathy (University of Washington)
- Nicholas Price (Monash University)
- Emilio Salinas (Wake Forest School of Medicine)
- Nate Sawtell (Columbia University)
- Sam Sober (Emory University)
- Gasper Tkacik (Institute for Science and Technology, Austria)
- Angela J Yu (University of California, San Diego)



### *Reviewers:*

Ana Amador, Ehsan Arabzadeh, Stephen Baccus, Pamela Baker, Dan Bendor, Vikas Bhandawat, Farran Briggs, Scott Brincat, Randy Bruno, Lars Buesing, Tim Buschman, Ioana Carcea, Stijn Cassenaer, Maurice Chacron, Samarth Chandra, Eugenia Chiappe, Paul Cisek, Ruben Coen-Cagli, Carina Curto, Jaime de la Rocha, Michael DeWeese, Kamran Diba, Jan Drugowitsch, Joshua Dudman, Alexander Ecker, Yasmine El-Shamayleh, Jacob Engelmann, Bernhard Englitz, Aldo Faisal, Jeremy Freeman, Makoto Fukushima, Julijana Gjorgjieva, Mark Goldman, Pedro Goncalves, Arnulf Graf, Sebastian Guderian, Adam Hantman, Matthew Harrison, Guillaume Hennequin, Greg Horwitz, Alexander Huk, Shantanu Jadhav, Yan Karklin, Zachary Kilpatrick, Sung Soo Kim, Minoru Koyama, Dwight Kravitz, Nikolaus Kriegeskorte, Arthur Leblois, Anthony Leonardo, Sukbin Lim, Matthieu Louis, Leo Lui, Cheng Ly, Jakob Macke, Olivier Marre, Timothee Masquelier, James Mazer, J Lucas McKay, Ferenc Mechler, Hugo Merchant, Jason Middleton, Paul Miller, Andrew Miri, Thierry Mora, Ruben Moreno-Bote, Adam Morris, Gabe Murphy, Ferdinando Mussa-Ivaldi, Thomas Naselaris, John O'Doherty, Anne-Marie Oswald, Joe Paton, Franco Pestilli, Biljana Petreska, Xaq Pitkow, Jason Prentice, Alexander Reyes, Robert Rosenbaum, Cristina Savin, David Schoppik, David Schwab, Walter Senn, John Serences, Jesper Sjostrom, Matt Smith, Joo-Hyun Song, Greg Stephens, Ian Stevenson, Gowan Tervo, Gelsy Torres-Oviedo, Wilson Truccolo, Tomoki Tsuchida, Naotsugu Tsuchiya, Srinivas Turaga, Mark van Rossum, Tim Vogels, Michael Wehr, Robert Wilson, Elizabeth Zavitz, Kechen Zhang, Shunan Zhang

### **Conference Support**

#### *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 2014.

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
- Evolved Machines
- National Science Foundation (NSF)

The 2014 recipients are:

Federico Carnevale, Julie Elie, Neir Eshel, Peiran Gao, Leon Gatys, Julijana Gjorgjieva, Thiago Gouvea, Maureen Hagan, Liberty Hamilton, Patrick Kaifosh, Ann Kennedy, Subhaneil Lahiri, Laura Lewis, Sukbin Lim, Robert Rasmussen, Patrick Sadtler, Jason Samonds, Cristina Savin, Jeffrey Seely, Alexandra Smolyanskaya, Merav Stern, Alan Veliz-Cuba

### **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 2014 recipients are:

Jeremie Barral, Helen Barron, Vahid Esmaeili, Aditya Gilra, Keren Haroush, Mark Ioffe, Sanjeev Khanna, Yoshiaki Ko, Veronika Koren, Matthias Kummerer, Oh-Sang Kwon, Eunjeong Lee, Yin Li, HaDi MaBouDi, Sarah Marzen, Feryal Mehraban Pour Behbahani, Mehrab Modi, Nader Nikbakht Nasrabadi, Megan Peters, Francesca Pizzorni-Ferrarese, Ofri Raviv, Daria Rylkova, Adam Snyder, Siyu Zhang

### **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 2014 NSF Cosyne Mentors are listed below, each followed by their mentee:

Yael Niv and Angela Radulescu, Joseph Paton and Asma Motiwala, Samuel Sober and Claire Tang, Michael Wehr and Michael Kyweriga



# Program

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

[http://cosyne.org/c/index.php?title=Cosyne2014\\_Program](http://cosyne.org/c/index.php?title=Cosyne2014_Program).

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

## Thursday, 27 February

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

### Session 1: Keynote

(Chair: Michael Long)

- 6:30 pm      [Circuits and strategies for skilled reach: From internal copies to gain control](#)  
Thomas Jessell, Columbia University (**invited**) . . . . . 25
- 7:30 pm      **Poster Session I**

## Friday, 28 February

- 7:30 am      **Continental breakfast**

### Session 2: Circuits I: From wiring to function

(Chair: Claudia Clopath)

- 8:30 am      [The functional organisation of synaptic connectivity in visual cortical microcircuits](#)  
Thomas Mrsic-Flogel, Universitat Basel (**invited**) . . . . . 26
- 9:15 am      [A unifying theory of receptive field formation](#)  
C. S. N. Brito, W. Gerstner, Ecole Polytechnique Federale de Lausanne . . . . . 42
- 9:30 am      [Learning to track moving stimuli with population codes](#)  
J. Makin, B. Dichter, P. Sabes, University of California, San Francisco . . . . . 46
- 9:45am      [Processing properties of medulla neurons define neural substrates of motion detection in Drosophila](#)  
D. Clark, R. Behnia, T. Clandinin, C. Desplan, Yale University . . . . . 31
- 10:00 am      **Coffee break**

## Program

---

### Session 3: Circuits II: Population recording

(Chair: Gasper Tkacik)

10:30am	<a href="#">Words, metrics, and a thesaurus for a neural population code</a> Elad Schneidman, Weizmann Institute ( <b>invited</b> ) . . . . .	26
11:15am	<a href="#">A theory of neural dimensionality and measurement</a> P. Gao, E. Trautmann, B. Yu, G. Santhanam, S. Ryu, K. Shenoy, S. Ganguli, Stanford University . . . . .	31
11:30am	<a href="#">Mapping the brain at scale</a> J. Freeman, N. Vladimirov, T. Kawashima, Y. Mu, D. V. Bennett, J. Rosen, C. Yang, L. Looger, M. Ahrens, Janelia Farm Research Campus . . . . .	32
11:45am	<a href="#">Whole-brain imaging of a stereotypic nervous system reveals temporal coordination of global activity</a> S. Kato, T. Schroedel, H. Kaplan, M. Zimmer, Institute of Molecular Pathology . . . . .	33
12:00 pm	<b>Lunch break</b>	

### Session 4: Circuits III: Network models

(Chair: Brent Doiron)

2:00pm	<a href="#">Self-monitoring of network activity in multi-state recurrent neural networks</a> S. Escola, L. Abbott, Columbia University . . . . .	33
2:15pm	<a href="#">The uncertain spike: a new spike-based code for probabilistic computation</a> C. Savin, S. Deneve, IST Austria . . . . .	34
2:30pm	<a href="#">Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons</a> S. Ostojic, Ecole Normale Supérieure . . . . .	34
2:45pm	<a href="#">Balanced cortical microcircuitry for spatial working memory based on corrective feedback control</a> S. Lim, M. Goldman, University of Chicago . . . . .	35
3:00pm	<a href="#">Intrinsic neuronal properties switch the mode of information transmission in networks</a> J. Gjorgjieva, R. Mease, W. Moody, A. L. Fairhall, Harvard University . . . . .	35
3:15 pm	<b>Coffee break</b>	

### Session 5: Navigation: From phenomenon to mechanism

(Chair: Stephanie Palmer)

3:45pm	<a href="#">Neural codes for 2-D and 3-D space in the hippocampal formation of bats</a> Nachum Ulanovsky, Weizmann Institute ( <b>invited</b> ) . . . . .	26
4:30pm	<a href="#">Nonlinear input integration in cortical pyramidal neurons during behavior</a> Jeffrey Magee, Janelia Farm Research Campus ( <b>invited</b> ) . . . . .	27
5:15pm	<a href="#">A compartment-specific synaptic integration and plasticity model for hippocampal memory</a> P. Kaifosh, J. Zaremba, N. Danielson, A. Losonczy, Columbia University . . . . .	36
5:30 pm	<b>Dinner break</b>	
7:30 pm	<b>Poster Session II</b>	

## Saturday, 1 March

7:30 am **Continental breakfast**

### Session 6: Behavior I: Dissecting innate movement

(Chair: Vivek Jayaraman)

8:30am	<a href="#">Digging for genes that contribute to behavioral evolution</a> Hopi Hoekstra, Harvard University ( <b>invited</b> ) . . . . .	27
9:15am	<a href="#">Stereotypy and the structure of behavioral space</a> G. Berman, D. Choi, U. Klibaite, W. Bialek, J. Shaevitz, Princeton University . . . . .	37
9:30am	<a href="#">Measuring and modeling the dynamics of the thermal memory of <i>C. elegans</i></a> I. Nemenman, K. Palanski, F. Bartumeus, W. Ryu, Emory University . . . . .	37
9:45am	<a href="#">A quantitative model for the sensorimotor integration underlying chemotaxis in the <i>Drosophila</i> larva</a> M. Louis, A. Gomez-Marin, A. Schulze, V. Rajendran, G. Lott, M. Musy, J. Sharpe, M. Venkadesan, P. Ahammad, S. Druckmann, V. Jayaraman, Center for Genomic Regulation . . . . .	38
10:00 am	<b>Coffee break</b>	

### Session 7: Behavior II: Motor learning

(Chair: Samuel Sober)

10:30am	<a href="#">Generating and shaping novel action repertoires</a> Rui Costa, Champalimaud Centre for the Unknown ( <b>invited</b> ) . . . . .	27
11:15am	<a href="#">Modeling enhanced and impaired learning with enhanced plasticity</a> S. Lahiri, T. B. Nguyen-Vu, G. Q. Zhao, A. Suvrathan, H. Lee, S. Ganguli, C. J. Shatz, J. L. Raymond, Stanford University . . . . .	39
11:30am	<a href="#">Neural constraints on learning</a> P. Sadtler, K. Quick, M. Golub, S. Ryu, B. Yu, A. Batista, University of Pittsburgh . . . . .	39
11:45 am	<b>Lunch break</b>	

### Session 8: Behavior III: Motor performance

(Chair: Leslie Osborne)

2:00pm	<a href="#">Some new thoughts about planning, learning and skill in the motor domain</a> John Krakauer, Johns Hopkins University ( <b>invited</b> ) . . . . .	28
2:45pm	<a href="#">State-space models for cortical-muscle transformations</a> J. Seely, M. Kaufman, C. J. Cueva, L. Paninski, K. Shenoy, M. Churchland, Columbia University . . . . .	40
3:00pm	<a href="#">Spatial organization of synchronous cell assemblies in HVC</a> J. Markowitz, Y. Lim, G. Guitchounts, T. Gardner, Boston University . . . . .	41
3:15 pm	<b>Coffee break</b>	

### Session 9: Reward: Learning and prediction

(Chair: Kenway Louie)

3:45pm	<a href="#">Learning task representations for reinforcement learning</a> Yael Niv, Princeton University ( <b>invited</b> ) . . . . .	28
--------	---	----

## Program

---

4:30pm	<a href="#">Arithmetic of dopamine prediction errors: subtraction with scaled inhibition</a> N. Eshel, J. Tian, N. Uchida, Harvard University . . . . .	41
4:45pm	<a href="#">Nucleus basalis cholinergic neurons broadcast precisely timed reinforcement signals</a> B. Hangya, S. Ranade, A. Kepecs, Cold Spring Harbor Laboratory . . . . .	42
5:00pm	<a href="#">Neural activity predicts social interaction decisions during prisoner's dilemma games in primates</a> K. Haroush, Z. Williams, Harvard Medical School . . . . .	42
5:15 pm	<b>Dinner break</b>	
7:30 pm	<b>Poster Session III</b>	

## Sunday, 2 March

7:30 am **Continental breakfast**

### Session 10: Decision making

(Chair: David Freedman)

8:30am	<a href="#">A multisensory decision task exposes mixed selectivity in posterior parietal cortex</a> Anne Churchland, Cold Spring Harbor Laboratory ( <b>invited</b> ) . . . . .	28
9:15am	<a href="#">Dissociated functional significance of choice-related activity across the primate dorsal stream</a> J. Yates, L. Katz, I. M. Park, J. W. Pillow, A. Huk, University of Texas at Austin . . . . .	43
9:30am	<a href="#">Task specific feedforward component of choice-related activity in MT</a> A. Smolyanskaya, S. G. Lomber, R. Born, University of Pennsylvania . . . . .	43
9:45am	<a href="#">Adaptive decision-making in a dynamic world</a> Joshua Gold, University of Pennsylvania ( <b>invited</b> ) . . . . .	29
10:30 am	<b>Coffee break</b>	

### Session 11: Modulation: (Dis)Engaging sensory processing

(Chair: Adam Kepecs)

11:00am	<a href="#">Long-range and local circuits for top-down modulation of visual cortical processing</a> S. Zhang, M. Xu, T. Kamigaki, S. Jenvay, Y. Dan, University of California, Berkeley . . . . .	44
11:15am	<a href="#">Cortical modulation of synaptic dynamics and sensory responses in auditory cortex during movement</a> D. Schneider, A. Nelson, R. Mooney, Duke University . . . . .	44
11:30am	<a href="#">The role of local inhibitory interneurons in stimulus-specific adaptation in primary auditory cortex</a> R. Natan, L. Mwilambwe-Tshilobo, M. Geffen, University of Pennsylvania . . . . .	45
11:45am	<a href="#">Thalamic reticular nucleus controls arousal through modulation of local thalamocortical dynamics</a> L. Lewis, J. Voigts, F. Flores, M. Wilson, M. Halassa, E. Brown, Massachusetts Institute of Technology . . . . .	46
12:00 pm	<b>Lunch break</b>	



**Session 12: Perception**

(Chair: Justin Gardner)

2:00pm	<a href="#">Bayesian active sensing in the categorization of visual patterns</a> S. C. Yang, M. Lengyel, D. Wolpert, Cambridge University . . . . .	46
2:15pm	<a href="#">Both spatial and temporal codes shape texture perception</a> H. Saal, J. Lieber, A. Weber, S. Bensmaia, University of Chicago . . . . .	47
2:30pm	<a href="#">The macaque face processing system</a> Doris Tsao, California Institute of Technology ( <b>invited</b> ) . . . . .	29
3:15 pm	<b>Closing remarks</b>	

## Poster Session I

7:30 pm Thursday 27th February

I-1. Coding consequences of activity propagation from sensory and artificial neural stimulation Garrett B Stanley, Christopher Rozell, Daniel Millard, Georgia Institute of Technology . . . . .	47
I-2. Auditory cortex mediates perceptual gap detection. Aldis Weible, Alexandra Moore, Christine Liu, Michael Wehr, University of Oregon . . . . .	48
I-3. Simultaneous optogenetic manipulation and calcium imaging in freely moving <i>C. elegans</i> Andrew Leifer, Frederick Shipley, Christopher Clark, Mark Alkema, Princeton University . . . . .	49
I-4. Modeling optogenetic manipulation of neural circuits in macaque visual cortex Michael Avery, Jonathan Nassi, John Reynolds, Salk Institute for Biological Studies . . . . .	49
I-5. Amplification of synchrony in the <i>Drosophila</i> antennal lobe James Jeanne, Rachel Wilson, Harvard University . . . . .	50
I-6. Effects of temporal information on perceptual decision making Guang-Yu Robert Yang, Hyojung Seo, Xiao-Jing Wang, Daeyeol Lee, New York University . . . . .	50
I-7. Neural computations in the auditory system of fruit flies — from sensory coding to behavior Jan Clemens, Cyrille Girardin, Pip Coen, Barry Dickson, Mala Murthy, Princeton University . . . . .	51
I-8. Adaptation in an auditory-motor feedback loop contributes to generating repeated vocal sequences Kristofer Bouchard, Jason Wittenbach, Michael Brainard, Dezhe Jin, University of California, San Francisco	51
I-9. Information flow through a sensorimotor circuit: Spatial orientation in <i>C. elegans</i> Eduardo Izquierdo, Paul L. Williams, Randall D. Beer, Indiana University . . . . .	52
I-10. Basal forebrain mechanisms of sleep-wake regulation Min Xu, Siyu Zhang, Yang Dan, University of California, Berkeley . . . . .	53
I-11. An olfactory cocktail party - Detection of single odorants within mixtures in rodents. Dan Rokni, Vivian Hemmelder, Vikrant Kapoor, Venkatesh N Murthy, Harvard University . . . . .	53
I-12. Optimality, not simplicity governs visual decision-making Wei Ji Ma, Shan Shen, New York University . . . . .	54
I-13. Changes in prefrontal circuit organization increase repetitive network activity in models of autism Francisco Luongo, Meryl Horn, Vikaas Sohal, University of California, San Francisco . . . . .	54
I-14. Attention and uncertainty during reward contingency learning R Becket Ebitz, Eddy Albarraan, Alireza Soltani, Tirin Moore, Stanford University . . . . .	55
I-15. The optimal allocation of attentional resource Nuwan de Silva, Wei Ji Ma, New York University . . . . .	55
I-16. Surprise-based learning: neuromodulation by surprise in multi-factor-learning-rules Mohammadjavad Faraji, Kerstin Preuschoff, Wulfram Gerstner, Ecole Polytechnique Federale de Lausanne	56
I-17. Neural mechanisms of object-based attention Daniel Baldauf, Robert Desimone, Massachusetts Institute of Technology . . . . .	56
I-18. Characterizing the impact of category uncertainty on human auditory categorization behavior Adam Gifford, Yale Cohen, Alan Stocker, University of Pennsylvania . . . . .	57
I-19. Visual categorization reflects second order generative statistics Feryal Mehraban Pour Behbahani, Aldo A. Faisal, Imperial College London . . . . .	57
I-20. Continuous psychometric-neurometric comparison in a perceptual decision making task Tiago S Gouvea, Tiago Monteiro, Asma Motiwala, Joe Paton, Champalimaud Neuroscience Programme	58
I-21. Dissecting the contributions of area MSTd and VIP to heading perception Kaushik Lakshminarasimhan, Sheng Liu, Eliana M Klier, Yong Gu, Gregory C. DeAngelis, Xaq Pitkow, Dora E. Angelaki, Baylor College of Medicine . . . . .	58
I-22. Computational model-based tDCS selectively enhances guilt-aversion based prosocial behavior Tsuyoshi Nihonsugi, Masahiko Haruno, Gifu Shotoku . . . . .	59

I-23. Stochastic behavior in rats and its gating by the medial prefrontal cortex Gowan Tervo, Mikhail Proskurin, Maxim Manakov, Mayank Kabra, Alison Vollmer, Kristin Branson, Alla Karpova, Janelia Farm Research Campus . . . . .	60
I-24. The role of rat medial prefrontal cortex and superior colliculus in a rapid rule-switching task Chunyu A Duan, Jeffrey Erlich, Katherine Fitch, Carlos Brody, Princeton University . . . . .	60
I-25. Information-theoretic basis of making decisions in the absence of new stimuli Adam Calhoun, Sreekanth Chalasani, Tatyana Sharpee, University of California, San Diego . . . . .	61
I-26. Humans categorize two visual features of a single object sequentially Yul Kang, Michael Shadlen, Columbia University . . . . .	61
I-27. Expected information signals in the posterior cingulate cortex David Barack, Jean-Francois Gariépy, Michael Platt, Duke University . . . . .	62
I-28. Ventral pallidum neurons signal an information-induced bias in handling reward uncertainty Ethan Bromberg-Martin, Okihide Hikosaka, National Eye Institute . . . . .	62
I-29. On the neural dynamics of perceptual decision-making under temporal uncertainty Federico Carnevale, Omri Barak, Victor de Lafuente, Ranulfo Romo, Nestor Parga, Universidad Autonoma de Madrid . . . . .	63
I-30. Hierarchical competitions subserving multi-attribute choice. Laurence Hunt, Ray Dolan, Timothy Behrens, University College London . . . . .	64
I-31. A Bayesian perspective on flexibly responding to stochastic and non-stationary tasks: a role for str Nicholas Franklin, Michael Frank, Brown University . . . . .	64
I-32. Graded representation of evidence accumulation and decision outcome in lateral intraparietal cortex Robbe L.T. Goris, Roozbeh Kiani, New York University . . . . .	65
I-33. Adaptive value coding and the optimality of choice Kenway Louie, Ryan Webb, Paul Glimcher, New York University . . . . .	65
I-34. An ideal observer solution for a perceptual decision making task Paul Masset, Joshua Sanders, Justin Kinney, Adam Kepecs, Watson School of Biological Sciences . . . . .	66
I-35. Probing for patterns: Blending reward and information in a non-stationary POMDP Jeffrey Cockburn, Michael Frank, Brown University . . . . .	66
I-36. Probing the sources of suboptimality in human Bayesian inference Luigi Acerbi, Sethu Vijayakumar, Daniel Wolpert, University of Edinburgh . . . . .	67
I-37. Neural computations underlying dynamic decisions with disparate value types Vassilios Christopoulos, Paul Schrater, California Institute of Technology . . . . .	67
I-38. Visual decisions in the presence of stimulus and measurement correlations Manisha Bhardwaj, Sam Carroll, Wei Ji Ma, Kresimir Josic, University of Houston . . . . .	68
I-39. Good noise or bad noise? The role of correlated variability in a probabilistic inference framework Ralf Haefner, Jozsef Fiser, Brandeis University . . . . .	68
I-40. The origin and structure of behavioral variability in perceptual decision-making Jan Drugowitsch, Valentin Wyart, Etienne Koechlin, University of Geneva . . . . .	69
I-41. Human intracranial electrophysiology supports a heuristic model of perceptual confidence Yoshiaki Ko, Megan A. K. Peters, Hakwan Lau, Columbia University . . . . .	70
I-42. Decision-making networks using spiking neurons Jeff Orchard, Rafal Bogacz, University of Waterloo . . . . .	70
I-43. Activity of the anterior and posterior cingulate cortex during adaptive learning. Yin Li, Matthew Nassar, Joshua Gold, University of Pennsylvania . . . . .	71
I-44. Effects of D2R blockage in the dorsal striatum on perceptual inference and reinforcement learning Eunjeong Lee, Moonsang Seo, Olga Dal Monte, Bruno Averbeck, NIMH/NIH . . . . .	71

## Posters I

---

I-45. Effector-specific decision-making in the parietal cortex Jan Kubanek, Lawrence Snyder, Washington University . . . . .	72
I-46. Low frequency modulation of the beta rhythm during movement selection in posterior parietal cortex David Hawellek, Yan Wong, Bijan Pesaran, New York University . . . . .	72
I-47. Neural mechanism of active learning based on hypothesis testing Rei Akaishi, Nils Kolling, Matthew Rushworth, University of Oxford . . . . .	73
I-48. The human sense of confidence obeys properties of a statistical confidence variable Joshua Sanders, Balazs Hangya, Adam Kepecs, Cold Spring Harbor Laboratory . . . . .	73
I-49. Fronto-parietal correlated discharges vary with performance during cognitive set shifting Rodrigo Salazar, Nicolas Dotson, Charles Gray, Montana State University . . . . .	74
I-50. Data-driven algorithms for learning state dynamics in natural decision-making behaviors Rich Pang, Adrienne L Fairhall, University of Washington . . . . .	74
I-51. Uncovering stimulus-induced network dynamics during narrative comprehension Erez Simony, Chris Honey, Janice Chen, Uri Hasson, Princeton University . . . . .	75
I-52. Measuring memory formation in the human brain using cross stimulus suppression Helen Barron, Raymond Dolan, Timothy Behrens, University College London . . . . .	75
I-53. Encoding Certainty in Bump Attractors Sam Carroll, Kresimir Josic, Zachary Kilpatrick, University of Utah . . . . .	76
I-54. Stable population coding for working memory in prefrontal cortex John Murray, Nicholas Roy, Ranulfo Romo, Christos Constantinidis, Alberto Bernacchia, Xiao-Jing Wang, New York University . . . . .	76
I-55. Parametric working memory in rats: sensory vs. Prefrontal cortex Vahid Esmaeili, Arash Fassihi, Athena Akrami, Mathew Diamond, International School for Advanced Studies	77
I-56. Bayes optimal spike adaptation predicts biophysical characteristics of neurons Alessandro Ticchi, Aldo A. Faisal, Imperial College London . . . . .	78
I-57. Low-dimensional functionality of connectome dynamics: Neuro-sensory integration in <i>C. elegans</i> James Kunert, Eli Shlizerman, J. Nathan Kutz, University of Washington . . . . .	78
I-58. Low-dimensional dynamics underlie auditory cortex responses for different brain states and stimuli Marius Pachitariu, Maneesh Sahani, Nicholas Lesica, Gatsby Computational Neuroscience Unit, UCL . . . . .	79
I-59. The dimensionality of dendritic computation Lei Jin, Bardia Behabadi, Bartlett W. Mel, University of Southern California . . . . .	79
I-60. A probabilistic modeling approach for uncovering neural population rotational dynamics Shamim Nemat, Scott Linderman, Zhe Chen, Harvard University . . . . .	80
I-61. Orthogonal representation of task parameters in higher cortical areas Wieland Brendel, Dmitry Kobak, Ranulfo Romo, Claudia Feierstein, Zachary Mainen, Christian Machens, Champalimaud Neuroscience Programme . . . . .	80
I-62. Co-fluctuation patterns of population activity in macaque primary visual cortex Benjamin Cowley, Matt Smith, Adam Kohn, Byron Yu, Carnegie Mellon University . . . . .	81
I-63. Neurons in V4 show quadrature pair selectivity along curved contours Ryan Rowekamp, Tatyana Sharpee, Salk Institute for Biological Studies . . . . .	82
I-64. Transition to chaos in heterogeneous networks Johnatan Aljadeff, Merav Stern, Tatyana Sharpee, Salk Institute for Biological Studies . . . . .	82
I-65. Stability of neuronal networks with homeostatic regulation Miha Pelko, Daniel Harnack, Antoine Chaillet, Yacine Chitour, Mark van Rossum, University of Edinburgh	82
I-66. Evidence towards surround suppression in perception of complex visual properties Daniel Leeds, John Pyles, Michael Tarr, Fordham University . . . . .	83

I-67. <a href="#">Structured chaos shapes population noise entropy in driven balanced networks</a> Guillaume Lajoie, Jean-Philippe Thivierge, Eric Shea-Brown, University of Washington . . . . .	83
I-68. <a href="#">Predicting noise correlations for non-simultaneously measured neuron pairs</a> Srinivas C Turaga, Lars Buesing, Adam Packer, Henry Dalgleish, Noah Pettit, Michael Hausser, Jakob Macke, Gatsby Computational Neuroscience Unit, UCL . . . . .	84
I-69. <a href="#">Statistical complexity of neural spike trains</a> Sarah Marzen, Michael DeWeese, Jim Crutchfield, University of California, Berkeley . . . . .	84
I-70. <a href="#">Neuroembedology: dynamical relation reflects network structures</a> Satohiro Tajima, Taro Toyozumi, RIKEN Brain Science Institute . . . . .	85
I-71. <a href="#">Dynamics of excitatory-inhibitory neuronal networks with exponential synaptic weight distributions</a> Ramakrishnan Iyer, Nicholas Cain, Stefan Mihalas, Allen Institute for Brain Science . . . . .	85
I-72. <a href="#">Predictive coding in active and quiescent states</a> Veronika Koren, Sophie Deveve, David G.T. Barrett, Ecole Normale Supérieure . . . . .	86
I-73. <a href="#">A novel point process approach to the stability analysis of spiking recurrent networks</a> Farzad Farkhooi, Carl van Vreeswijk, Bernstein Center for Computational Neuroscience, Berlin . . . . .	86
I-74. <a href="#">Low-dimensional dynamics of somatosensory cortex: experiment and simulation</a> Cliff C Kerr, Marius Pachitariu, Jordan Iordanou, Joseph Francis, Maneesh Sahani, William Lytton, SUNY Downstate Medical Center . . . . .	87
I-75. <a href="#">On the brink of instability: network fragility in the epileptic cortex</a> Sridevi Sarma, Duluxan Sriharan, Johns Hopkins University . . . . .	87
I-76. <a href="#">Alpha coherence networks and the prerequisites for communication in simulated cortex</a> Stefan Berteau, Boston University . . . . .	88
I-77. <a href="#">Contextual influences on gamma oscillations in inhibition stabilized networks</a> Yashar Ahmadian, Kenneth Miller, Columbia University . . . . .	88
I-78. <a href="#">Dynamical criticality during Induction of anesthesia in human ECoG recordings</a> Leandro Alonso, Alex Proekt, Guillermo Cecchi, Marcelo Magnasco, Rockefeller University . . . . .	89
I-79. <a href="#">Modeling neural responses in the presence of unknown modulatory inputs</a> Neil C Rabinowitz, Robbe L.T. Goris, Johannes Balle, Eero Simoncelli, New York University . . . . .	89
I-80. <a href="#">The temporal structure of a random network near criticality and human ECoG dynamics</a> Rishidev Chaudhuri, Biyu He, Xiao-Jing Wang, New York University . . . . .	90
I-81. <a href="#">How brain areas can predict behavior, yet have no influence</a> Xaq Pitkow, Kaushik Lakshminarasimhan, Alexandre Pouget, Baylor College of Medicine . . . . .	91
I-82. <a href="#">Opening the black box: understanding how recurrent neural networks compute</a> David Sussillo, Omri Barak, Stanford University . . . . .	91
I-83. <a href="#">Dynamics of networks of excitatory and inhibitory units with sparse, partially symmetric couplings</a> Daniel Marti, Nicolas Brunel, Srđjan Ostojic, Group of Neural Theory, ENS . . . . .	92
I-84. <a href="#">Volume transmission as a new homeostatic mechanism</a> Yann Sweeney, Jeanette Hellgren Kotaleski, Matthias H Hennig, University of Edinburgh . . . . .	92
I-85. <a href="#">Network dynamics of spiking neurons with adaptation</a> Moritz Deger, Tilo Schwalger, Richard Naud, Wulfram Gerstner, Ecole Polytechnique Federale de Lausanne	93
I-86. <a href="#">Connection-type specific biases make random network models consistent with cortical recordings.</a> Tim Vogels, Christian Tamm, Michael Avermann, Carl Petersen, Wulfram Gerstner, University of Oxford . . . . .	93
I-87. <a href="#">Fast sampling in recurrent neural circuits</a> Mate Lengyel, Guillaume Hennequin, Laurence Aitchison, University of Cambridge . . . . .	94
I-88. <a href="#">Relating synaptic efficacies and neuronal separation in a simulated layer 2/3 cortical network</a> Daniel Miner, Jochen Triesch, Frankfurt Institute for Advanced Studies . . . . .	94

## Posters I

---

I-89. <a href="#">Encoding and learning the timing of events in neural networks</a>	
Alan Veliz-Cuba, Zachary Kilpatrick, Kresimir Josic, University of Houston . . . . .	95
I-90. <a href="#">Efficient fitting of large-scale neural models</a>	
Jascha Sohl-Dickstein, Niru Maheswaranathan, Benjamin Poole, Surya Ganguli, Stanford University . . . . .	95
I-91. <a href="#">Effects of noise injection in artificial neural networks</a>	
Benjamin Poole, Jascha Sohl-Dickstein, Surya Ganguli, Stanford University . . . . .	96
I-92. <a href="#">Unsupervised emergence of continuous tuning curves and sensory maps in spiking neural networks</a>	
Robert Guetig, Max Planck Institute of Experimental Medicine . . . . .	96
I-93. <a href="#">Learning multi-stability in plastic neural networks</a>	
Friedemann Zenke, Everton J Agnes, Wulfram Gerstner, Ecole Polytechnique Federale de Lausanne . . . . .	97
I-94. <a href="#">Interplay between short- and long-term plasticity in cell-assembly formation</a>	
Naoki Hiratani, Tomoki Fukai, RIKEN Brain Science Institute . . . . .	97
I-95. <a href="#">A Hopfield net trained on images matches retinal spike statistics and encodes images efficiently</a>	
Christopher Hillar, Sarah Marzen, Urs Koester, Kilian Koepsell, Redwood Center . . . . .	98
I-96. <a href="#">One rule to grow them all? Activity-dependent regulation for robust neural circuit specification</a>	
Timothy O’Leary, Eve Marder, Brandeis University . . . . .	98
I-97. <a href="#">The dynamics of variability in nonlinear recurrent circuits</a>	
Guillaume Hennequin, Mate Lengyel, University of Cambridge . . . . .	99
I-98. <a href="#">Specialization within populations of neural feature detectors</a>	
Julia Hillmann, Robert Guetig, Max Planck Institute of Experimental Medicine . . . . .	99
I-99. <a href="#">Role of long-range horizontal connections in visual texture classification</a>	
Felix Bauer, Matthias Kaschube, Frankfurt Institute for Advanced Studies . . . . .	100
I-100. <a href="#">The brain as social network: self-organization of a dynamic connectome</a>	
Michael Buice, Nicholas Cain, Ramakrishnan Iyer, Mike Hawrylycz, Stefan Mihalas, Allen Institute for Brain Science . . . . .	100

## Poster Session II

7:30 pm Friday 28th February

<a href="#">II-1. Examining signal and noise correlation in a recurrent network of neurons in culture</a> Jeremie Barral, Tatjana Tchumatchenko, Alexander Reyes, New York University . . . . .	101
<a href="#">II-2. Local feed-forward dynamics of scale-invariant neuronal avalanches</a> Shan Yu, Andreas Klaus, Hongdian Yang, Dietmar Plenz, National Institute of Mental Health/NIH . . . . .	101
<a href="#">II-3. Decorrelation in networks of spiking neurons: from microscopic to macroscopic magnitudes</a> Jesus Manrique, Nestor Parga, Departamento Fisica Teorica UAM . . . . .	102
<a href="#">II-4. Learning-based crossmodal suppression of ongoing activity in primary cortical areas of the awake rat</a> Jozsef Fiser, Benjamin White, Central European University . . . . .	102
<a href="#">II-5. The sign rule and beyond: boundary effects, flexibility, and noise correlations in population codes</a> Yu Hu, Joel Zylberberg, Eric Shea-Brown, University of Washington . . . . .	103
<a href="#">II-6. Noise correlations induced by computation</a> Ingmar Kanitscheider, Ruben Coen-Cagli, Alexandre Pouget, University of Geneva . . . . .	104
<a href="#">II-7. Biophysically inspired cell assembly detection at multiple scales</a> Yazan N Billeh, Michael T Schaub, Costas A Anastassiou, Mauricio Barahona, Christof Koch, California Institute of Technology . . . . .	104
<a href="#">II-8. Decoding thalamic inputs from cortical activity</a> Audrey Sederberg, Stephanie E Palmer, Jason MacLean, University of Chicago . . . . .	105
<a href="#">II-9. Topological analysis of hippocampal correlations reveals non-random, low-dimensional clique topology</a> Vladimir Itskov, Carina Curto, Eva Pastalkova, Chad Giusti, University of Nebraska, Lincoln . . . . .	105
<a href="#">II-10. A novel rate-correlation relationship might significantly improve population coding in retina</a> Joel Zylberberg, Maxwell Turner, Jon Cafaro, Fred Rieke, Eric Shea-Brown, University of Washington . . . . .	106
<a href="#">II-11. Structure of neuronal correlation during eye movement planning in FEF</a> Sanjeev Khanna, Adam Snyder, Matt Smith, University of Pittsburgh . . . . .	107
<a href="#">II-12. Spike count correlation relates to the phase and amplitude of EEG oscillations measured at the scalp</a> Adam Snyder, Cory Willis, Matt Smith, University of Pittsburgh . . . . .	107
<a href="#">II-13. Efficient coding theory in nonlinear networks with noise</a> Jerome Tubiana, Zhuo Wang, Daniel Lee, Haim Sompolinsky, Ecole Normale Superieure . . . . .	108
<a href="#">II-14. Determining the optimal distributions of neural population codes</a> Wentao Huang, Kechen Zhang, Johns Hopkins University . . . . .	108
<a href="#">II-15. Transformation of optimal motion encoding in the mouse early visual system</a> Alfred Kaye, Edward M Callaway, Tatyana Sharpee, Salk Institute for Biological Studies . . . . .	109
<a href="#">II-16. Statistical inference for directed phase coupling in neural oscillators</a> HaDi MaBouDi, Hideaki Shimazaki, Mehdi Abouzari, Shun-ichi Amari, Hamid Soltanian-Zadeh, Institute for Research in Fundamental Sciences . . . . .	109
<a href="#">II-17. Nonlinear transfer of signal and noise correlations in cortical networks</a> Dmitry Lyamzin, Samuel Barnes, Tara Keck, Nicholas Lesica, University College London . . . . .	110
<a href="#">II-18. Existence of multiple bursting oscillators in cultured neural network and transitions among them</a> June Hoan Kim, Ryouon Heo, Joon Ho Choi, Kyoung Lee, Korea University . . . . .	110
<a href="#">II-19. Extracting a signal in noise: Effect of bursts and tonic spikes on detection and discriminability</a> Clarissa Shephard, Christian Waiblinger, Cornelius Schwarz, Garrett B Stanley, Georgia Institute of Technology . . . . .	111

## Posters II

---

<a href="#">II-20. Decoding analog multiunit activity in a brain-machine interface for individuals with tetraplegia</a>	
Wasim Malik, John Donoghue, Leigh Hochberg, Harvard Medical School . . . . .	112
<a href="#">II-21. Maximum Variance Differentiation (MVD) explains the transformation from IT to Perirhinal cortex</a>	
Marino Pagan, Eero Simoncelli, Nicole Rust, University of Pennsylvania . . . . .	112
<a href="#">II-22. Predictive coding with hodgkin-huxley-type neurons</a>	
Michael Schwemmer, Sophie Deneve, Eric Shea-Brown, Adrienne L Fairhall, Ohio State University . . . . .	113
<a href="#">II-23. Zipf's law and criticality in multivariate data without fine-tuning</a>	
David Schwab, Ilya Nemenman, Pankaj Mehta, Princeton University . . . . .	114
<a href="#">II-24. Heaps' law as a tuning principle for retinal population codes</a>	
Bas van Opheusden, Gasper Tkacik, Olivier Marre, Michael Berry, New York University . . . . .	114
<a href="#">II-25. Fast population coding in recurrent networks: single cell vs. network mechanisms</a>	
Rainer Engelken, Michael Monteforte, Fred Wolf, Max Planck Institute for Dynamics and Self-Organization	115
<a href="#">II-26. Neural population coding of multiple stimuli</a>	
Emin Orhan, Wei Ji Ma, New York University . . . . .	115
<a href="#">II-27. How much signal is there in the noise?</a>	
Leon Gatys, Alexander S Ecker, Tatjana Tchumatchenko, Matthias Bethge, University of Tuebingen . . . . .	116
<a href="#">II-28. Neuronal population decoding can account for perceptual lightness illusions</a>	
Douglas A Ruff, David H H Brainard, Marlene Cohen, University of Pittsburgh . . . . .	116
<a href="#">II-29. Adaptation of intrinsic circuit correlations and synergy in retinal synchrony</a>	
Ann Hermundstad, Olivier Marre, Stephanie E Palmer, Michael Berry, Gasper Tkacik, Vijay Balasubramanian, University of Pennsylvania . . . . .	117
<a href="#">II-30. A data-driven model-free population receptive field estimation framework</a>	
Francesca Pizzorni-Ferrarese, Jonas Larsson, University of London . . . . .	118
<a href="#">II-31. Robustness of criticality under different stimulus conditions</a>	
Mark Ioffe, Gasper Tkacik, William Bialek, Michael Berry, Princeton University . . . . .	118
<a href="#">II-32. Mapping large scale retinal population activity with high density multielectrode arrays</a>	
Sahar Pirmoradian, Gerrit Hilgen, Oliver Muthman, Alessandro Maccione, Upinder Bhalla, Luca Berdonini, Evelyne Sernagor, Matthias H Hennig, University of Edinburgh . . . . .	119
<a href="#">II-33. Latent variable models (almost) always display signatures of criticality</a>	
Laurence Aitchison, Nicola Corradi, Peter Latham, Gatsby Computational Neuroscience Unit, UCL . . . . .	119
<a href="#">II-34. Generation of sequences through reconfiguration of ongoing activity in neural networks: A model of c</a>	
Kanaka Rajan, Christopher Harvey, David Tank, Princeton University . . . . .	120
<a href="#">II-35. A method for fitting linear neural models to extracellular electrophysiology data</a>	
Corinne Teeter, Christof Koch, Stefan Mihalas, Allen Institute for Brain Science . . . . .	121
<a href="#">II-36. Direct and indirect pathways of the basal ganglia exert opposing but not identical control of motor</a>	
Ian Oldenburg, Bernardo Sabatini, Harvard University . . . . .	121
<a href="#">II-37. Spatial learning by Drosophila melanogaster in virtual reality</a>	
TJ Florence, Charles Zuker, Michael Reiser, Columbia University . . . . .	122
<a href="#">II-38. Redundant hippocampal spatial coding offsets competition between interference and generalization</a>	
Alexander Keinath, Joshua T Dudman, Isabel Muzzio, University of Pennsylvania . . . . .	122
<a href="#">II-39. Phase comparison of functionally distinct inputs in CA1 pyramidal cells during locomotion</a>	
Katie Bittner, Jeffrey Magee, Janelia Farm Research Campus . . . . .	123
<a href="#">II-40. Representation and generalization in Drosophila olfactory learning</a>	
Ann Kennedy, Peter Wang, Daisuke Hattori, Richard Axel, L.F. Abbott, Columbia University . . . . .	123



[II-41. The functional role of randomness in olfactory processing](#)  
 Kamesh Krishnamurthy, Ann Hermundstad, Thierry Mora, Aleksandra Walczak, Venkatesh N Murthy, Charles Stevens, Vijay Balasubramanian, University of Pennsylvania . . . . . 124

[II-42. Two types of representation of sensory information in cortical feedback axons in awake mice](#)  
 Gonzalo Otazu, Hong Goo Chae, Dinu Albeanu, Cold Spring Harbor Laboratory . . . . . 125

[II-43. Long-distance inhibition by short axon cells gate the output of the olfactory bulb.](#)  
 Arkarup Bandyopadhyay, Fred Marbach, Francesca Anselmi, Matthew Koh, Martin B. Davis, Hassana K. Oyibo, Priyanka Gupta, Dinu Albeanu, Cold Spring Harbor Laboratory . . . . . 125

[II-44. Learning hierarchical representations of natural auditory scenes](#)  
 Wiktor Mlynarski, Max Planck Institute for Mathematics in the Sciences . . . . . 126

[II-45. Invariant selectivity of auditory neurons due to predictive coding](#)  
 Izzet Yildiz, Brian Fischer, Sophie Deneve, Ecole Normale Superieure . . . . . 126

[II-46. Bayesian perception of the pitch of non-stationary natural sounds](#)  
 Vincent Adam, Maneesh Sahani, Gatsby Computational Neuroscience Unit, UCL . . . . . 127

[II-47. Learning noise-invariant auditory receptive fields for speech recognition](#)  
 Phillip Schafer, Dezhe Jin, Pennsylvania State University . . . . . 127

[II-48. Sensitive periods for functional connectivity in auditory cortical circuits](#)  
 Liberty Hamilton, Jascha Sohl-Dickstein, Shaowen Bao, University of California, Berkeley . . . . . 128

[II-49. Population coding of complex sounds in the avian auditory system](#)  
 Mike Schachter, Tyler Lee, Julie Elie, Friedrich Sommer, Frederic Theunissen, University of California, Berkeley . . . . . 128

[II-50. Changes in encoding of communication signals by populations of neurons in the auditory cortex](#)  
 Isaac Carruthers, Ryan Natan, Andrew Jaegle, Laetitia Mwilambwe-Tshilobo, Diego Laplagne, Maria Geffen, University of Pennsylvania . . . . . 129

[II-51. Clustering of cortical responses to natural sounds reveals organization for pitch and speech](#)  
 Samuel Norman-Haignere, Joshua McDermott, Nancy Kanwisher, Massachusetts Institute of Technology 130

[II-52. Characterizing central auditory neurons in \*Drosophila\*](#)  
 Allison Baker, Alex Vaughan, Bruce Baker, Rachel Wilson, Harvard Medical School . . . . . 130

[II-53. Synaptic computation of interaural level differences in rat auditory cortex](#)  
 Michael Kyweriga, Whitney Stewart, Carolyn Cahill, Michael Wehr, University of Oregon . . . . . 131

[II-54. Near-optimallinear decoding for marginalization: application to heading estimation](#)  
 HyungGoo R Kim, Xaq Pitkow, Ryo Sasaki, Dora E. Angelaki, Gregory C. DeAngelis, University of Rochester 131

[II-55. Integration of visual and tactile signals in behaving rats](#)  
 Nader Nikbakht, Azadeh Tafreshiha, Rodrigo Q. Quiroga, Davide Zoccolan, Mathew Diamond, International School for Advanced Studies . . . . . 132

[II-56. 3D surface tilt estimation in natural scenes from image cue gradients](#)  
 Johannes Burge, Brian McCann, Wilson Geisler, University of Texas at Austin . . . . . 132

[II-57. A new database for natural motion: movie data and analysis of natural motion statistics](#)  
 Stephanie E Palmer, Jared Salisbury, Heather Yee, Hanna Torrence, Don Ho, David Schwab, University of Chicago . . . . . 133

[II-58. Higher-order stimulus correlations drive a sparse representation in mouse V1 well-suited for discrim](#)  
 Emmanouil Froudarakis, Philipp Berens, R. James Cotton, Alexander S Ecker, Fabian H. Sinz, Dimitri Yatsenko, Peter Saggau, Matthias Bethge, Andreas S. Tolias, Baylor College of Medicine . . . . . 134

[II-59. Representation and reconstruction of natural scenes in the primate retina](#)  
 Alexander Heitman, Martin Greschner, Greg Field, Peter Li, Daniel Ahn, Alexander Sher, Alan Litke, E. J. Chichilnisky, University of California, San Diego . . . . . 134

[II-60. A selective color-opponent pathway in the mouse retina](#)  
 Maximilian Joesch, Markus Meister, Harvard University . . . . . 135

II-61. [Early representations of whisker self-motion revealed using closed-loop control in anesthetized rats](#)  
Avner Wallach, Knarik Bagdasarian, Ehud Ahissar, Weizmann Institute of Science . . . . . 135

II-62. [An optimal object-tracking model provides a unifying account of motion and position perception](#)  
Oh-sang Kwon, Duje Tadin, David Knill, University of Rochester . . . . . 136

II-63. [Binding by synchrony in complex-valued deep neural networks](#)  
David P Reichert, Thomas Serre, Brown University . . . . . 136

II-64. [Interplay between sensory statistics and neural circuit constraints in visual motion estimation](#)  
James E Fitzgerald, Damon Clark, Harvard University . . . . . 137

II-65. [Selectivity for shape and blur in area V4](#)  
Timothy Oleskiw, Amy Nowack, Anitha Pasupathy, University of Washington . . . . . 138

II-66. [Decoding what types of information are in the macaque face patch system](#)  
Ethan Meyers, Mia Borzello, Winrich Freiwald, Doris Tsao, Massachusetts Institute of Technology . . . . . 138

II-67. [Efficient coding can account for the existence of sustained and transient bipolar cell types](#)  
Tao Hu, Dmitri B Chklovskii, Texas A&M University . . . . . 139

II-68. [Characterizing retinal network adaptation in Drosophila using models of mean and contrast adaptation](#)  
Daniel Coca, Carlos Luna Ortiz, Stephen Billings, Mikko Juusola, University of Sheffield . . . . . 139

II-69. [It cortex contains a general-purpose visual object representation](#)  
Ha Hong, Daniel Yamins, Najib Majaj, James DiCarlo, Massachusetts Institute of Technology . . . . . 140

II-70. [Interaction of color, motion and contrast in macaque V4 population activity and behavior](#)  
Ivo Vanzetta, Maria Bermudez, Thomas Deneux, Frederic Barthelemy, Guillaume S Masson, INT - CNRS UMR7289 . . . . . 141

II-71. [Evolutionary convergence in computation of local motion signals in monkey and dragonfly](#)  
Eyal Nitzany, Gil Menda, Paul Shamble, James R Golden, Ron Hoy, Jonathan Victor, Cornell University . . . . . 141

II-72. [Pop-out visual search of moving targets in the archer fish](#)  
Mor Ben-Tov, Opher Donchin, Ohad Ben-Shahar, Ronen Segev, Ben Gurion University . . . . . 142

II-73. [Bayesian inference for low rank receptive fields](#)  
Mijung Park, Jonathan W Pillow, University of Texas at Austin . . . . . 142

II-74. [Lateral feature-level interactions in anterior visual cortex of monkeys and humans](#)  
Chou Hung, Chang Mao Chao, Li-feng Yeh, Yueh-peng Chen, Chia-pei Lin, Yu-chun Hsu, Georgetown University Medical Center . . . . . 143

II-75. [Quantifying and modeling the emergence of object recognition in the ventral stream](#)  
Darren Seibert, Daniel Yamins, Ha Hong, James DiCarlo, Justin Gardner, Massachusetts Institute of Technology . . . . . 144

II-76. [Using deep networks to learn non-linear receptive field models of V2 neurons](#)  
Michael Oliver, Jack Gallant, University of California, Berkeley . . . . . 144

II-77. [Four-layer sparse coding model of natural images that reproduces tuning properties in V1, V2 and V4](#)  
Haruo Hosoya, Aapo Hyvarinen, ATR . . . . . 145

II-78. [Low-dimensional structure of neural correlations in cortical microcircuits](#)  
Dimitri Yatsenko, Kresimir Josic, Alexander S Ecker, Emmanouil Froudarakis, R. James Cotton, Andreas S. Tolias, Baylor College of Medicine . . . . . 145

II-79. [Intrinsic networks in visual cortex are mediated by local and inter-areal rhythmic synchronization](#)  
Chris Lewis, Conrado Bosman, Thilo Womelsdorf, Pascal Fries, Ernst Strungmann Institute . . . . . 146

II-80. [Spectral noise shaping: a multipurpose coding strategy in dorsal and ventral visual pathways](#)  
Grant Mulliken, Robert Desimone, Massachusetts Institute of Technology . . . . . 146

[II-81. Spatiotemporal integration of optic flow and running speed in V1](#)  
 Maneesh Sahani, Marius Pachitariu, Adil Khan, Jasper Poort, Ivana Orsolich, Georg Keller, Sonja Hofer, Thomas Mrsic-Flogel, Gatsby Computational Neuroscience Unit, UCL . . . . . 147

[II-82. Learning related changes of stimulus representations in mouse visual cortex](#)  
 Jasper Poort, Adil Khan, Marius Pachitariu, Abdellatif Nemri, Ivana Orsolich, Julija Krupic, Marius Bauza, Georg Keller, Maneesh Sahani, Thomas Mrsic-Flogel, Sonja Hofer, University College London . . . . . 148

[II-83. Relationship between natural image statistics and lateral connectivity in the primary visual cortex](#)  
 Philipp Rudiger, Jean-Luc Stevens, Bharath Chandra Talluri, Laurent Perrinet, James Bednar, University of Edinburgh . . . . . 148

[II-84. Can retinal ganglion cell dipoles seed iso-orientation domains in the visual cortex?](#)  
 Manuel Schottdorf, Stephen J. Eglon, Fred Wolf, Wolfgang Keil, Max Planck Institute for Dynamics and Self-Organization . . . . . 149

[II-85. Cortical maps of ON and OFF visual responses are the substrate of the cortical orientation map](#)  
 Jens Kremkow, Jianzhong Jin, Stanley J Komban, Yushi Wang, Reza Lashgari, Jose-Manuel Alonso, State University of New York . . . . . 150

[II-86. V1 laminar-specific activity patterns evoked by receptive field, near and far surround stimulation](#)  
 Maryam Bijanzadeh, Lauri Nurminen, Jennifer Ichida, Sam Merlin, Kenneth Miller, Alessandra Angelucci, Moran Eye center . . . . . 150

[II-87. Stereoscopic surface disambiguation and interpolation in the macaque primary visual cortex](#)  
 Jason M Samonds, Brian Potetz, Christopher Tyler, Tai Sing Lee, Carnegie Mellon University . . . . . 151

[II-88. Visuotopic anisotropy in functional connectivity among local cortical sites in human visual areas](#)  
 Jungwon Ryu, Sang-Hun Lee, Seoul National University . . . . . 151

[II-89. Visualizing the similarity and pedigree of neuronal ion channel models for NEURON.](#)  
 William Podlaski, Rajnish Ranjan, Alexander Seeholzer, Henry Markram, Wulfram Gerstner, Tim Vogels, Ecole Polytechnique Federale de Lausanne . . . . . 152

[II-90. Scalable semi-supervised framework for activity mapping of the leech nervous system](#)  
 Edward Frady, William Kristan, University of California, San Diego . . . . . 153

[II-91. The cellular origin of neuronal avalanches](#)  
 Tim Bellay, Andreas Klaus, Saurav Seshadri, Dietmar Plenz, National Institutes of Health . . . . . 153

[II-92. Microcircuit activity is patterned topologically and reveals features of underlying connectivity](#)  
 Brendan Chambers, Alexander Sadovsky, Jason MacLean, University of Chicago . . . . . 154

[II-93. Scalable ROI detection for calcium imaging](#)  
 David Pfau, Jeremy Freeman, Misha Ahrens, Liam Paninski, Columbia University . . . . . 154

[II-94. Investigating mechanisms of drug action at the hippocampal ensemble level in behaving mice.](#)  
 Tamara Berdyeva, Stephani Otte, Leah Aluisio, Laurie Burns, Yaniv Ziv, Christine Dugovic, Kunal Gosh, Mark J Schnitzer, Tim Lovenberg, Pascal Bonaventure, Janssen LLC . . . . . 155

[II-95. Phase coding in posterior parietal cortex predicts neural suppression in a gaze-anchoring task](#)  
 Maureen Hagan, Bijan Pesaran, New York University . . . . . 156

[II-96. PV+ inhibition regulates feedforward activity in auditory cortex.](#)  
 Alexandra Moore, Michael Wehr, University of Oregon . . . . . 156

[II-97. Observing the unobservable: a note on stop-signal reaction-time distributions](#)  
 Tobias Teichert, Vincent Ferrera, University of Pittsburgh . . . . . 157

[II-98. A simplified Bayesian model accounts for procedure-related variability in perceptual thresholds](#)  
 Ofri Raviv, Yonatan Loewenstein, Merav Ahissar, Hebrew University . . . . . 157

[II-99. A behavioral tracking paradigm for estimating visual sensitivity with dynamic internal models](#)  
 Kathryn Bonnen, Johannes Burge, Jacob L Yates, Jonathan W Pillow, Lawrence K Cormack, University of Texas at Austin . . . . . 158

**Posters II**

---

[II-100. Modeling and exploiting co-adaptation in brain-computer interfaces](#)  
Josh Merel, Antony Qian, Roy Fox, Liam Paninski, Columbia University . . . . . 158

## Poster Session III

7:30 pm Saturday 1st March

<a href="#">III-1. Behavioral and neural evidence that mice infer probabilistic models for timing</a> Yi Li, Joshua T Dudman, Janelia Farm Research Campus . . . . .	159
<a href="#">III-2. High-frequency activity in the medial orbitofrontal cortex encodes positive outcomes</a> Ashwin Ramayya, Ashwini Sharan, James Evans, Michael Sperling, Timothy Lucas, Michael Kahana, University of Pennsylvania . . . . .	159
<a href="#">III-3. Altered neuronal avalanches at the population and cell level in a rat model of schizophrenia</a> Saurav Seshadri, Tim Bellay, Samantha S Chou, Dietmar Plenz, National Institute of Mental Health . . . . .	160
<a href="#">III-4. Modeling attention-induced drop of noise correlations by inhibitory feedback</a> Tatjana Wiese, Marlene Cohen, Brent Doiron, Carnegie Mellon University . . . . .	160
<a href="#">III-5. A single net excitation strength measure broadens the scope of the 2-layer model</a> Lei Jin, Bartlett W. Mel, University of Southern California . . . . .	161
<a href="#">III-6. Low-dimensional models of neural population recordings with complex stimulus selectivity</a> Evan Archer, Jonathan W Pillow, Jakob Macke, Max Planck Institute . . . . .	162
<a href="#">III-7. Scalable nonparametric models for binary spike patterns</a> Il Memming Park, Evan Archer, Kenneth Latimer, Jonathan W Pillow, University of Texas at Austin . . . . .	162
<a href="#">III-8. High frequency deep brain stimulation elicits neural restoration via loop-based reinforcement</a> Sabato Santaniello, Erwin B. Montgomery, John T. Gale, Sridevi Sarma, Johns Hopkins University . . . . .	163
<a href="#">III-9. Theory for the firing statistics of neurons driven by non-Poissonian synaptic input</a> Tilo Schwalger, Benjamin Lindner, Ecole Polytechnique Federale de Lausanne . . . . .	163
<a href="#">III-10. Error correction in the retinal population code</a> Jason Prentice, Olivier Marre, Mark Ioffe, Gasper Tkacik, Michael Berry, Princeton University . . . . .	164
<a href="#">III-11. The nature of cortical variability</a> I-Chun Lin, Michael Okun, Matteo Carandini, Kenneth Harris, University College London . . . . .	165
<a href="#">III-12. From sounds to meaning: neural representation of calls in the avian auditory cortex</a> Julie Elie, Frederic Theunissen, University of California, Berkeley . . . . .	165
<a href="#">III-13. Stimulus-specific adaptation in the auditory system — models and data</a> Israel Nelken, Leila Khouri, Bshara Awwad, Itai Hershenhoren, Amit Yaron, Tohar Yarden, Hebrew University . . . . .	166
<a href="#">III-14. Modeling sound pulse counting with phasic inhibition and nonlinear integration</a> Richard Naud, Dave Houtman, Gary Rose, Andre Longtin, University of Ottawa . . . . .	166
<a href="#">III-15. How do insects see motion? Using connectomics, physiology and behavior to constrain possible models</a> Arjun Bharioke, Dmitri B Chklovskii, Janelia Farm Research Campus . . . . .	167
<a href="#">III-16. Predicting IT and V4 neural responses with performance-optimized neural networks</a> Daniel Yamins, Ha Hong, Darren Seibert, James DiCarlo, Massachusetts Institute of Technology . . . . .	167
<a href="#">III-17. Adaptation in V1 populations reflects optimal inference of natural movie statistics</a> Michael Snow, Ruben Coen-Cagli, Odelia Schwartz, Albert Einstein College of Medicine . . . . .	168
<a href="#">III-18. Heisenberg: a tool for mapping complex cell receptive fields</a> Urs Koester, Bruno Olshausen, Redwood Center . . . . .	168
<a href="#">III-19. Are spikes unitary signals? How axonal noise produces synaptic variability</a> Ali Neishabouri, Aldo A. Faisal, Imperial College London . . . . .	169
<a href="#">III-20. Optimal population codes with limited input information have finite tuning-curve widths</a> Guillaume Dehaene, Jeffrey Beck, Alexandre Pouget, University of Geneva . . . . .	169
<a href="#">III-21. Optimal compensation for neuron death</a> David G.T. Barrett, Sophie Deneve, Christian Machens, Ecole Normale Supérieure . . . . .	170

## Posters III

---

III-22. Optimal sensory coding by populations of ON and OFF neurons Julijana Gjorgjieva, Markus Meister, Haim Sompolinsky, Harvard University . . . . .	170
III-23. Distinguishing different efficient coding principles from tuning curve statistics Zhuo Wang, Alan Stocker, Daniel Lee, University of Pennsylvania . . . . .	171
III-24. Second order phase transition describes maximally informative encoding in the retina David Kastner, Stephen A Baccus, Tatyana Sharpee, Ecole Polytechnique Federale de Lausanne . . . . .	172
III-25. Sensor placement for sparse sensory decision-making Bingni Brunton, Annika Eberle, Bradley Dickerson, Steven Brunton, J. Nathan Kutz, Tom Daniel, University of Washington . . . . .	172
III-26. The role of environmental feedback in transitions from passive to active brain processing Christopher Buckley, Taro Toyozumi, RIKEN Brain Science Institute . . . . .	173
III-27. Calcium-dependence of vesicle replenishment at photoreceptor synapses Caitlyn Parmelee, Carina Curto, University of Nebraska, Lincoln . . . . .	173
III-28. Uncertainty modulates timing of neural interaction in PMd during reach planning Pavan Ramkumar, Brian Dekleva, Paul Wanda, Hugo Fernandes, Lee E Miller, Konrad Koerding, Rehabilitation Institute of Chicago . . . . .	174
III-29. Bayesian inference of neural connectivity via approximate message passing Alyson Fletcher, Sundeeep Rangan, University of California, Santa Cruz . . . . .	175
III-30. Reconstructing functional connectivity in complete neural population by randomized sparse sampling Yuriy Mishchenko, Toros University . . . . .	175
III-31. Classical conditioning via inference over observable situation contexts Nisheeth Srivastava, Paul Schrater, University of Minnesota . . . . .	176
III-32. Tractable, optimal high dimensional Bayesian inference Madhu Advani, Surya Ganguli, Stanford University . . . . .	176
III-33. Short-term facilitation as a normative consequence of presynaptic spike-rate adaptation Simone Carlo Surace, Jean-Pascal Pfister, University of Bern . . . . .	177
III-34. Learning, inference, and replay of hidden state sequences in recurrent spiking neural networks Dane Corneil, Emre Neftci, Giacomo Indiveri, Michael Pfeiffer, Ecole Polytechnique Federale de Lausanne . . . . .	177
III-35. Unsupervised identification of excitatory and inhibitory populations from multi-cell recordings Lars Buesing, Gatsby Computational Neuroscience Unit, UCL . . . . .	178
III-36. The size-weight illusion is not anti-Bayesian after all Megan A. K. Peters, Ladan Shams, University of California, Los Angeles . . . . .	178
III-37. How well can saliency maps tell where people look Matthias Kummerer, Matthias Bethge, University of Tuebingen . . . . .	179
III-38. Sampling orientations at different contrasts Agnieszka Grabska-Barwinska, Peter Latham, Gatsby Computational Neuroscience Unit, UCL . . . . .	179
III-39. Sampling methods for spike inference from calcium imaging data Eftychios Pnevmatikakis, Josh Merel, Liam Paninski, Columbia University . . . . .	180
III-40. Predictive balance of cortical membrane potentials by STDP of inhibition Christian Albers, Maren Westkott, Klaus Pawelzik, University of Bremen . . . . .	180
III-41. Comparing roles of feedback and feedforward inhibition in sparsening of sensory codes Pavel Sanda, Tiffany Kee, Nitin Gupta, Mark Stopfer, Maxim Bazhenov, University of California, Riverside . . . . .	181
III-42. Persistent activity through multiple mechanisms in a spiking network that solves DMS tasks Jason Fleischer, The Neurosciences Institute . . . . .	181

III-43. Oscillations emerging from noise-driven networks with electrical synapses and subthreshold resonance	
Claudia Clopath, Tatjana Tchumatchenko, Imperial College London . . . . .	182
III-44. Self-supervised learning of neural integrator networks	
Erik Nygren, Robert Urbanczik, Walter Senn, University of Bern . . . . .	182
III-45. Nonlinear threshold dynamics implements multiplicative adaptation in pyramidal neurons	
Christian Pozzorini, Skander Mensi, Oliver Hagens, Wolfram Gerstner, Ecole Polytechnique Federale de Lausanne . . . . .	183
III-46. Asymmetry of electrical synapses: measurement and impact in small networks	
Julie Haas, Joe C Brague, Jessica Sevetson, Lehigh University . . . . .	183
III-47. Balanced networks in space	
Robert Rosenbaum, Jonathan Rubin, Brent Doiron, University of Pittsburgh . . . . .	184
III-48. Derivation of leaky integrate-and-fire and STDP from a signal processing cost function	
Dmitri B Chklovskii, Tao Hu, Janelia Farm Research Campus . . . . .	184
III-49. Summation and phase decorrelation of odor responses in an olfactory bulb microcircuit model	
Aditya Gilra, Upinder Bhalla, National Centre for Biological Sciences . . . . .	185
III-50. Dynamics of olfactory bulb output generates concentration-invariant cortical odor representations	
Merav Stern, L.F. Abbott, Kevin M. Franks, Columbia University . . . . .	186
III-51. Inferring synaptic conductances from spike trains with a point process encoding model	
Kenneth Latimer, E. J. Chichilnisky, Fred Rieke, Jonathan W Pillow, University of Texas at Austin . . . . .	186
III-52. Emergent connectivity in a spiking model of visual cortex with attentional feedback	
Philip Meier, Dimitry Fisher, Jayram Moorkanikara Nageswaran, Botond Szatmary, Eugene Izhikevich, Brain Corporation . . . . .	187
III-53. A spiking neural network for autoencoder learning	
Kendra Burbank, University of Chicago . . . . .	187
III-54. Evidence for critical slowing down in neurons near spiking threshold	
Christian Meisel, Andreas Klaus, Christian Kuehn, Dietmar Plenz, National Institute of Mental Health . . . . .	188
III-55. Self-organization in balanced state networks by STDP and homeostatic plasticity	
Felix Effenberger, Anna Levina, Jurgen Jost, Max Planck Institute for Mathematics in the Sciences . . . . .	189
III-56. Predicting in vivo statistics from in vitro cell parameters in a detailed prefrontal cortex model	
Joachim Hass, Loreen Hertaeg, Tatiana Golovko, Daniel Durstewitz, Bernstein Center for Computational Neuroscience, Heidelberg-Mannheim . . . . .	189
III-57. Firing-rate dynamics from spiking network models with low-rank connectivity	
Brian DePasquale, Mark Churchland, L.F. Abbott, Columbia University . . . . .	190
III-58. CA1 activity sequences emerge after reorganization of network correlation structure during learning	
Mehrab Modi, Ashesh Dhawale, Upinder Bhalla, National Centre for Biological Sciences . . . . .	190
III-59. Input-driven unsupervised learning in recurrent neural networks	
Alireza Alemi-Neissi, Carlo Baldassi, Nicolas Brunel, Riccardo Zecchina, Human Genetics Foundation . . . . .	191
III-60. Normalized Hebbian learning develops both simple and complex receptive fields from naturalistic vide	
Thomas Miconi, Jedediah Singer, Gabriel Kreiman, The Neurosciences Institute . . . . .	192
III-61. A normative framework of single-neuron computation yields Oja-like learning with realistic features	
Cengiz Pehlevan, Tao Hu, Zaid J. Towfic, Dmitri B Chklovskii, Janelia Farm Research Campus . . . . .	192
III-62. Sequence generation by spatio-temporal cycles of inhibition	
Jonathan Cannon, Jeffrey Markowitz, Nancy Kopell, Timothy Gardner, Boston University . . . . .	193
III-63. Adaptive shaping of feature selectivity in the awake rodent vibrissa system	
He Zheng, Douglas Ollerenshaw, Daniel Millard, Qi Wang, Garrett B Stanley, Georgia Institute of Technology	193

III-64. Computational model of optimal changes in auditory cortical receptive fields during behavioral tasks Michael Carlin, Mounya Elhilali, Johns Hopkins University . . . . .	194
III-65. Graded memories in balanced attractor networks Dylan Festa, Guillaume Hennequin, Mate Lengyel, University of Cambridge . . . . .	194
III-66. Interplay between learning-rate control and uncertainty minimization during one-shot causal learning Sang Wan Lee, John O'Doherty, Shinsuke Shimojo, California Institute of Technology . . . . .	195
III-67. A self-consistent theory for STDP and spiking correlations in neuronal networks. Gabriel K Ocker, Ashok Litwin-Kumar, Brent Doiron, University of Pittsburgh . . . . .	196
III-68. Generalization of memories by spatial patterning of protein synthesis Cian O'Donnell, Terrence Sejnowski, Salk Institute for Biological Studies . . . . .	196
III-69. A normative theory of structural and intrinsic plasticity in dendrites Balazs Ujfalussy, Mate Lengyel, University of Cambridge . . . . .	197
III-70. Contribution of cerebellar Golgi cells to learned motor timing during the vestibulo-ocular reflex Thomas Chartrand, Grace Q. Zhao, Jennifer L. Raymond, Mark Goldman, University of California, Davis . . . . .	197
III-71. Stochastic synchronization in the Purkinje cells of the cerebellum Sergio Verduzco-Flores, University of Colorado at Boulder . . . . .	198
III-72. Two forms of information found in local field potentials of the hippocampus Gautam Agarwal, Kenji Mizuseki, Antal Berenyi, Gyorgy Buzsaki, Friedrich Sommer, Redwood Center for Theoretical Neuroscience . . . . .	198
III-73. Cluster analysis of sharp-wave ripple field potential signatures in the macaque hippocampus Juan Ramirez-Villegas, Nikos Logothetis, Michel Besserve, Max Planck Institute for Biological Cybernetics . . . . .	199
III-74. Priors formed over multiple time scales facilitate performance in a visual search task Nikos Gekas, Aaron Seitz, Peggy Series, University of Edinburgh . . . . .	200
III-75. A multi-step decision task elicits planning behavior in rats Kevin Miller, Jeffrey Erlich, Charles Kopec, Matthew Botvinick, Carlos Brody, Princeton University . . . . .	200
III-76. Dissociable systems for reward and content updating revealed using cross stimulus suppression Erie Boorman, Vani Rajendran, Timothy Behrens, University of Oxford . . . . .	201
III-77. Working memory contributions to reinforcement learning impairments in schizophrenia. Anne Collins, James Gold, James Waltz, Michael Frank, Brown University . . . . .	201
III-78. Temporal stimulus segmentation and decision making via spike-based reinforcement learning Giancarlo La Camera, Robert Urbanczik, Walter Senn, State University of New York at Stony Brook . . . . .	202
III-79. A link between corticostriatal plasticity and risk taking in humans David Arkadir, Angela Radulescu, Deborah Raymond, Susan B Bressman, Pietro Mazzoni, Yael Niv, Columbia Neurological Institute . . . . .	202
III-80. Kalman filter modeling of reinforced behavior: Application to pitch and tempo learning in songbirds Anja T. Zai, Florian Blattler, Anna E. Stepien, Alessandro Canopoli, Richard H.R. Hahnloser, ETH Zurich . . . . .	203
III-81. A role for habenula in dopamine neurons' responses to reward omission Ju Tian, Naoshige Uchida, Harvard University . . . . .	204
III-82. The interaction between learning and mood Eran Eldar, Yael Niv, Princeton University . . . . .	204
III-83. Striatal neuronal activity of macaque tracks local and global changes during natural habit formation Theresa Desrochers, Ken-ichi Amemori, Ann Graybiel, Brown University . . . . .	205
III-84. A signal transduction model for adenosine-induced indirect corticostriatal plasticity Naoto Yukinawa, Kenji Doya, Junichiro Yoshimoto, OIST Graduate University . . . . .	205
III-85. In vivo two-photon imaging of cortico-cerebellar mossy fiber synapses: functional properties Daria Rylkova, Ye Zengyou, David Linden, Johns Hopkins University . . . . .	206



III-86. [Efficient coding of visual motion signals in the smooth pursuit system](#)  
 Leslie Osborne, Bing Liu, University of Chicago . . . . . 206

III-87. [Corresponding neural signatures of movement and imagined movement in human motor cortex](#)  
 Chethan Pandarinath, Vikash Gilja, Christine Blabe, Leigh Hochberg, Krishna Shenoy, Jaimie Henderson, Stanford University . . . . . 207

III-88. [Towards mapping the genesis of behavior in a complete nervous system](#)  
 Alexander Katsov, Cornelia Bargmann, Rockefeller University . . . . . 208

III-89. [Precise temporal encoding in a vocal motor system](#)  
 Claire Tang, Kyle Srivastava, Diala Chehayeb, Ilya Nemenman, Samuel Sober, University of California, San Francisco . . . . . 208

III-90. [Conserved and independent population activity between arm and Brain-Machine Interface control](#)  
 Hagai Lalazar, L.F. Abbott, Eilon Vaadia, Columbia University . . . . . 209

III-91. [Single-trial prediction of slow or fast reaction from MEG brain activity before movement](#)  
 Ryu Ohata, Kenji Ogawa, Hiroshi Imamizu, Cognitive Mechanisms Laboratories, ATR . . . . . 209

III-92. [Neuronal network dynamics within and between frontal and parietal cortex in the monkey](#)  
 Benjamin Wellner, Steve Suway, Hans Scherberger, German Primate Center . . . . . 210

III-93. [Motor intention is marked by LFP beta amplitude suppression in a person with paralysis using a BCI](#)  
 Tomislav Milekovic, Beata Jarosiewicz, Anish Sarma, Leigh Hochberg, John Donoghue, Brown University 210

III-94. [Unveiling neural code with a low dimensional model of birdsong production](#)  
 Ana Amador, Yonatan Sanz Perl, Gabriel B. Mindlin, Daniel Margoliash, University of Buenos Aires . . . . 211

III-95. [Stereotyped, distributed networks revealed by whole-brain activity maps in behaving zebrafish](#)  
 Claudia Feierstein, Ruben Portugues, Florian Engert, Michael Orger, Champalimaud Neuroscience Programme . . . . . 212

III-96. [How do rats and humans use sensory information to time their actions?](#)  
 Adina Drumea, Mathew Diamond, International School for Advanced Studies . . . . . 212

III-97. [Dynamic motor-related input to visual interneurons in Drosophila](#)  
 Anmo Kim, Gaby Maimon, Rockefeller University . . . . . 213

III-98. [Motor cortical activity predicts behavioral corrections following last-minute goal changes](#)  
 Katherine Ames, Krishna Shenoy, Stanford University . . . . . 213

III-99. [Dynamic range adaptation in motor cortical neurons](#)  
 Robert Rasmussen, Andrew Schwartz, Steven Chase, University of Pittsburgh . . . . . 214

III-100. [Pushing in the wrong direction: optogenetic perturbation misaligns with motor cortical dynamics](#)  
 Daniel O'Shea, Werapong Goo, Paul Kalanithi, Ilka Diester, Karl Deisseroth, Krishna Shenoy, Stanford University . . . . . 214



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

This PDF can be downloaded at: [http://cosyne.org/cosyne14/Cosyne2014\\_program\\_book.pdf](http://cosyne.org/cosyne14/Cosyne2014_program_book.pdf)





# Abstracts

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

## T-1. Circuits and strategies for skilled reach: From internal copies to gain control

Thomas Jessell  
Columbia University

The planning and execution of forelimb movement represents one of the more refined attributes of the mammalian motor system, permitting animals to reach for, and manipulate objects with dexterity and precision. Reaching movements rely on both descending commands and sensory feedback, although insight into the spinal circuitry and strategies used to integrate these inputs remains sketchy. Motor commands require validation of the accuracy of the initial plan, yet sensory feedback is burdened by a temporal delay that limits its role in the on-line refinement of movement. Moreover the broad dynamic range inherent in sensory feedback demands an effective means of scaling sensory gain to motor output. We have explored the organization and function of spinal circuits assigned these two challenges through a combination of genetics, physiology, anatomy and theory in the mouse. Internal copy pathways permit timely updating of motor output, but their intricate association with motor circuitry has precluded evaluation of their impact on motor performance. Cervical propriospinal interneurons (PNs) have the potential to convey an internal copy of pre-motor signals by virtue of their dual excitatory innervation of forelimb motor neurons and pre-cerebellar neurons in the lateral reticular nucleus. We find that PNs comprise a genetically-accessible subset of cervical V2a excitatory interneurons and show that targeted ablation of this neuronal set selectively perturbs reaching. More tellingly, optogenetic activation of the PN internal copy projection recruits a fast cerebellar feedback loop that modulates forelimb motor neuron activity and impairs reaching kinematics. These findings place PNs at the core of an internal copy pathway assigned to the rapid updating of motor output during reaching behavior. Models of feedback control have argued that the stability of limb movement requires tight constraints on the gain of proprioceptive sensory input. In principle, presynaptic inhibition at sensory-motor synapses offers an effective means of regulating sensory gain, but in the absence of selective control over presynaptic inhibitory interneurons it has not been possible to address whether, or how, this specialized form of inhibition influences motor behavior. We have used Gad2 as a genetic entry point to manipulate presynaptic inhibitory interneurons, revealing that their optogenetic activation suppresses sensory transmitter release. Moreover genetic stripping of presynaptic inhibitory boutons from sensory terminals uncovers a stereotypic forelimb motor tremor that has its origins in feedback-driven damped harmonic oscillation, severely degrading the fidelity of goal-directed reaches. These findings add experimental weight to the idea that the fine-tuned regulation of sensory input gain is essential for the smooth operation of goal-directed motor behaviors. More generally, our findings point to the potential of merging mouse motor behavioral analysis and molecular genetics in the interrogation of circuits for skilled reach.

## **T-2. The functional organisation of synaptic connectivity in visual cortical microcircuits**

Thomas Mrsic-Flogel  
Universität Basel

The brain's ability to extract information from the sensory environment depends on the computational capabilities of precisely connected neuronal networks in the cerebral cortex. Despite variations between cortical areas and species, neurons receiving similar inputs or sharing similar response properties often aggregate into fine-scale functional networks. In the mature primary visual cortex (V1), for example, neurons are not randomly connected, but form specific local circuits, which may influence the neuronal selectivity for different features of visual scenes, such as their orientation or location. Using in vivo two-photon imaging in mouse V1 and patch clamp recordings in slices of the same tissue, we are investigating the precision by which neurons with different functional properties connect to each other, and how this relates to information processing at the network level.

## **T-3. Words, metrics, and a thesaurus for a neural population code**

Elad Schneidman  
Weizmann Institute

Information is carried in the brain by the joint spiking patterns of large groups of noisy, unreliable neurons. The nature of correlations and noise in these neural population codes set the limits on the nature of information processing and computation in the brain. I will show that in the vertebrate retina the 'vocabulary' of neural populations is strongly correlated, and that the structure of the population code is shaped by a sparse network of high-order interactions between neurons. I will then show that maximum entropy models of network noise correlations allow us to characterize encoding at the population level, and build a neural 'thesaurus', measuring the similarity between population activity patterns based on the content they carry. This thesaurus reveals that the code is organized in clusters of synonymous activity patterns that are similar in meaning, but may be different in structure.

## **T-4. Neural codes for 2-D and 3-D space in the hippocampal formation of bats**

Nachum Ulanovsky  
Weizmann Institute

The work in our lab focuses on understanding the neural basis of spatial memory and spatial cognition in freely-moving, freely behaving mammals, employing the echolocating bat as a novel animal model. I will describe our recent studies, including: (i) recordings of 3-D head-direction cells in the presubiculum of crawling bats, as well as recordings from hippocampal 3-D place cells in freely-flying bats, using a custom neural telemetry system, which revealed an elaborate 3-D spatial representation in the mammalian brain; and (ii) recordings of 'grid cells' in the bat's medial entorhinal cortex, in the absence of theta oscillations which strongly argues against the prevailing computational model of grid formation. I will also describe our recent studies of spatial memory and navigation of fruit bats in the wild, using micro-GPS devices, which revealed outstanding navigational abilities and provided the first evidence for a large-scale 'cognitive map' in a mammal.

## T-5. Nonlinear input integration in cortical pyramidal neurons during behavior

Jeffrey Magee<sup>1,2</sup>

<sup>1</sup>Janelia Farm Research Campus

<sup>2</sup>Howard Hughes Medical Institute

Most principal neurons in cortical microcircuits integrate multiple classes or modes of excitatory input from distinct brain regions. These different inputs can be categorized as separate sensory modes, motor and sensory, internal and external, feedforward and feedback or predicted and actual. In pyramidal cell based circuits these different input streams are spatially separated with one input mainly innervating the proximal regions of the neurons and the other the most distal regions. We are using various in vivo recording and circuit manipulation techniques to examine the properties and functions of this input interaction in different circuits of the neocortex and hippocampus during simple behaviors. Thus far we have found that the integration of these distinct input streams proceeds in a nonlinear fashion. In both neocortical layer 5 and hippocampal CA1 pyramidal neurons a dendritic calcium plateau potential is evoked when sufficient distal input (motor or entorhinal input) is coincident with back propagating action potentials evoked by the proximal input (sensory or hippocampal input). There is a complex theta phase dependence to plateau potential initiation in the CA1 cells. In the end, the additional current mediated by the dendritic plateau increases the firing rate of the pyramidal cells in what may be a form of gain modulation mediated by the multiplicative dendritic input interaction.

## T-6. Digging for genes that contribute to behavioral evolution

Hopi Hoekstra

Harvard University

Understanding how and which genes affect ecologically important behaviors remains a major challenge in biology. To address this goal, we are capitalizing on natural variation in burrowing behavior within and between species of deer mice (genus *Peromyscus*). We have found that several, but not all, deer mouse species construct stereotyped burrows, and that dramatic differences in burrow architecture evolve rapidly. Moreover, deer mice reared in captivity recapitulate their natural burrowing behaviors under controlled laboratory conditions, and cross-fostering experiments suggest that burrowing has a large genetic component. We have been focusing on two sister-species, the deer mouse (*Peromyscus maniculatus*) and the oldfield mouse (*P. polionotus*). While the deer mouse constructs a simple burrow with a single, short tunnel leading to a nest chamber, the oldfield mouse constructs complex burrows that are long and composed of multiple tunnels, including an “escape tunnel.” To understand how changes at the genetic level lead to these innate behavioral differences, we currently are using several complementary approaches. Here, I will talk about our efforts to: (1) characterize the differences in mouse behavior that lead to differences in burrow shape, (2) identify the genetic basis of burrow size and shape using a forward-genetic approach, and (3) understand the neurobiological pathways involved in this natural behavior. Together, we expect these studies will elucidate the links between evolution, genetics, neurobiology and behavior in wild mice.

## T-7. Generating and shaping novel action repertoires

Rui Costa

Champalimaud Centre for the Unknown

Some actions are innate or prewired (such as swallowing or breathing). Others are learned anew throughout life, likely through a process of trial and selection. We used electrophysiology, imaging and optogenetics in behaving

animals to investigate how novel self-paced actions are generated, and how actions that lead to particular outcomes are then selected. We uncovered that dopamine is critical for the initiation of novel self-paced actions, and for the modulation of variability in cortico-basal ganglia circuits. On the contrary, dopamine appears to be less critical for ongoing action performance, or for the execution of well-learned actions. As new actions are learned, trial-to-trial neural and behavior variability decrease. Plasticity in basal ganglia circuits is important for the decrease in action variability, and for the selection of particular behavioral and neural patterns during learning. With increasing practice, actions become organized into chunks, and neural correlates of motor chunking emerge in basal ganglia circuits. Basal ganglia sub-circuits differentially modulate the initiation and execution of these action chunks. These data offer interesting working hypotheses about the neural bases of action variability and selection during learning.

### **T-8. Some new thoughts about planning, learning and skill in the motor domain**

John Krakauer  
Johns Hopkins University

Data will be shown that suggest new ways to think about the role of the reaction time in planning, about the contributions of explicit processes to visuo-motor adaptation, about the role of reinforcement in long-term motor memory and about what we mean by motor skill.

### **T-9. Learning task representations for reinforcement learning**

Yael Niv  
Princeton University

In recent years ideas from the computational field of reinforcement learning (RL) have revolutionized the study of learning in the brain, famously providing new, precise theories about the effects of dopamine on learning in the basal ganglia. However, the first ingredient in any RL algorithm is a representation of the task as a sequence of states. Where do these state representations come from? In this talk I will first argue, and demonstrate using behavioral experiments, that animals and humans learn the structure of a task, thus forming a state space through experience. I will then present some results regarding the algorithms that they may use for such learning, and their neural implementation.

### **T-10. A multisensory decision task exposes mixed selectivity in posterior parietal cortex**

Anne Churchland  
Cold Spring Harbor Laboratory

Neurons within a single cortical structure likely support a dizzying array of behaviors. This is especially true for areas with diverse sensory and motor inputs such as the posterior parietal cortex (PPC). Two competing frameworks might explain how PPC neurons leverage diverse inputs to support multiple behaviors. First, dedicated neural populations within PPC might constitute subnetworks, each supporting a restricted set of behaviors. If so, functional clusters of neurons with selectivity for particular parameters should be evident. Alternatively, a neural population within PPC might support multiple behaviors with orthogonal patterns of activity across the population



supporting each one. If so, no clustering should be evident because single neurons would tend to have “mixed selectivity” for random combinations of task parameters. To distinguish these frameworks, we recorded PPC neurons in 4 rats trained to report decisions about the rate of a sequence of sensory events. In interleaved trials, sensory events were auditory, visual or multisensory. Many individual neurons were strongly selective for choice, stimulus modality or both. A novel analysis tested whether selectivity for task parameters clustered in dedicated subnetworks or was distributed randomly. We observed a marked absence of clustering of response patterns: individual neurons reflect random mixtures of choice and modality signals. Dedicated “choice” or “modality” neurons occur only as often as would be expected by chance given randomly distributed parameters. Importantly, linear decoders could nonetheless predict the animal’s choice or the stimulus modality successfully, indicating that mixed selectivity does not prevent reliable decoding. These observations argue that the same population of PPC neurons is used to support multiple behaviors, and that this may be achieved by using different linear combinations of the same neurons.

## T-11. Adaptive decision-making in a dynamic world

Joshua Gold

University of Pennsylvania

Many of our decisions are based on an accumulation of noisy information over time. In a dynamic world, this accumulation process must itself be adaptive, because changes can occur that render past information irrelevant to future decisions. For example, historical yields from a fruit tree that has since died no longer predict future yields. A history of stable stock prices can become irrelevant after a major change in corporate leadership. I will talk about ongoing work in my lab that has begun to identify the neural mechanisms responsible for making effective decisions in these kinds of dynamic environments. This work is based on four complementary approaches: 1) the development of an ideal-observer model to recognize and respond appropriately to environmental change-points under certain conditions, plus a systematically reduced version of the model to a simple analytic form that makes specific predictions about the underlying neural computations; 2) human psychophysics that allows us to quantify how well human performance matches predictions of our models; 3) measurements of pupil diameter that allow us to test specific hypotheses about the relationship between the computations described in our models and functions of the pupil-linked arousal system; and 4) measurements of neural activity in the brainstem noradrenergic nucleus locus coeruleus, which is associated with non-luminance-mediated changes in pupil diameter, and one of its primary cortical targets, the anterior cingulate cortex (ACC), of monkeys performing change-point tasks. The results suggest that arousal systems play sophisticated computational roles in how we make effective decisions in a dynamic world.

## T-12. The macaque face processing system

Doris Tsao

California Institute of Technology

How the brain distills a representation of meaningful objects from retinal input is one of the central challenges of systems neuroscience. Functional imaging experiments in the macaque reveal that one ecologically important class of objects, faces, is represented by a system of six discrete, strongly interconnected regions. Electrophysiological recordings show that these ‘face patches’ have unique functional profiles. By understanding the distinct visual representations maintained in these six face patches, the sequence of information flow between them, and the role each plays in face perception, we can gain new insights into hierarchical information processing in the brain.

## T-13. A unifying theory of receptive field formation

Carlos S N Brito  
Wulfram Gerstner  
Ecole Polytechnique Federale de Lausanne

CARLOS.STEIN@EPFL.CH  
WULFRAM.GERSTNER@EPFL.CH

Developmental models derived from the principle of sparse coding have been shown to reproduce receptive field (RF) formation as it is observed in different cortical areas. More recent models have been able to successfully capture RF formation while also satisfying important biological constraints, such as locality of the plasticity rule and signaling through spikes. Since a growing number of different models successfully explain RF formation, the question arises of what are the indispensable mechanisms for feature learning. Here, we present a theoretical approach that answers this question. We first reformulate the sparse coding algorithm to obtain a phenomenological model accounting for Hebbian plasticity and neurons with nonlinear activation functions. Our analytical result demonstrates that sparse coding can be reinterpreted as a biologically plausible model. We then conclude that it can be alternatively interpreted in the framework of contrast function maximization. It has previously been shown that half of all possible contrast functions are solutions for the recovery of sparse features. Since the RFs observed in the primary visual cortex (V1) are sparse features of natural images, we concluded that many possible plasticity rules exist that will be able to capture RF formation in V1. In particular, we demonstrate that the contrast functions associated with sparse coding, ICA, BCM theory and several spiking networks are possible solutions to the problem. Thus one can predict a priori that all these models will successfully capture RF formation. Moreover, one can predict the specific shape of the RFs developed by each model. Using our unifying theory, we finally identify the requirements that a model (or a neural system) should fulfill in order to implement feature learning.

## T-14. Learning to track moving stimuli with population codes

Joseph Makin<sup>1</sup>  
Benjamin Dichter<sup>1,2</sup>  
Philip Sabes<sup>1</sup>

MAKIN@PHY.UCSF.EDU  
BEN.DICHTER@BERKELEY.EDU  
SABES@PHY.UCSF.EDU

<sup>1</sup>University of California, San Francisco

<sup>2</sup>University of California, Berkeley

Tracking moving objects, including one's own body, is a fundamental ability of higher organisms, playing a central role in many perceptual and motor tasks. While it is unknown how the brain learns to follow and predict the dynamics of objects, it is known that this process of state estimation can be learned purely from the statistics of noisy observations. When the dynamics are simply linear with additive Gaussian noise, the optimal solution is the well-known Kalman filter (KF), the parameters of which can be learned via latent-variable density estimation (the EM algorithm). The brain does not, however, directly manipulate matrices and vectors, but instead appears to represent probability distributions with population codes; nor are stimulus dynamics always linear. We show that a recurrent neural network's modified form of an exponential family harmonium (EFH) that takes a probabilistic population code as input can learn to estimate the state of mildly nonlinear dynamics. After observing a series of population codes noisily encoding a moving object, the network learns to represent the velocity of the object and forms nearly optimal predictions about the position at the next time-step. This result builds on our previous work showing that a similar network can learn to perform multisensory integration and coordinate transformations for static stimuli. We also examine the receptive fields of the network's units through training and observe that position tuning precedes velocity tuning, which makes qualitative predictions about the developing and learning brain.

## T-15. Processing properties of medulla neurons define neural substrates of motion detection in *Drosophila*

Damon Clark<sup>1</sup>  
 Rudy Behnia<sup>2</sup>  
 Thomas Clandinin<sup>3</sup>  
 Claude Desplan<sup>2</sup>

DAMON.CLARK@YALE.EDU  
 RB141@NYU.EDU  
 TRC@STANFORD.EDU  
 CD38@NYU.EDU

<sup>1</sup>Yale University  
<sup>2</sup>New York University  
<sup>3</sup>Stanford University

The brain processes spatiotemporal changes in luminance to extract visual motion cues. The algorithms employed in this task and the neural circuits that implement them have been the focus of intense research for more than 50 years. An influential correlation-based model, the Hassenstein-Reichardt correlator (HRC) relies on the differential filtering of two spatial input channels, delaying one input signal with respect to the other. This delay allows selective amplification of the output only when the delayed and non-delayed signals coincide in time, signaling motion. Recent work in flies has identified two parallel motion channels specialized for detecting either moving light or dark edges. Each of these pathways requires two critical processing steps to be applied to incoming signals: differential delay between the spatial input channels in each pathway, and asymmetric treatment of light and dark contrast signals. While the neural substrates that define the input and output channels of the light and dark edge circuits have been identified, the neural substrates that implement these two critical processing steps remain elusive. Using *in vivo* patch-clamp recordings, we show that four medulla interneurons exhibit these processing properties. The interneurons Mi1 and Tm3 respond selectively to brightness increments, with the response of Mi1 delayed relative to Tm3. Conversely, Tm1 and Tm2 respond selectively to brightness decrements, with the response of Tm1 delayed compared to Tm2. HRC models that are constrained by these measurements produce outputs consistent with previously measured tuning properties of motion detectors in flies. We therefore propose that Mi1 and Tm3 perform critical processing of the delayed and non-delayed input arms of the correlator responsible for the detection of light edges, while Tm1 and Tm2 play analogous roles in the detection of moving dark edges.

## T-16. A theory of neural dimensionality and measurement

Peiran Gao<sup>1</sup>  
 Eric Trautmann<sup>1</sup>  
 Byron Yu<sup>2</sup>  
 Gopal Santhanam<sup>3</sup>  
 Stephen Ryu<sup>1</sup>  
 Krishna Shenoy<sup>1</sup>  
 Surya Ganguli<sup>1</sup>

PRGAO@STANFORD.EDU  
 ETRAUT@STANFORD.EDU  
 BYRONYU@CMU.EDU  
 GOPAL@NERUR.COM  
 SEOULMAN@STANFORD.EDU  
 SHENOY@STANFORD.EDU  
 SGANGULI@STANFORD.EDU

<sup>1</sup>Stanford University  
<sup>2</sup>Carnegie Mellon University  
<sup>3</sup>Google

In many experiments, neuroscientists tightly control behavior, record many trials, and obtain trial-averaged firing rates from hundreds of neurons in circuits containing millions of behaviorally relevant neurons. Dimensionality reduction has often shown that such datasets are strikingly simple; they can be described using a much smaller number of dimensions (principal components (PCs)) than the number of recorded neurons, and the resulting projections onto these components yield a remarkably insightful dynamical portrait of circuit computation. This ubiquitous simplicity raises several profound and timely conceptual questions. What is the origin of this simplicity and its implications for the complexity of brain dynamics? Would neuronal datasets become more complex if we recorded more neurons? How and when can we trust dynamical portraits obtained from only hundreds of neurons

in circuits containing millions of neurons? We present a theory that answers these questions, and test it using data from reaching monkeys. We derive a theoretical upper bound on the dimensionality of data. Our bound has a natural interpretation as a quantitative measure of task complexity. Interestingly, the dimensionality of motor cortical data is close to this bound, indicating neural activity is as complex as possible, given task constraints. Our theory provides a general analytic framework to ascertain whether neural dimensionality is constrained by task complexity or intrinsic brain dynamics, furthering our ability to interpret large-scale datasets. We also describe sufficient conditions on PCs underlying neural activity so that low dimensional dynamical portraits remain unchanged as we record more neurons, and show that they are satisfied by motor cortical data. This theory yields a picture of the neural measurement process as a random projection of neural dynamics, conceptual insights into how we can reliably recover dynamical portraits in such under-sampled measurement regimes, and quantitative guidelines for the design of future experiments.

## T-17. Mapping the brain at scale

Jeremy Freeman<sup>1,2</sup>  
Nikita Vladimirov<sup>1,2</sup>  
Takashi Kawashima<sup>1,2</sup>  
Yu Mu<sup>1,2</sup>  
Davis V Bennett<sup>1,2</sup>  
Josh Rosen<sup>3</sup>  
Chao-Tsung Yang  
Logan Looger<sup>1,2</sup>  
Misha Ahrens<sup>1,2</sup>

FREEMANJ11@JANELIA.HHMI.ORG  
VLADIMIROVN@JANELIA.HHMI.ORG  
KAWASHIMAT@JANELIA.HHMI.ORG

YANGC@JANELIA.HHMI.ORG

AHRENSM@JANELIA.HHMI.ORG

<sup>1</sup>Janelia Farm Research Campus

<sup>2</sup>Howard Hughes Medical Institute

<sup>3</sup>University of California, Berkeley

The behavior of an animal reflects computations across its entire nervous system, involving the coordinated activity of thousands of neurons within and across multiple brain areas. New imaging technologies [1] allow us to monitor neural function during naturalistic behavior at unprecedented scales. But the data are quickly outpacing the capabilities of ordinary analytical approaches: a half hour experiment can yield a terabyte or more. These data are complex and high-dimensional, and we want to understand how their structure evolves over both space and time. Recent developments in distributed computing on clusters are making rich data analytics at the terabyte or petabyte scale tractable. By adapting a novel, open-source platform for distributed computing, Spark [2], we have developed a library of analyses and interactive visualizations for revealing spatiotemporal structure in large-scale calcium imaging data. With Spark, we leverage both the computing power of a distributed environment and its distributed memory, performing in seconds complex, iterative computations that would be intractable or impossibly slow with existing methods. We apply these tools to the larval zebrafish. Calcium responses from nearly all neurons ( $\sim 10^9$  voxels reflecting the activity of  $\sim 10^5$  neurons) are recorded in a fish expressing GCaMP5 pan-neuronally using light-sheet microscopy. For the first time, whole-brain imaging is performed while the animal is behaving in response to visual stimuli. Our analyses yield computational brain maps at single-cell resolution. Anatomically structured populations of neurons are found to play distinct roles in stimulus encoding and the generation of motor behavior. The evolution of whole-brain activity through a low dimensional space can be monitored on a trial-by-trial basis and linked directly to behavior. The open-source analytical framework we offer thus holds promise for turning brain activity mapping efforts into interpretable results.

## T-18. Whole-brain imaging of a stereotypic nervous system reveals temporal coordination of global activity

Saul Kato  
Tina Schroedel  
Harris Kaplan  
Manuel Zimmer

Institute of Molecular Pathology

KATO@IMP.AC.AT  
TINA.SCHROEDEL@IMP.AC.AT  
HARRIS.KAPLAN@IMP.AC.AT  
MANUEL.ZIMMER@IMP.AC.AT

We have recently demonstrated a novel platform for whole-brain imaging in *C. elegans* (Schroedel et al., 2013), yielding the opportunity to simultaneously record the activity of nearly every neuron of a nervous system with stereotyped anatomy and a well-characterized connectome. We have also developed methods for automatically extracting and registering individual neural signals across animals to allow high-throughput generation and comparison of whole-brain activity datasets under varying stimuli and conditions and under pharmacological and genetic perturbations. Linear dimensionality reduction reveals attractor-like dynamics across several motor and interneuron classes. By complementing whole-brain imaging with single-neuron imaging in freely moving animals, we have discovered neural activity patterns that correspond to sequences of well-defined locomotory states, allowing the decoding of the instantaneous behavioral state of an animal from its neural activity in single trials. This revealed that behavioral state transitions correspond to traversals between attractors through compact regions of state space. A surprising finding is that the dominant components of these dynamics are represented across more than 30% of the neurons in the *C. elegans* brain. We are investigating the influence of chemosensory circuits on the dynamics of these premotor groups. We also sought to understand detailed network function by performing a comprehensive search for low-order non-linear, temporally causal transformations within activity datasets, thereby building a functional connectivity map. This approach reveals directed relationships not evident from correlation analysis, the majority of which form a subset of the anatomical connectome. We are further studying to what extent all of the connections can be excited through varied stimuli and environments. For the first time, we can observe the time evolution of an entire, explicitly mapped nervous system as it parses stimuli, integrates it with a dynamic internal state, and generates coordinated behavioral sequences, and we present first insights into how this real-time computation unfolds.

## T-19. Self-monitoring of network activity in multi-state recurrent neural networks

Sean Escola  
L.F. Abbott  
Columbia University

SEAN.ESCOLA@GMAIL.COM  
LFA2103@COLUMBIA.EDU

Neural networks can operate in multiple computational states corresponding to different tasks or subtasks. Transitioning between these states may be influenced by a network's own activity. For example, monkeys have reaction time increases if motor cortex is disturbed with microstimulation immediately before the go cue, suggesting that the transition between the delay and movement states has been postponed (Churchland and Shenoy, *J Neurophysiol*, 2007). This could reflect a self-monitoring system that recognizes the post-microstimulation activity as noisy, and thus delays transitioning until appropriate movement preparatory activity is recovered. To study transitioning, we trained a recurrent neural network to possess multiple non-fixed-point states, each associated with a particular trajectory produced by the network's linear readout unit. Inputs drive the network to transition between these states. We then included additional readout units ('transition opportunity detectors') whose roles are to monitor network activity and act as gates that permit transitioning only when the network activity is within some region of firing-rate space. Depending on the detectors' tuning, a wide range of network behaviors can be produced including no transitioning, deterministic cycles of states, and both Markovian and non-Markovian stochastic transitioning (in the presence of noise). Thus we show that networks can monitor their activity and influence their state transitioning as in the microstimulation experiments. Furthermore, we show that in large

networks, with high probability, random projections will provide a library of detectors from which those needed to effect a desired transitioning behavior can be chosen. These units tile firing rate space in a state-dependent manner, and, as state-dependent dynamics evolve, their population response resembles the condition-dependent time cells seen in hippocampus (Eichenbaum, in preparation) and posterior parietal cortex (Harvey et al., Nature, 2012), suggesting that these heretofore unexplained results may be features of random connectivity in the setting of neural dynamics.

### **T-20. The uncertain spike: a new spike-based code for probabilistic computation**

Cristina Savin<sup>1</sup>  
Sophie Deneve<sup>2</sup>

CSAVIN@IST.AC.AT  
SOPHIE.DENEVE@ENS.FR

<sup>1</sup>Institute of Science and Technology, Austria

<sup>2</sup>Ecole Normale Supérieure

Ambiguity and noise plague virtually all neurally-relevant computation. To overcome these limitations, uncertainty needs to be represented and dealt with appropriately. Indeed, behavioural evidence suggests the brain represents and computes using probability distributions. The neural underpinnings of these processes, however, remain unclear (Fiser, 2010; Pouget, 2013). In particular, while previous models have shown how spiking neural networks could represent a probability over single real-valued variables (Zemel, 1998), or the joint probability of many binary random variables (Buesing, 2011), it's less clear how neural circuits could encode high-dimensional real-valued distributions, such as those required in classic models of probabilistic computation in the brain (e.g sparse coding in V1). Here we propose a circuit model for representing, and computing with, such distributions. The core idea underlying our representation is that the network activity evolves through recurrent dynamics such that samples from the desired distribution can be linearly decoded from the neural responses. As such, our model inherits the computational benefits of a sampling-based representation (Fiser, 2010). Furthermore, our proposal overcomes some of the potential shortcomings of classic sampling-based codes, by making the representation robust to neuronal damage and enabling an instantaneous representation of uncertainty (because it allows several independent samples to be encoded simultaneously in a neural population). As a consequence of the spike-based encoding used (Boerlin, 2013), the model reproduces classic properties of neural responses, such as stimulus tuning, Poisson neural variability, low pairwise correlations. It also predicts nontrivial dependencies in sub-threshold activity as a function of the neurons' stimulus selectivity. Lastly, our model highlights some of the challenges associated with interpreting neural activity in relation to behavioural uncertainty.

### **T-21. Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons**

Srdjan Ostojic<sup>1,2</sup>

SRDJAN.OSTOJIC@ENS.FR

<sup>1</sup>Ecole Normale Supérieure

<sup>2</sup>Inserm U960

Balanced networks of excitatory and inhibitory neurons are believed to play the role of fundamental units of computation in the cortex. Asynchronous activity in such networks constitutes the primary medium for the propagation and the processing of information. Here we show that an unstructured, sparsely connected network of integrate-and-fire neurons can display two fundamentally different types of asynchronous activity. For weak synaptic couplings, the network at rest is in the well-studied, classical asynchronous state in which individual neurons fire irregularly at constant rates. For strong couplings, we find that the network at rest displays a novel type of heterogeneous asynchronous activity, in which the firing rates of individual neurons fluctuate strongly in time and across neurons. The two types of asynchronous resting states possess vastly different computational

properties. In the classical asynchronous state, temporally varying inputs lead to a highly redundant response of different neurons that favors information transmission but strongly limits the computational capacity of the network. In the heterogeneous asynchronous state, the incoming stimulus interacts with the internal dynamics, so that the response of different neurons to the input strongly vary. This variability in the population deteriorates the transmission of information, but provides a rich substrate for non-linear processing of the stimuli as performed in decision-making and categorization.

## T-22. Balanced cortical microcircuitry for spatial working memory based on corrective feedback control

Sukbin Lim<sup>1</sup>  
Mark Goldman<sup>2</sup>

SUKBIN@UCHICAGO.EDU  
MSGOLDMAN@UCDAVIS.EDU

<sup>1</sup>University of Chicago

<sup>2</sup>University of California, Davis

A hallmark of working memory is the ability to maintain graded representations of both the spatial location and amplitude of a memorized stimulus. Previous work has identified a neural correlate of spatial working memory in the persistent maintenance of spatially specific patterns of neural activity. How such activity is maintained by neocortical circuits remains unknown. Traditional models of working memory maintain analog representations of either the spatial location or the amplitude of a stimulus, but not both. Furthermore, although most previous models require local excitation and lateral inhibition to maintain spatially localized persistent activity stably, the substrate for widespread inhibitory feedback pathways is unclear. Here, we suggest an alternative model for spatial working memory that is capable of maintaining analog representations of both the spatial location and amplitude of a stimulus, and that does not rely on long-range feedback inhibition. Previously, we considered networks with no spatial structure in their connectivity and showed that networks can maintain graded levels of persistent activity through a corrective feedback signal: when recurrent excitatory and inhibitory inputs are balanced in strength, but offset in time, net recurrent inputs generate negative-derivative feedback that counteracts slippage in activity [1]. Here, we extend the negative-derivative feedback model to networks with a functionally columnar structure of recurrently connected excitatory and inhibitory neural populations. The resulting networks can temporally integrate inputs at any spatial location, are robust against commonly considered perturbations in network parameters, and, when implemented in a spiking model, generate irregular neural firing characteristic of that observed experimentally during persistent activity. This work suggests balanced excitatory-inhibitory memory circuits implementing corrective negative feedback as a substrate for spatial working memory and together with previous work on parametric memory and integration in spatially uniform networks [1], suggests negative-derivative feedback as a common principle underlying different types of working memory.

## T-23. Intrinsic neuronal properties switch the mode of information transmission in networks

Julijana Gjorgjieva<sup>1</sup>  
Rebecca Mease<sup>2</sup>  
William Moody<sup>3</sup>  
Adrienne L Fairhall<sup>3</sup>

GJORGJIEVA@FAS.HARVARD.EDU  
REBECCA.MEASE@LRZ.TUM.DE  
PROFBILL@U.WASHINGTON.EDU  
FAIRHALL@U.WASHINGTON.EDU

<sup>1</sup>Harvard University

<sup>2</sup>Technische Universitaet Munchen

<sup>3</sup>University of Washington

During a critical period of early development, waves of spontaneous electrical activity cross the cerebral cortex. During this same period, the intrinsic properties of cortical neurons undergo substantial changes. We use mod-

eling to investigate the potential role of experimentally observed changes in intrinsic neuronal properties in determining the ability of developing cortical networks to propagate waves of activity. We model feedforward networks with multiple layers of neurons to study information flow in the typical direction of wave propagation. We show that changes in intrinsic properties strongly affect the ability of such networks to represent and transmit information on multiple timescales. With properties modeled on those observed at early stages of development, neurons are relatively insensitive to rapid fluctuations and tend to fire synchronously in response to large-amplitude wave-like events. Following developmental changes in voltage-dependent conductances, these same neurons become efficient encoders of fast input fluctuations over few layers, but lose the ability to transmit slower, population-wide input variations across many layers. Depending on the neurons' intrinsic properties, noise plays different roles in modulating neuronal input-output curves, which can dramatically impact network transmission. We qualitatively predict network dynamics using a mean field approach that incorporates the neurons' noisy input-output curves and connectivity in the network. We propose that the observed changes in intrinsic properties can contribute to the transition from spontaneous wave propagation in developing cortex to sensitivity to local input fluctuations in more mature networks, priming cortical networks to become capable of processing functionally relevant stimuli. This work underscores the significance of simple changes in conductance parameters in governing how neurons represent and propagate information, and suggests a role for noise in switching the mode of information transmission from global to one in which computations are scaled to local activity.

## **T-24. A compartment-specific synaptic integration and plasticity model for hippocampal memory**

Patrick Kaifosh  
Jeffrey Zaremba  
Nathan Danielson  
Attila Losonczy  
Columbia University

PWK2108@COLUMBIA.EDU  
JDZ2110@COLUMBIA.EDU  
NBD2111@COLUMBIA.EDU  
AL2856@COLUMBIA.EDU

We combine experimental and theoretical approaches to study hippocampal circuit mechanisms underlying memory. Experimentally, we perform two-photon Ca<sup>2+</sup> imaging in area CA1 of awake head-fixed mice. By targeting Cre-dependent viral vectors to specific brain regions and imaging in single cytoarchitectonic layers, we are able to monitor the activity of anatomically and genetically identified hippocampal circuit components. We image long-range axons providing inputs to CA1, axons and somata of locally projecting interneurons within CA1, and the dendrites and somata of CA1 pyramidal cells providing the primary output of the hippocampus. Our experimental work is complemented by development of a theoretical model of the hippocampal circuit mechanisms for episodic memory. We investigate the computational implications of nonlinear interactions between distal dendrite-targeting direct entorhinal inputs and proximal dendrite-targeting tri-synaptic loop inputs to CA3 and CA1 pyramidal cells. Analyzing networks of binary neurons with interacting proximal and distal compartments, we find that these interactions can lead to plasticity and network dynamics enabling later recall of information-rich memory engrams, with the propensity of a memory to be recalled depending on its recency, salience, and similarity to the recall cues. Based on memory performance considerations, our model predicts that different distal and proximal input levels to CA3 and CA1 cause CA3 representations to be more sensitive to small changes in hippocampal input. Indeed, when simulating networks of rate neurons with spatially tuned inputs, we find remapping effects that are consistent with known differences in place field rate remapping between CA3 and CA1. According to our model, these differences result from differences in CA3 and CA1 circuitry and are not specific to remapping between similar environments. We are therefore applying our two-photon imaging methods to test for similar differences in the long-term stability of spatial representations between CA3 and CA1 of head-fixed mice on a treadmill.



## T-25. Stereotypy and the structure of behavioral space

Gordon Berman  
Daniel Choi  
Ugne Klibaite  
William Bialek  
Joshua Shaevitz  
Princeton University

GBERMAN@PRINCETON.EDU  
DMCHOI@PRINCETON.EDU  
KLIBAITE@PRINCETON.EDU  
WBIALEK@PRINCETON.EDU  
SHAEVITZ@PRINCETON.EDU

Most experiments in the neurobiology of behavior rely upon the concept of stereotypy, that animals frequently engage in movements that are performed often and with great similarity. While these actions are often the basis for mapping neural circuits and understanding the effects of genetic manipulations, stereotypy is usually defined in an ad hoc manner, thereby limiting the sensitivity and repeatability of subsequent analyses. Moreover, the underlying assumption that an animal's behavior can be described in terms of discrete states at all typically remains unverified. Here, we introduce a technique for mapping the behavioral repertoire of organisms, relying only upon the underlying structure of the animals' postural movements. Starting with only raw movies of a freely behaving organism, our approach generates a representation of behavior where precise definitions for stereotyped movements naturally emerge. Applying our method to movies of six closely-related species of freely-behaving drosophilids, we find that a wide variety of non-stereotyped and stereotyped behaviors, spanning a wide range of time scales. We observe the subtle behavioral differences between these species, identifying the some of the effects of phylogenetic history on behavior. Moreover, we find that the transitions between the observed behaviors display a hierarchical syntax, with similar behaviors likely to transition between each other, but with a long time scale of memory – information about the organism's current behavioral state is found to persist for up to thousands of transitions in the future. This suggests possible mechanisms of neural control. In recent work, we have also applied our techniques to other organisms, such as mice, ants, and humans, opening the possibility to explore a wide range of questions in animal behavior, from speciation, to aging, to the neural control of movements.

## T-26. Measuring and modeling the dynamics of the thermal memory of *C. elegans*

Ilya Nemenman<sup>1</sup>  
Konstantine Palanski<sup>2</sup>  
Frederic Bartumeus<sup>3</sup>  
William Ryu<sup>2</sup>

ILYA.NEMENMAN@EMORY.EDU  
PALANSKI@GMAIL.COM  
FBARTU@CEAB.CSIC.ES  
WRYU@PHYSICS.UTORONTO.CA

<sup>1</sup>Emory University

<sup>2</sup>University of Toronto

<sup>3</sup>Centre for Advanced Studies of Blanes

*C. elegans* has the capacity to learn associatively. For example, *C. elegans* associates temperature with food and performs thermotaxis towards this temperature when placed on a spatial thermal gradient. However, very little is understood about how *C. elegans* acquires this thermal memory. We have developed a novel droplet-based microfluidic assay to measure the dynamics of the thermal memory of *C. elegans*. Individual animals are placed in an array of microdroplets on a microscope slide, and a linear temperature gradient of 0.5 deg/cm is applied to the array. By measuring the position of the swimming *C. elegans* in the droplets, we show that they can perform thermotaxis. By calculating the temperature bias of the position over time, we quantify the worm's thermal preference (and therefore its thermal memory) and measure the dynamics of this memory when the animals are exposed to different conditions of feeding and starvation. Over a time scale of hours, we find that the thermal preference of wild-type worms decays and actually become inverted before disappearing. This biphasic result can be explained with a mathematical model of operant conditioning consisting of habituation to the experienced temperature (with a shorter time scale of an hour or two) and an avoidance memory reinforced by the absence of food, with a longer time scale of 5-10 hours. By assaying specific mutants with defects in the insulin

signaling pathway at the ligand (*ins-1*), receptor (*daf-2*), and kinase (*age-1*) level, we show that habituation and the avoidances memories are genetically separable and that certain mutations can increase negative association rates. We argue that such multiscale memory may result in a more efficient search for food in temporally varying environments.

## T-27. A quantitative model for the sensorimotor integration underlying chemotaxis in the *Drosophila* larva

Matthieu Louis<sup>1</sup>

Alex Gomez-Marin<sup>2</sup>

Aljoscha Schulze<sup>1</sup>

Vani Rajendran<sup>1</sup>

Gus Lott<sup>3,4</sup>

Marco Musy<sup>1</sup>

James Sharpe<sup>1</sup>

Madhu Venkadesan<sup>5</sup>

Parvez Ahammad<sup>3,4</sup>

Shaul Druckmann<sup>3,4</sup>

Vivek Jayaraman<sup>3,4</sup>

<sup>1</sup>Center for Genomic Regulation

<sup>2</sup>Champalimaud Neuroscience Programme

<sup>3</sup>Janelia Farm Research Campus

<sup>4</sup>Howard Hughes Medical Institute

<sup>5</sup>National Centre for Biological Sciences

MLOUIS@CRG.EU

AGOMEZMARIN@GMAIL.COM

ALJOSCHA.SCHULZE@CRG.EU

VANI.R88@GMAIL.COM

GUSLOTT@YARCOM.COM

MARCO.MUSY@CRG.EU

JAMES.SHARPE@CRG.EU

MV@NCBS.RES.IN

AHAMMADP@JANELIA.HHMI.ORG

DRUCKMANN@JANELIA.HHMI.ORG

VIVEK@JANELIA.HHMI.ORG

Behavioral strategies employed for chemotaxis have been studied across phyla, but the neural computations underlying odor-guided orientation responses remain poorly understood. Combining electrophysiology, optogenetics, quantitative behavioral experiments and computational modeling, we study how dynamical olfactory signals experienced during unconstrained behavior in odor gradients are processed by first-order sensory neurons of the *Drosophila* larva, and how this information is converted into orienting behavior. The activity of a single larval olfactory sensory neuron (OSN) is sufficient to direct larval chemotaxis (Louis et al., 2008). By characterizing the activity of a targeted OSN in response to naturalistic changes in stimulus intensity, we report the surprising degree of information processing taking place in first-order olfactory neurons. Using an ordinary-differential-equation formalism, we find that the nonlinear response characteristics of an OSN can be described by a general class of regulatory motifs that includes the incoherent feedback loop. This quantitative model reproduces two basic operations carried out by the OSN on the stimulus: temporal differentiation on short timescales and temporal integration on longer timescales, which in concert enable the OSN to detect features related to changes in the stimulus acceleration. The model explains how these operations permit the OSN to perform similar computations over a wide range of stimulus intensity. Using optogenetics in virtual reality experiments, we studied the sensorimotor integration underlying the navigation of odor gradients. We examined how increases in the activity of an OSN can suppress turning, whereas abrupt decreases in activity can initiate turning. We explain how the conversion of OSN activity into a choice between alternative behavioral states — running or turning — can be accounted for by a simple stochastic model. Altogether, our work clarifies the link between neural computations at the sensory periphery and decision-making processes that direct navigation in a sensory landscape.

## T-28. Modeling enhanced and impaired learning with enhanced plasticity

Subhaneil Lahiri  
 T.D. Barbara Nguyen-Vu  
 Grace Q. Zhao  
 Aparna Suvrathan  
 Han-Mi Lee  
 Surya Ganguli  
 Carla J. Shatz  
 Jennifer L. Raymond  
 Stanford University

SULAHIRI@STANFORD.EDU  
 BNGUYENV@STANFORD.EDU  
 GQZHAO@STANFORD.EDU  
 APARNASU@STANFORD.EDU  
 HANMILEE@STANFORD.EDU  
 SGANGULI@STANFORD.EDU  
 CSHATZ@STANFORD.EDU  
 JENR@STANFORD.EDU

The widely held belief that synaptic plasticity mediates learning suggests the tantalizing possibility that we might enhance learning by enhancing plasticity. However, many experiments involving animals with genetically enhanced LTP or LTD reveal conflicting results; some show that enhanced plasticity actually impairs learning, while others show it enhances learning. Theoretical principles allowing us to reconcile such conflicting data remain elusive. We developed a mouse in which parallel-fiber Purkinje cell synapse LTD is enhanced by genetically knocking out the Class-I major histocompatibility molecules H2-Db and H2-Kb. We then tested these knockout (KO) mice in a gain-up vestibular oculomotor reflex (VOR) learning task. Interestingly, we found such enhanced LTD mice showed impaired gain up learning. However, when such KO mice were first pre-trained using a gain down learning task, they subsequently showed enhanced gain up learning. However, this same pre-training in WT in contrast impaired subsequent gain up learning. These results indicate that whether learning is impaired or enhanced depends upon a complex interplay between plasticity rates and prior experience. Using computational analysis, we uncovered principles of synaptic plasticity that can account for this complex learning dynamics. Importantly, we found that widely used classical models of synaptic plasticity in which LTP/LTD simply increases/decreases a scalar synaptic efficacy, are unable to account for behavior. Instead, we required more complex, metaplastic synaptic models with multiple internal states, allowing the propensity for LTP/LTD to depend upon the prior history of LTP/LTD events. In particular, we found that a form of synaptic stubbornness in which repeated LTD biases the synapse towards internal states from which further LTD or LTP becomes more difficult, could naturally account for both impaired and enhanced learning phenotypes in WT and KO. More generally, our modeling provides a theoretical framework for resolving conflicting experimental results relating perturbations of synaptic plasticity to learning.

## T-29. Neural constraints on learning

Patrick Sadtler<sup>1</sup>  
 Kristin Quick<sup>1</sup>  
 Matt Golub<sup>2</sup>  
 Stephen Ryu<sup>3</sup>  
 Byron Yu<sup>2</sup>  
 Aaron Batista<sup>1</sup>

PSADTLER@PITT.EDU  
 KRISTINMQUICK@GMAIL.COM  
 GOLUB.MATTHEW@GMAIL.COM  
 SEOULMANMD@GMAIL.COM  
 BYRONYU@CMU.EDU  
 APB10@PITT.EDU

<sup>1</sup>University of Pittsburgh  
<sup>2</sup>Carnegie Mellon University  
<sup>3</sup>Stanford University

The neural basis of learning has been studied at multiple levels, from synapses to individual neurons and cortical maps. Yet, relatively little is known about how activity patterns shown by populations of neurons change in an organized manner during learning. To quantify how the activity of many neurons co-modulate, we define a multi-neuronal firing rate space, where each axis represents the firing rate of one neuron. A given direction (i.e., a vector) in this space represents a co-modulation pattern. This work investigated whether there are constraints on the co-modulation patterns a subject can generate. We studied this question using a brain-computer interface

(BCI), which allowed us to directly specify which co-modulation patterns we would like the subject to show. We implanted a monkey with a 96-channel electrode array in the primary motor cortex. The subject modulated his neural activity to control the velocity of a cursor on a center-out BCI task. Each session, the subject first used an intuitive BCI mapping (i.e., the relation between neural activity and cursor velocity) to control the cursor. We then perturbed the mapping guided by the intrinsic manifold (IM), which is a subspace of the multi-neuronal firing rate space defined by the frequently generated co-modulation patterns. Some perturbations required the monkey to combine the basis co-modulation patterns of the IM in new ways. Other perturbations required the monkey to generate new co-modulation patterns outside the IM. We hypothesized that it would be easier for the subject to learn to control within-manifold perturbations than outside-manifold perturbations. Across all sessions, the data supported this hypothesis. This suggests that the IM reflects constraints imposed by the underlying circuitry, which prevented the monkey from readily generating co-modulation patterns outside the IM. This work provides a novel network-level explanation for why learning some tasks might be easier than others.

### T-30. State-space models for cortical-muscle transformations

Jeffrey Seely<sup>1</sup>

Matthew Kaufman<sup>2</sup>

Christopher J Cueva<sup>1</sup>

Liam Paninski<sup>1</sup>

Krishna Shenoy<sup>3</sup>

Mark Churchland<sup>1</sup>

JSSEELY@GMAIL.COM

ANTIMATT@GMAIL.COM

CCUEVA@GMAIL.COM

LIAM@STAT.COLUMBIA.EDU

SHENOY@STANFORD.EDU

MC3502@COLUMBIA.EDU

<sup>1</sup>Columbia University

<sup>2</sup>Cold Spring Harbor Laboratory

<sup>3</sup>Stanford University

A common desire in systems neuroscience is to infer models that link recorded inputs and outputs. Input-output models can be static—mapping inputs at each time to outputs at that time—or dynamic, where outputs may depend, via internal dynamics, on the history of inputs. A paradigmatic example of input/output transformations is from motor cortex to muscle activity. Over the last 45 years, the nature of this transformation has been hotly debated. Does motor cortex code kinematic parameters, which are then transformed into muscle activity by the spinal cord and other subcortical structures? Or does motor cortex produce muscle activity directly? Both possibilities might involve static or dynamic transformations. Prior studies have sometimes fit muscle activity with a temporally filtered version of neural activity, consistent with the possibility that spinal dynamics may provide filtering. However, such studies have not applied standard system identification techniques to determine (1) whether dynamics provide additional explanatory power or (2) the nature of any underlying dynamics. We therefore compared the performance of a linear state-space model with that of a static model. Analysis was based on data collected from one monkey performing a delayed reach task (180 single unit recordings from motor cortex, 8 EMG recordings from muscles, 27 reach conditions). The state-space model was identified using subspace identification followed by an optimization routine. The dynamical model fit 85% of the output variance using 11 input dimensions, translating to 130 system parameters. The static model required 42 input dimensions to achieve 85% fit, translating to 336 model parameters. Thus, the state-space model can fit outputs using many fewer model parameters. Generalization performance was also superior for the state space models. Finally, examination of the underlying dynamics revealed that the improved performance was due not just to exponential filtering but to more complex dynamics as well.

## T-31. Spatial organization of synchronous cell assemblies in HVC

Jeffrey Markowitz<sup>1</sup>  
Yoonseob Lim<sup>1</sup>  
Grigori Guitchounts<sup>2</sup>  
Timothy Gardner

JMARKOW@CNS.BU.EDU  
YSLIM@BU.EDU  
GUITCHOUNTS@FAS.HARVARD.EDU  
TIMOTHYG@BU.EDU

<sup>1</sup>Boston University  
<sup>2</sup>Harvard University

'Time cells' appear throughout the vertebrate forebrain, across different species, brain areas, and behavioral contexts. These cells manifest the remarkable property of firing at precise time-lags relative to the onset of a behavior; yet, what exactly drives this property in any species or context remains unknown. To date, the most precise time cells found in vertebrates are pre-motor neurons in the songbird motor cortical area HVC (used as a proper name). These cells fire a burst of action potentials at just one or two locations in a song, serving to produce muscular sequences in the vocal organ, and also to guide exploratory trial and error learning. To understand the network mechanisms that generate time cells, we recorded both single units and local field potentials (LFPs) in HVC. We show that interneurons in HVC fire synchronously over 100-175 micron length scales, and that time cells fire in the gaps between bursts of interneuron activity. Using a novel, fully automated segmentation algorithm based on time-frequency contours, we also found LFP phase locking and enhanced cell firing probability near acoustic boundaries. These results, along with recent in vitro data, suggest a model in which the clock rate and temporal ordering of time cells are defined by release from local inhibition. This opens up new models for robust sequence generation based on networks with minimally structured connectivity.

## T-32. Arithmetic of dopamine prediction errors: subtraction with scaled inhibition

Neir Eshel  
Ju Tian  
Naoshige Uchida  
Harvard University

NEIR\_ESHEL@HMS.HARVARD.EDU  
JUTIAN@FAS.HARVARD.EDU  
UCHIDA@MCB.HARVARD.EDU

Dopamine (DA) neurons are thought to encode reward prediction error, or the difference between actual and expected reward. How do DA neurons calculate this response? Traditional models have assumed the following form:  $\text{Prediction error} = \text{actual reward} - \text{expected reward}$ . However, the computation at the heart of this equation's subtraction has never been verified experimentally. It is possible that expectation divides DA activity, a common neural computation that would transform downstream readout. To determine how DA neurons respond to expectation, we recorded from light-identified DA neurons in the ventral tegmental area while mice performed a classical conditioning task. The task included three odors, which predicted reward at 10%, 50%, or 90% probability. Surprisingly, we found that expectation had a divisive effect, reducing reward responses by 8%, 40%, and 75%, respectively. To determine if the division in the population holds for individual neurons, we modified the task. On some trials, an odor predicted reward with 100% probability, but the reward size varied randomly. On other trials, rewards of various sizes were delivered unexpectedly, without odor expectation. Using this task, we measured how the same expectation affected DA responses across a range of rewards. We found that expectation subtracted from DA responses: regardless of reward size, expectation removed an average of ~3 spikes/s. How do we reconcile individual neurons with the population? Although each neuron showed a subtractive effect, the size of this effect varied substantially. Across DA neurons, the magnitude of the subtractive effect was proportional to the neuron's mean reward response. Thus, a neuron's 'responsiveness' (i.e., response to free rewards) was a good predictor of the magnitude of subtraction due to expectation. This process of scaled inhibition is sufficient to convert individual subtraction to population division. Our results allow us to compactly describe DA neuron activity at the population and individual levels.

### T-33. Nucleus basalis cholinergic neurons broadcast precisely timed reinforcement signals

Balazs Hangya  
Sachin Ranade  
Adam Kepecs

BALAZS.CSHL@GMAIL.COM  
RANADES@CSHL.EDU  
KEPECS@CSHL.EDU

Cold Spring Harbor Laboratory

The nucleus basalis (NB) is a vitally important yet poorly understood forebrain neuromodulatory system that provides cholinergic projections to the entire neocortex. NB is thought to play significant roles in cognitive functions such as learning and attention, and its degeneration parallels the cognitive decline in patients with Alzheimer's disease and Parkinson's dementia. Despite its importance, the functional roles of NB are contentious. There are no electrophysiological recordings from identified cholinergic neurons during behavior, which would be necessary to disambiguate their function from intermingled GABAergic and glutamatergic projection neurons. We used a combination of cell-type-specific optogenetic and electrophysiological tools with high-throughput rodent psychophysics that allowed us, for the first time, to record from identified NB neurons during behavior. By expressing ChR2 in cholinergic neurons in the auditory NB of ChAT-Cre mice, we identified them in extracellular recordings based on their short latency response to photostimulation. We trained head-fixed mice on an auditory sustained attention task involving temporal uncertainty. Mice had to promptly detect faint 'go' tones presented at randomized times over a background noise and respond by licking a water port while withhold licking to 'no-go' tones of a different pitch. Remarkably, cholinergic neurons were phasically activated following delivery of the reinforcer, either air puff or water reward. Surprisingly, almost all cholinergic neurons recorded during the task (n=13/15 neurons in 4 mice) showed a similar fast response to the air puff with a median latency of 17.5 ms and a very low jitter (median jitter = 2.2 ms). Responses to reward were more variable. Such fast responses of cholinergic neurons have not been reported and cast into doubt the prevailing notion of neuromodulators as being slow-acting. These data also provide support to the idea of acetylcholine being a 'broadcast signal' involved in identifying salient events, possibly mediating cortical plasticity and learning.

### T-34. Neural activity predicts social interaction decisions during prisoner's dilemma games in primates

Keren Haroush  
Ziv Williams

KHAROUSH@GMAIL.COM  
ZWILLIAMS@PARTNERS.ORG

Harvard Medical School

Social interactions are highly dynamic, requiring animals not only to understand how their choices affect their own personal outcomes but also how their choices affect the outcomes of other animals in a group or how these animals may consequently respond. The neuronal basis for such decision making remains poorly understood. To investigate this processes at the neuronal level, multiple pairs of Rhesus monkey performed an iterative Prisoner's dilemma game in which each animal could either cooperate or defect in order to receive reward. Cooperation led to the highest mutual reward whereas defection led to the highest individual reward. If both monkeys defected, both received the least possible amount. Four monkeys performed this task, of which two underwent neuronal recordings from the dorsal anterior cingulate cortex (dACC). We find that numerous neurons signaled the planned selection of cooperation over defection in relation to both one's own current and prior selections and to mutual selections. Remarkably, many neurons were also highly predictive of the opponent monkey's likely future response, and were independent of their own motor response or receipt of reward. Neural population activity was predictive of self and other's choices on a trial-by-trial basis. Recording up to 100 neurons simultaneously, we trained a linear decoder to predict the monkeys selections based on population activity. We find that activity predicted up to 66% of the recorded monkey's own current choices and, remarkably, predicted the other monkeys yet unknown choices on up to 79.4% of trials. Consistent with these physiological observations, focal disruption of dACC activity led to marked degradation in cooperation between monkeys. These findings provide new insight into the

single-neuronal process by which animals make social economic decisions and anticipate the intentions of others as well as point to cingulate dysfunction as a potential locus of difficulties with such social interactions.

## T-35. Dissociated functional significance of choice-related activity across the primate dorsal stream

Jacob Yates  
Leor Katz  
Il Memming Park  
Jonathan W Pillow  
Alexander Huk  
University of Texas at Austin

JACOBY8S@GMAIL.COM  
LEOR.KATZ@GMAIL.COM  
MEMMING@AUSTIN.UTEXAS.EDU  
PILLOW@MAIL.UTEXAS.EDU  
HUK@UTEXAS.EDU

It is commonly assumed that weak choice-related fluctuations in the responses of MT neurons are indicative of their role in encoding sensory information used for direction discrimination, while larger choice-related correlations in LIP neurons signal their role in evidence accumulation for actually forming decisions. Here we show that the meaning of choice-related activity in area LIP is inconsistent with the notion of area LIP integrating signals from MT in forming decisions. We recorded from multiple well-isolated neurons in MT and LIP simultaneously, while macaques performed a reverse-correlation direction-discrimination task. On each trial, the subject viewed a patch of dynamic flickering gabors, some of which pulsed in one direction or the other at various times throughout a 1.16 msec viewing duration; the subject then communicated their choice with an eye movement. The motion patch was inside the RF of a cluster of MT neurons, and one of the saccade targets was placed within the RF of a cluster of LIP neurons. MT neurons responded in a direction-selective way that was weakly correlated with choices (choice probability  $\sim 0.54$ ); LIP neurons ramped in a choice-selective way with higher choice probabilities (CP  $\sim 0.65$ ). MT CP was large early and then faded; LIP CP increased over time as the psychophysical weighting of the stimulus lessened. We also inactivated either MT or LIP with muscimol while recording from the other area. Inactivating MT devastated psychophysical performance; inactivating LIP had no measurable effect (although it did bias free choices in a control task). Taken together, these results suggest that LIP receives a post-decisional 'cognitive efference' signal that is not critical for performing direction-discrimination, but which may be either epiphenomenal or relevant to some other computation not directly addressed within conventional perceptual decision making tasks.

## T-36. Task specific feedforward component of choice-related activity in MT

Alexandra Smolyanskaya<sup>1</sup>  
Stephen G Lomber<sup>2</sup>  
Richard Born<sup>3</sup>

ASMOLY@SAS.UPENN.EDU  
STEVE.LOMBER@UWO.CA  
RBORN@HMS.HARVARD.EDU

<sup>1</sup>University of Pennsylvania  
<sup>2</sup>University of Western Ontario  
<sup>3</sup>Harvard Medical School

The activity of individual cortical neurons can be predictive of an animal's decisions, a relationship quantified with choice probability (CP) or detect probability (DP). Rather than reflecting a neuron's causal contribution, these signals likely reflect the correlation structure between each neuron and others involved in the decision. In turn, these correlations likely result from shared input due to some combination of feedforward and feedback influences. The relative contributions of these sources remain unknown. Here, we directly measured the contribution of feedforward pathways to decision-related signals. We recorded neuronal activity in visual area MT while monkeys performed motion and depth detection tasks. During each session, we reversibly inactivated V2 and V3, which are a major source of input to MT and convey more information about depth than motion. This procedure spares

other major inputs to MT, such as those from V1, leaving MT robustly responsive during inactivation. During V2/V3 inactivation MT neurons became less predictive of decisions during the depth task (DP change = -40%,  $p < 0.01$ ) but not the motion task (DP change = +14%, not significant), demonstrating that this functionally specialized feedforward input makes a large task-specific contribution to choice-related signals. To better understand how specialized feedforward inputs can selectively contribute to DP during one task and not another we simulated the effects of removing such an input in the context of several standard decision models. We found that we could reproduce task-specific changes in DP only when we implemented a model that involves the comparison of activity between two pools of neurons, a model more commonly used for discrimination tasks. A one-pool model that is typically used for detection tasks could not account for our results. This finding suggests that the neural circuitry for detection tasks may be more similar to that for discrimination tasks than previously thought.

## **T-37. Long-range and local circuits for top-down modulation of visual cortical processing**

Siyu Zhang<sup>1</sup>  
Min Xu<sup>1,2</sup>  
Tsukasa Kamigaki<sup>1</sup>  
Sean Jenvey<sup>1</sup>  
Yang Dan<sup>1</sup>

ZHANG\_SIYU@BERKELEY.EDU  
XU\_MIN@BERKELEY.EDU  
TSUKASAR@BERKELEY.EDU  
SEANERS13@GMAIL.COM  
YDAN@BERKELEY.EDU

<sup>1</sup>University of California, Berkeley  
<sup>2</sup>Howard Hughes Medical Institute

Top-down modulation of sensory processing is a prominent cognitive process that allows the animal to select sensory information most relevant to the current task, but the underlying circuit mechanisms remain poorly understood. We took advantage of the genetic and viral tools available for mouse brain and dissected the underlying neuronal circuits involved in top-down modulation. Using both anterograde and retrograde tracing, we identified a region of mouse frontal cortex — the cingulate cortex (Cg) — innervating both primary visual cortex (V1) and superior colliculus (SC), which may be an anatomical analogue of frontal eye field (FEF) in primate. Optogenetic activation of Cg neurons strongly enhanced V1 neuron responses ( $n = 38$ ,  $P = 1 \times 10^{-4}$ ) and improved the mouse's performance of a visual discrimination task ( $n = 5$ ,  $P = 0.005$ ). Focal activation of Cg axons in V1 caused a response increase at the activation site but decrease at nearby locations ( $P < 10^{-5}$ ). From brain slice recording, we found that Cg axons directly innervate pyramidal neurons and local GABAergic interneurons (SOM+, PV+ and VIP+ neurons) in V1. To investigate the roles of the local interneurons in top-down modulation, we optogenetically inactivated each of them in vivo. Optogenetic inactivation of the PV+ interneurons caused similar increases at center and surround, while SOM+ neuron inactivation caused the largest increase at surround, converting the surround suppression into a slight facilitation. Inactivation of the VIP+ neurons, in contrast, blocked the center facilitation without affecting surround suppression, indicating a specific role of these interneurons in enhancing the visual responses at the site of Cg axon activation. These findings reveal how long-range cortico-cortical projections act through local microcircuits to exert spatially specific top-down modulation of sensory processing.

## **T-38. Cortical modulation of synaptic dynamics and sensory responses in auditory cortex during movement**

David Schneider  
Anders Nelson  
Richard Mooney  
Duke University

SCHNEIDER@NEURO.DUKE.EDU  
NELSON@NEURO.DUKE.EDU  
MOONEY@NEURO.DUKE.EDU

As animals engage with their environment, sensory and motor areas of the brain exchange information to facilitate



perception and to update motor planning. Movement-related signals that infiltrate the auditory system, particularly at the cortical level, are believed to be critical for auditory-motor tasks including speech and musicianship, and dysfunction of the circuitry conveying these signals is thought to underlie complex auditory hallucinations such as those that characterize psychosis. Despite the postulated roles of movement-related signals in normal and disordered hearing, the source of these signals, their effect on auditory cortical synaptic activity, and the circuit through which they exert their effects remain enigmatic. Here we provide evidence that during movement, the motor cortex provides a disinhibitory signal to the auditory cortex that quickly modulates synaptic dynamics of auditory cortical pyramidal neurons and suppresses responses to sounds. Performing *in vivo* intracellular physiology in freely moving mice and in head-fixed mice on a treadmill, we found that movements were preceded and accompanied by stereotyped changes in membrane potential dynamics, including decreased variability and slight depolarization. Sensory-evoked responses were weaker in the auditory cortex during movement relative to rest, and coincided with decreases in excitability and input resistance, suggesting that movement engages shunting inhibition within the auditory cortex. Indeed, experimentally activating inhibitory neurons within the auditory cortex of resting mice was sufficient to reproduce movement-like membrane potential dynamics. Moreover, experimentally suppressing motor cortical projection neurons — a subset of which synapse primarily onto inhibitory interneurons in the auditory cortex — caused auditory cortical neurons to revert to rest-like membrane potential dynamics during locomotion. These experiments support a cortical corollary discharge circuit that is activated during a variety of natural behaviors to suppress auditory cortical responses to sounds.

### **T-39. The role of local inhibitory interneurons in stimulus-specific adaptation in primary auditory cortex**

Ryan Natan  
Laetitia Mwilambwe-Tshilobo  
Maria Geffen  
University of Pennsylvania

RNATAN@MAIL.MED.UPENN.EDU  
LMWILAMBWE10@GMAIL.COM  
MGEFFEN@MAIL.MED.UPENN.EDU

Neurons in the auditory cortex exhibit stimulus-specific adaptation (SSA), a property in which responses to frequently occurring stimuli are suppressed, while those to rare stimuli are relatively enhanced. SSA is thought to underlie a range of auditory processing mechanisms, such as novelty detection and adaptation to background sounds. While functional aspects of SSA in the primary auditory cortex (A1) have been extensively characterized, little is known about the neuronal mechanisms that control its generation and propagation. Here, we tested how activity of interneurons within A1 affects SSA. We recorded extracellular spiking activity across all cortical layers in A1 of lightly anesthetized head-fixed mice, using linearly arranged multi-site electrodes. Mice were exposed to classical oddball stimuli designed to reveal SSA in neuronal spiking activity. Almost all recorded units exhibited SSA. Using an optogenetic technique, we selectively suppressed parvalbumin (PV+) or somatostatin (SST+) positive interneuron spiking activity during interleaved trials. Inactivation of either PV+ or SST+ interneurons led to reduced SSA. To investigate how interneurons influence the propagation of SSA within A1, we compared SSA magnitude in granular, supra- and infra-granular layers. We found that suppression of PV+ reduced SSA across all layers, while suppression of SST+ reduced SSA in only the granular and infra-granular layers. These findings shed light on recent studies, which found that SSA in the non-lemniscal pathway is dependent upon cortical activity. We show that SSA magnitude can be attenuated differentially between deep and superficial cortical layers by specific interneuronal subtypes. Together, these results suggest that local inhibitory interneurons may independently control the strength of adaptive signals sent from the layers of A1 to their respective cortical and subcortical targets.

## T-40. Thalamic reticular nucleus controls arousal through modulation of local thalamocortical dynamics

Laura Lewis<sup>1</sup>  
 Jakob Voigts<sup>1</sup>  
 Francisco Flores<sup>1</sup>  
 Matthew Wilson<sup>1</sup>  
 Michael Halassa<sup>1</sup>  
 Emery Brown<sup>2,1</sup>

LDLEWIS@MIT.EDU  
 JVOIGTS@MIT.EDU  
 FJFLORES@MIT.EDU  
 MWILSON@MIT.EDU  
 MHALASSA@GMAIL.COM  
 ENB@NEUROSTAT.MIT.EDU

<sup>1</sup>Massachusetts Institute of Technology

<sup>2</sup>Harvard University

During sleep, general anesthesia, and some types of coma, cortical neurons exhibit rhythmic activity in the delta (1-4 Hz) range. Delta rhythms are associated with periodic silencing of cortical neurons, and previous studies have suggested that they are a candidate mechanism for decreased arousal. However, the circuit mechanism that produces the delta rhythm and its causal link to decreased arousal are not known. Here we show that tonic optogenetic activation of thalamic reticular nucleus (TRN), the major source of inhibitory input to thalamus, rapidly produces delta waves in the associated cortex. Importantly, the induced delta waves are associated with a significant reduction in arousal state, decreasing the animals' motor activity and increasing non-REM sleep. The induced oscillations were limited to local cortical areas, demonstrating that delta generation is supported by corticotopic circuitry. Furthermore, the induced oscillations resemble those seen in sleep, as cortical units undergo periods of silence that are phase-locked to the delta wave. Activating TRN thus produced both physiological and behavioral signs of sleep. Single unit recordings in TRN and dorsal thalamus demonstrated that stimulation activated only a subset of TRN units but consistently suppressed thalamic firing. This result suggests a complex circuit mechanism in which mutual inhibition within TRN leads to increased firing in a subset of neurons, decreasing thalamic input to cortex and thereby producing cortical delta waves and decreased arousal. In summary, we identify a causal mechanism for induction of delta waves, show that their spatial spread is controlled by the extent of TRN activation, and demonstrate that the same circuit also decreases behavioral signs of arousal. These results show an essential role for TRN and thalamic inhibition in regulating sleep and arousal, and elucidate the thalamocortical circuitry that controls local cortical state.

## T-41. Bayesian active sensing in the categorization of visual patterns

Scott Cheng-Hsin Yang  
 Mate Lengyel  
 Daniel Wolpert  
 University of Cambridge

SCHY2@CAM.AC.UK  
 M.LENGYEL@ENG.CAM.AC.UK  
 WOLPERT@ENG.CAM.AC.UK

Categorizing visual patterns, such as whether a patch of fur belongs to a zebra or a cheetah, is a time-constrained process. Due to the fovea, the information extracted depends on the sequence of eye movements chosen to scan the scene. Passive precomputed scan paths are suboptimal as the optimal eye movement depends on past information gathered about the actual scene. The optimal Bayesian active sensor (BAS) uses knowledge of the statistics of different visual patterns and the evidence accumulated from previous saccades to compute the probability of each pattern category and, hence, where to look to achieve the maximal reduction in categorization error. Previous studies of eye movements have only examined active sensing in simple visual search in which randomly arranged objects are either targets or non-targets. In contrast, here we examine visual pattern categorization in structured scenes (ie. stripy vs. patchy images), in which information from multiple locations in the image must be integrated in order to solve the task. In our experiment, participants had to distinguish between two visual pattern categories (generated by different 2D Gaussian processes). We used a gaze-contingent display in which a small aperture was revealed at each fixation location. On each trial, we controlled the number of fixations (5-25) before the categorization decision. We compared participants' categorization performance under these active

eye movements to several passive revealing strategies spanning from random to optimal as computed by BAS to isolate categorization performance independent from sensory, selection and motor noise that may affect the selection of the eye movements. We show that our participants' eye movements are quantitatively close to the optimal strategy once motor and selection noise is included. Moreover, using only the performance data from our participants, BAS predicts the qualitatively different and individualistic scan paths for each visual category observed in our participants.

## T-42. Both spatial and temporal codes shape texture perception

Hannes Saal  
Justin Lieber  
Alison Weber  
Sliman Bensmaia  
University of Chicago

HSAAL@UCHICAGO.EDU  
JUSTINLIEBER@UCHICAGO.EDU  
AIWEBER@UW.EDU  
SLIMAN@UCHICAGO.EDU

A fundamental question in sensory neuroscience is how patterns of neuronal activity convey information about objects and culminate in perception. Here, we present evidence that the tactile perception of texture relies on a variety of complementary codes, both across neurons and across time, that in combination account for perceptual experience. During texture exploration the human fingertip undergoes complex mechanical transformations, which culminate in a spatial distribution of strains and stresses in the skin. At the same time, movement between skin and surface elicits low-amplitude, high-frequency surface waves that propagate across the finger skin. These different aspects of skin deformation — some spatial, other temporal — are transduced by different receptors embedded in the skin. We measure both the microstructure of a diverse set of 55 everyday textures and the vibrations elicited in the skin of human and non-human primates as these textures are scanned across their fingertips. In paired neurophysiological recordings, we record the responses of individual mechanoreceptive afferents to texture presentations. We find that the coarse spatial layout of textures is represented spatially across populations of one population of afferents, while texture-elicited skin vibrations are encoded in the responses of two others. These vibrations are systematically shaped by the texture microstructure, the fingerprint layout, and the scanning speed. Importantly, vibrations convey information about texture at spatial scales smaller than can be resolved spatially given the innervation density of the fingertip and therefore this vibrational mechanism greatly extends the range of tangible texture. Vibration-sensitive afferents respond with precise spiking patterns to textures and these spiking sequences complement information carried by spatial patterns of afferent activation. Finally, we show that both spatial and temporal codes need to be combined to account for perceptual judgments of texture. The perception of texture thus arises from the combination of diverse codes across all mechanoreceptor afferent classes.

## I-1. Coding consequences of activity propagation from sensory and artificial neural stimulation

Garrett B Stanley<sup>1,2</sup>  
Christopher Rozell<sup>1,2</sup>  
Daniel Millard<sup>1,2</sup>

GARRETT.STANLEY@BME.GATECH.EDU  
CROZELL@GATECH.EDU  
DMILLARD6@GATECH.EDU

<sup>1</sup>Georgia Institute of Technology

<sup>2</sup>Emory University

Artificial stimulation techniques afford the ability to replace lost function or interrogate neural circuits by acting as surrogate inputs to the brain. The design of artificial stimulation patterns must produce functionally relevant downstream neural responses. However, behavioral (Masse and Cook, 2010) and electrophysiological (Logothetis et al., 2010) evidence suggest that sensory and artificial stimuli activate neural circuits through distinct neural mechanisms with different neural response properties as the activation propagates. In novel experiments, we

directly measured the propagation of neural activity generated by sensory and artificial stimuli and establish a theoretical framework for the optimal design of artificial stimulation signal sets. Specifically, we quantified the amplitude, variability, and spatial spread of the cortical response to whisker, electrical, and optical stimuli using voltage sensitive dye imaging in the vibrissa region of the rodent primary somatosensory cortex. The variability in the cortical response to whisker deflections increased as a function of the response amplitude, whereas electrical and optical stimuli exhibited high variability only around the threshold response amplitude. Spatially, the electrical stimulation activated an increasingly large region of the cortical space, while whisker and optogenetic inputs produced spatially focused responses. The different response properties across sensory and artificial stimuli have profound implications for producing functionally relevant surrogate inputs into neural circuits. To overcome the distinct nonlinear mappings of sensory and artificial stimuli, we maximize the classification performance of an ideal observer of cortical activity through optimal signal set design based upon an information theoretic distance measure. Through maximization of the pairwise, nearest neighbor distances within the signal set, the average classification performance was identical for models of sensory and artificial stimuli in the face of the distinct nonlinear mappings. The combined experimental and computational framework forms the basis for the design of optimal and functionally relevant surrogate inputs into neural circuits.

## **I-2. Auditory cortex mediates perceptual gap detection.**

Aldis Weible  
Alexandra Moore  
Christine Liu  
Michael Wehr  
University of Oregon

ALDIS@UONEURO.UOREGON.EDU  
ALEXANDRAKAYEHARTMAN@GMAIL.COM  
CLIU10@UOREGON.EDU  
WEHR@UOREGON.EDU

Understanding speech in the presence of background noise often becomes increasingly difficult with age. These age-related speech processing deficits reflect impairments in temporal acuity. Gap detection is a model for temporal acuity in speech processing, in which a gap inserted in white noise acts as a cue that attenuates subsequent startle responses. Lesion studies have shown that auditory cortex is necessary for the detection of brief gaps, and auditory cortical neurons respond to the end of the gap with a characteristic burst of spikes called the gap response. However, it remains unknown whether or how gap responses play a causal role in gap detection. Here we tested this by optogenetically silencing the activity of somatostatin- or parvalbumin-expressing inhibitory interneurons, or CaMKII-expressing excitatory neurons, in auditory cortex of behaving mice during specific epochs of a gap detection protocol. Silencing interneurons during the post-gap interval increased gap responses and enhanced gap detection. Silencing excitatory cells during this interval decreased gap responses and attenuated gap detection. Silencing activity preceding the gap had the opposite behavioral effects, whereas prolonged silencing across both intervals had no effect on gap detection. In addition to confirming the essential role of cortex, we demonstrate here for the first time a causal relationship between post-gap spiking activity and gap detection. Furthermore, our results suggest that gap detection involves an ongoing comparison of pre- and post-gap spiking activity. Finally, we propose a simple, yet biologically plausible neural circuit model that successfully reproduces each of these neural and behavioral results.

### I-3. Simultaneous optogenetic manipulation and calcium imaging in freely moving *C. elegans*

Andrew Leifer<sup>1</sup>  
Frederick Shipley<sup>1</sup>  
Christopher Clark<sup>2</sup>  
Mark Alkema<sup>2</sup>

LEIFER@PRINCETON.EDU  
FSHIPLEY@PRINCETON.EDU  
CHRISTOPHER.CLARK@UMASSMED.EDU  
MARK.ALKEMA@UMASSMED.EDU

<sup>1</sup>Princeton University

<sup>2</sup>University of Massachusetts

A fundamental goal of systems neuroscience is to probe the dynamics of neural activity that drive behavior. Here we present an instrument to simultaneously manipulate and monitor neural activity in freely behaving *C. elegans* while observing behavior. We use the instrument to directly observe the relation between sensory stimuli, interneuron activity and locomotion in the mechanosensory circuit. We induce neural activity by illuminating the optogenetic protein Channelrhodopsin (ChR2) and we monitor calcium dynamics by illuminating the optical calcium indicator GCaMP3 and a reference mCherry and observing their fluorescence. A transgenic animal is made to express ChR2 in one set of neurons and GCaMP3 and mCherry in another set. By patterning our illumination to independently illuminate either or both sets of neurons, the system can persistently observe calcium dynamics from one neuron while transiently activating another set in an awake and behaving animal. The worm crawls freely on agar while custom computer vision software identifies the location of targeted neurons and generates illumination patterns in real-time. The software adjusts mirrors on a digital micromirror device (DMD) to project patterned illumination onto only targeted neurons. The system independently illuminates neurons expressing GCaMP3 without activating ChR2 in neighboring neurons so long as the two targets are sufficiently spatially separated. We demonstrate independent illumination of three sets of neurons, AVA, ALM and AVM that are >50 microns apart in a moving worm. To observe calcium transients, GCaMP3 and mCherry fluorescence signals are imaged simultaneously on a CMOS camera via an additional imaging path. To our knowledge, this is the first instrument of its kind. We measured calcium transients in interneuron AVA in response to a 2.7 second optical activation of touch neurons ALM, AVM or both. We observe that calcium levels of AVA are negatively correlated with reversal velocity, regardless of stimulus.

### I-4. Modeling optogenetic manipulation of neural circuits in macaque visual cortex

Michael Avery  
Jonathan Nassi  
John Reynolds

AVERYM@UCI.EDU  
NASSI@SALK.EDU  
REYNOLDS@SALK.EDU

Salk Institute for Biological Studies

In recent years, a host of optogenetic tools have been developed, enabling experimentalists to manipulate neuronal activity with high temporal precision and cell-type specificity. This positions us to causally perturb the cortical circuit in ways that, when linked with computational modeling techniques, make it possible to test specific hypotheses about cortical circuits. Using optogenetics in the awake macaque, recent research in our lab has shown that activation of pyramidal neurons results in both facilitation and suppression of visually-evoked neuronal responses. These effects can be additive, multiplicative, or reflect competitive interactions between visual and laser-evoked responses (Figure 1A). We find that many of these changes can be accounted for by a simple normalization model (Figure 1B), in which laser activation of excitation results in indirect activation of inhibition, modulating the gain of neurons. This model, however, is abstract, and does not adequately account for the dynamics of the neuronal response. We have therefore developed a Hodgkin-Huxley model in which we explicitly model light-activated opsin conductances (Figure 1C). The model makes testable predictions about changes in spiking response dynamics across different temporal patterns of laser stimulation. For example, the model predicts that phase-randomized, slow oscillating inputs leads to increases in Fano Factor and coefficient of variation (Figure 1D). These findings

provide insight into how temporally-patterned optogenetic stimulation leads to changes in neuronal dynamics. More importantly, this illustrates how computational neural modeling can be combined with optogenetics to gain insight into cortical information processing.

## I-5. Amplification of synchrony in the *Drosophila* antennal lobe

James Jeanne<sup>1</sup>  
Rachel Wilson<sup>2</sup>

JAMES\_JEANNE@HMS.HARVARD.EDU  
RACHEL\_WILSON@HMS.HARVARD.EDU

<sup>1</sup>Harvard University

<sup>2</sup>Harvard Medical School

Synchronous activation of populations of neurons occurs frequently during normal brain activity, including sensory coding. In contrast to the expanding knowledge of how sensory stimuli are encoded by synchronous population activity, little is known about how downstream circuits decode this synchrony. This is largely due to the difficulty in controlling precise spike timing in large neural populations while identifying and monitoring activity in common postsynaptic neurons. We overcame these challenges by exploiting the orderly anatomy of the *Drosophila* antennal lobe, the first stage of olfactory processing in the brain. The antennal lobe is organized into discrete processing channels known as glomeruli, wherein ~60 olfactory receptor neurons (ORNs) synapse onto several projection neurons (PNs), which send olfactory information to higher brain areas. We expressed Channelrhodopsin-2 in all ORNs that project into a single glomerulus and designed light pulses that elicited exactly one spike per ORN. Short, intense pulses drove highly synchronous ORN spikes, while longer, weaker pulses drove less synchronous ORN spikes. We then made targeted loose-patch recordings of pairs of postsynaptic PNs. Both stimuli drove PN spikes with considerably tighter synchrony than in ORNs but with similar spike counts. Feedforward network theory (Diesmann, et al., 1999) demonstrates that synchrony amplification results from high convergence ratios, strong synapses, and large integration windows, which conspire to increase the onset slope of integrated postsynaptic currents. Steeper slopes traverse the spike threshold range faster, reducing susceptibility to noise. Consistent with this theory, whole-cell PN recordings revealed that higher synchrony stimuli indeed drove integrated post-synaptic currents with steeper onset slopes. These observations suggest that one important mode of decoding synchrony is to amplify and propagate it to higher brain areas. We are currently investigating neurons postsynaptic to these PNs to determine how they respond to high- and low-synchrony PN spiking.

## I-6. Effects of temporal information on perceptual decision making

Guang-Yu Robert Yang<sup>1</sup>  
Hyojung Seo<sup>2</sup>  
Xiao-Jing Wang<sup>1</sup>  
Daeyeol Lee<sup>2</sup>

GYANG.NEURO@GMAIL.COM  
HYOJUNG.SEO@YALE.EDU  
XJWANG@NYU.EDU  
DAEYEOL.LEE@YALE.EDU

<sup>1</sup>New York University

<sup>2</sup>Yale University

The ability to time various behavioral responses appropriately is crucial across a broad range of contexts. For example, humans and animals can anticipate the timing of behaviorally relevant sensory stimuli, and adjust their responses accordingly. However, how temporal uncertainty of sensory evidence influences perceptual decision remains poorly understood. Here we show that temporal information can strongly affect decision making strategy by altering the dynamics of decision threshold and sensory gain. To investigate the effect of temporal information, monkeys performed a novel color discrimination task with variable stimulus onset time. A non-informative pre-stimulus precedes the actual stimulus, making it difficult to detect the onset of weak sensory signals. Temporal information of the stimulus is manipulated by changing the duration of the pre-stimulus. Behavioral analyses of monkeys' reaction time revealed a strong influence of temporal information on the speed and accuracy of

monkeys' performance. In computational modeling with variants of drift-diffusion model, changes in the dynamics of decision parameters could account for the systematic influence of stimulus timing. Within a trial, sensory gain and noise gradually increase over time. Across trials, anticipation of delayed stimulus elevates decision threshold and sensory gain, but does not affect noise. Our task provides an efficient way to alter the speed-accuracy-tradeoff of monkeys, and allows studies of temporal dynamics of decision parameters over time. Our findings also suggest a possible neural mechanism that adjusts the animal's decision making strategy according to temporal information.

## I-7. Neural computations in the auditory system of fruit flies — from sensory coding to behavior

Jan Clemens<sup>1</sup>  
Cyrille Girardin<sup>1</sup>  
Pip Coen<sup>1</sup>  
Barry Dickson<sup>2,3</sup>  
Mala Murthy<sup>1</sup>

CLEMENSJAN@GOOGLEMAIL.COM  
CYRILLE.GIRARDIN@UNI-KONSTANZ.DE  
PIPCOEN@GMAIL.COM  
BARRY.DICKSON@IMP.AC.AT  
MMURTHY@PRINCETON.EDU

<sup>1</sup>Princeton University

<sup>2</sup>Janelia Farm Research Campus

<sup>3</sup>Howard Hughes Medical Institute

In *Drosophila*, acoustic signals produced by males play an important role in the female's decision to mate. The songs of different species differ both in spectral and temporal characteristics, and these features are thought to be important for female decision-making. Flies perceive sounds using mechanosensory neurons in the antenna that project to a brain structure called the antennal mechanosensory and motor center (AMMC). There, second-order auditory neurons pool receptor inputs. To unveil the neural computations underlying auditory processing during courtship, we fit models to electrophysiological recordings of the receptor neuron population and of second-order neurons in the brain. We find that receptor neurons encode both spectral and temporal features of song with high fidelity, whereas second-order neurons primarily represent temporal patterns present in song. We find that quadratic models — normally used to describe higher-order neurons in the visual or auditory system of vertebrates — are required to explain the encoding of sound among the population of receptor neurons. In contrast, simple linear-nonlinear models with adaptation are sufficient to explain responses of brain neurons. We then use quantitative behavioral assays to show how the output of brain neurons could be read-out and combined to produce the behavioral selectivity for short and long time-scale features of fly song. A population analysis of feature selectivity in the brain suggests simple computations on the outputs of early brain neurons that can replicate the behavioral selectivity for short temporal features of song. Also, the computations performed by auditory neurons — low-pass filtering and adaptation — can explain aspects of the behavioral selectivity of fruit flies for the song structure on longer time scales. We thus provide a comprehensive description of the neural code for sound in the early auditory system of flies and propose how this neuronal representation produces behavioral selectivity for song at multiple time scales.

## I-8. Adaptation in an auditory-motor feedback loop contributes to generating repeated vocal sequences

Kristofer Bouchard<sup>1</sup>  
Jason Wittenbach<sup>2</sup>  
Michael Brainard  
Dezhe Jin

KRISTOFER.BOUCHARD@GMAIL.COM  
JUW217@PSU.EDU  
MSB@PHY.UCSF.EDU  
DJINPSU@GMAIL.COM

<sup>1</sup>University of California, San Francisco

<sup>2</sup>Pennsylvania State University

The majority of ethological behaviors are not produced in isolation, but occur in the context of statistical sequences. For example, the song of Bengalese finches (Bf) is composed of probabilistically sequenced syllables, which can exhibit Markovian and non-Markovian transitions. Despite the ubiquity of sequencing in natural behaviors, the neural mechanisms that generate complicated motor sequences are largely unknown. In many song-birds, including the Bf, a prominent form of non-Markovian sequencing is long strings of syllable repetitions. Here, we test the hypothesis that stimulus-specific adaptation of auditory feedback contributes to the generation of repeated syllables. A branching chain neural network model with adapting auditory feedback to repeating syllable-chains produces repeat distributions similar to those observed in a large database of Bengalese finch songs. From the network model we derive an analytic model for repeat probability and show that it reproduces the repeat distributions of both the network model and the Bf song data. The analytic model predicts that long repeat syllables tend to have shorter durations, a feature found in both the data and network model. Removal of auditory feedback by deafening Bengalese finches rapidly and dramatically reduced the peak repeat number, confirming a key feature of the proposed mechanism. Furthermore, electrophysiological recordings in nucleus HVC of Bengalese finches show that auditory responses exhibit stimulus-specific adaptation to repeated presentations of the same syllable, providing evidence for another key feature of the proposed mechanism. Finally, we found that repeated syllables are louder than other syllables and elicit stronger auditory responses, suggesting that a threshold auditory feedback magnitude is required to generate long strings of repeated syllables. Together, these results implicate an adapting, positive auditory-feedback loop in the generation of repeated syllable sequences, and suggests that animals may learn to directly use normal sensory-feedback signals to guide behavioral sequence generation.

## I-9. Information flow through a sensorimotor circuit: Spatial orientation in *C. elegans*

Eduardo Izquierdo  
Paul L. Williams  
Randall D. Beer  
Indiana University

EDIZQUIE@INDIANA.EDU  
PLW.MAIL@GMAIL.COM  
RDBeer@INDIANA.EDU

Understanding how information about an external stimulus is processed by a nervous system is of great interest to neuroscience. Unfortunately, due to theoretical and experimental limitations this has been considered only in isolated parts of a nervous system (e.g., accuracy of sensory encoding). In this work, we characterize in detail the information flow through a complete sensorimotor circuit: from stimulus, to sensory neurons, to interneurons, to motor neurons, to muscles. Specifically, we apply a set of information dynamics tools to an ensemble of neuroanatomically grounded models of salt klinotaxis in *Caenorhabditis elegans*. The tools are based on a novel multivariate generalization of mutual information which extends Shannon's traditional concepts in several important ways. The unknown parameters of the model were optimized to reproduce the worm's behavior. Our analysis of the best solution reveals four key aspects of its operation: (1) There is a left/right information asymmetry in the circuit at the level of the AIY interneurons. (2) The AIZ gap junction plays a crucial role in the information transfer responsible for the symmetry in information necessary at the level of the motor neurons. (3) The SMB motor neurons display an information gating mechanism that plays an important role in the generation of a state-dependent response. (4) The overall time-lag of the information in the circuit is about 3/4 the duration of a full locomotion cycle. Further analysis reveals the crucial functional role this plays for steering. Despite the large variations observed in the parameter space of klinotaxis circuits, the information flow architecture revealed by our analysis is consistent throughout the ensemble. This suggests the information flow analysis captures the operation of the circuit at a level of description that is productive for understanding the general principles underlying the neural basis of this behavior.



## I-10. Basal forebrain mechanisms of sleep-wake regulation

Min Xu<sup>1,2</sup>  
Siyu Zhang<sup>1</sup>  
Yang Dan<sup>1</sup>

XU\_MIN@BERKELEY.EDU  
ZHANG\_SIYU@BERKELEY.EDU  
YDAN@BERKELEY.EDU

<sup>1</sup>University of California, Berkeley

<sup>2</sup>Howard Hughes Medical Institute

Regulation of sleep and wakefulness is essential for the survival of animals, and defects in the regulation cause hypersomnia, insomnia, and/or other associated mental illnesses. Basal forebrain (BF) has long been known to play an important role in sleep-wake regulation. In addition to the well-studied cholinergic neurons (ChAT+), there are GABAergic (PV+, SST+) as well as glutamatergic neurons (VGLUT2+) in the BF, but the roles of these neurons in sleep-wake regulation are still poorly understood. A critical step in understanding the functional roles of these neurons is to perform cell type-specific recording and manipulation in freely moving animals. In this study, we labeled specific types of BF neurons with light-activated channels (ChR2) for optogenetic identification and manipulation. We found that the firing rates of many BF neurons are highly modulated by sleep-wake states — awake, rapid-eye-movement (REM) sleep, and non-REM sleep (also called slow-wave sleep, 'SWS'). While ChAT+, VGLUT2+, and PV+ GABAergic neurons are wake and/or REM active, all SWS active neurons are SST+. To further explore how the BF local circuitry contributes to sleep-wake regulation, we used double transgenic mice to map the pairwise synaptic connections between cell types in brain slice. Together, these data led to a circuit diagram for BF regulation of sleep-wake states.

## I-11. An olfactory cocktail party - Detection of single odorants within mixtures in rodents.

Dan Rokni  
Vivian Hemmelder  
Vikrant Kapoor  
Venkatesh N Murthy  
Harvard University

DROKNI@GMAIL.COM  
V.HEMMELEDER@GMAIL.COM  
VIC.KAPOOR@GMAIL.COM  
VNMURTHY@FAS.HARVARD.EDU

In odorant rich environments, animals must be able to detect specific odorants of interest against varying backgrounds, which may occur at much higher concentrations than the searched target. Although this kind of task is generally thought to be part of the daily routine of many species, the limits of such olfactory figure-background separation in animal models remain unknown. Theoretical studies and human psychophysics indicate that the olfactory system has a very limited capacity to encode multiple odorants simultaneously or even to detect a specific target odorant within a mixture. Other studies suggested similar difficulties in rats, yet these were based on either spontaneous behavior with binary mixtures, or an assumption of generalization of task rules. These studies have contributed to the widely-held view that olfaction is primarily a synthetic sense in which odorant mixtures elicit emergent perceptions thereby making the perception of the individual components impossible. We designed a behavioral task to test the limits of figure ground separation in mice. We trained mice to detect the presence of a target odorant embedded in unpredictable and variable mixtures of up to 14 odorants. To test for a relationship between odorant similarity and separation difficulty, half the odorants in the pool were of the same chemical group. We found that mice can easily perform such a task. Performance was near perfect (95%) for stimuli composed of just one odorant and dropped to 85% in mixtures of 14 odorant. Background odorants of the same chemical group as the target had a strong negative effect on performance while other odorants had little effect. Our study offers new information on the ability of mice to solve the olfactory cocktail party problem, and paves the way for studies of neural circuits mediating this behavior in an experimentally accessible animal.

## I-12. Optimality, not simplicity governs visual decision-making

Wei Ji Ma<sup>1</sup>  
Shan Shen<sup>2</sup>

WEIJIMA@NYU.EDU  
SSHAN@CNS.BCM.EDU

<sup>1</sup>New York University

<sup>2</sup>Baylor College of Medicine

Many forms of human decision-making seem close to the ideal set by Bayesian optimality. At the same time, the brain might have incentives to use computationally simple and reasonably effective, though suboptimal decision rules; this has been called 'satisficing' (Simon) or 'ecological rationality' (Gigerenzer). A well-known example is the gaze heuristic people use to catch balls, but simple decision rules have also been extensively explored in visual perception. This raises the question of whether optimality or simplicity would prevail as an organizing principle of decision-making when the two are directly pitted against each other. The problem is that in previous studies, optimality was either ambiguous (too many choices for the modeler), or not tested, or indistinguishable from simplicity in the ability to fit data. We designed a novel visual search paradigm in which a) optimality is unambiguous, b) the optimal computation involves a triple marginalization and is by no standard simple, and c) simple rules cannot achieve near-optimal performance. We find that none of the simple rules describes human behavior well, but the optimal rule does, in spite of its high complexity. Differences are large and qualitative, not subtle at all. Our results suggest that in visual decision-making, optimality takes precedence over simplicity, and raise the question whether there are any circumstances at all under which simplicity is provably an organizing principle.

## I-13. Changes in prefrontal circuit organization increase repetitive network activity in models of autism

Francisco Luongo  
Meryl Horn  
Vikaas Sohal

FLUONGO@GMAIL.COM  
MERYLHORN@GMAIL.COM

University of California, San Francisco

Pathological changes in neocortical microcircuits are believed to instigate neuropsychiatric disorders including autism, but measuring such circuit-level changes remains challenging. Deficits associated with disorders such as autism often extend across multiple cognitive, behavioral, sensory, and affective modalities, suggesting that these diseases may comprise common information processing phenotypes at the level of neuronal networks. Such network-level phenotypes might be missed by studies of individual cells and instead require studies of datasets that probe patterns of network activity. Furthermore, disorders such as autism often have many causes; thus a major challenge is identifying common mechanisms that are conserved across multiple disease models and represent putative neural substrates for key behavioral and cognitive deficits. Here, we introduce novel methods for interrogating cortical circuits and find possible microcircuit endophenotypes associated with autism. Specifically, we use single-photon GCaMP imaging to measure patterns of activity generated by isolated prefrontal circuits in two etiologically distinct mouse models of autism: FMR1 knockout (KO) mice and mice exposed to valproic acid (VPA) in utero. In both models, we find enhanced functional interactions between prefrontal neurons and an increase in repetitive patterns of network activity that could plausibly contribute to perseveration and stereotyped behavior. Notably, these changes occur despite the absence of increased network excitability. We develop a novel computational analysis, which confirms that the enhanced interactions between neurons are sufficient to drive the observed increase in repetitive network activity. These results demonstrate a powerful new approach for identifying circuit-level changes in neuropsychiatric disease and linking microscopic changes (e.g. enhanced functional interactions) with their network-level consequences (e.g. an increase in repetitive patterns of activity).

## I-14. Attention and uncertainty during reward contingency learning

R Becket Ebitz<sup>1</sup>  
Eddy Albarran<sup>1</sup>  
Alireza Soltani<sup>2</sup>  
Tirin Moore<sup>1</sup>

REBITZ@GMAIL.COM  
ALBARRAN@STANFORD.EDU  
ASOLTANI@STANFORD.EDU  
TIRIN@STANFORD.EDU

<sup>1</sup>Stanford University

<sup>2</sup>Dartmouth College

Little is known about how attention is deployed during decision-making, particularly in naturalistic decision contexts with imperfect value knowledge and variable reward contingencies. Experimental evidence, frequently using targets with fixed reward values, suggests that attention is invariably directed towards the highest-value targets. However, disproportionately prioritizing high-value targets may not be the most efficient use of limited attentional resources when target values can only be estimated and require constantly updating. Therefore, we examined how attention is deployed when reward learning is required for performance of a value-guided decision-making task. We developed a behavioral paradigm that dissociates subjective target value from attentional deployment, utilizing an implicit measure of stimulus processing to quantify attention. Rhesus monkeys chose between probabilistically rewarded targets whose reward contingencies walked over trials. In order to maximize reward, monkeys had to choose targets on the basis of the reward history of each target. Though monkeys successfully tracked reward history, stimulus processing was not invariably enhanced for high-value targets. Instead, across two metrics of value uncertainty, we observed increases in stimulus processing as uncertainty about reward outcome increased. First, as the number of trials since the last selection of a target increased, stimulus processing increased. Stimulus processing was also enhanced with reward entropy, increasing in magnitude as the information content of the upcoming outcome increased. These results challenge the hypothesis that attention is solely determined by value. Rather, in a naturalistic, unpredictable environment, with only imperfect value knowledge, attention may be informed by value uncertainty. This observation provides empirical support for the idea that uncertainty may be able to influence learning through early adjustments in stimulus processing.

## I-15. The optimal allocation of attentional resource

Nuwan de Silva  
Wei Ji Ma

NUWAN12@GMAIL.COM  
WEIJIMA@NYU.EDU

New York University

The brain faces the fundamental challenge of making good perceptual and cognitive decisions using limited resources. Many studies have shown that directing attention to one location increases performance there but reduces performance at other locations. However, it is not known how attentional resource should be allocated in order to maximize performance. Here, we study a broad class of tasks that involve cueing: different stimuli (locations) have different probabilities of being probed. We assume that the expected utility at each location is a concave function of the amount of resource allocated, and that the total amount of resource is bounded above. For this broad class of tasks, we develop an efficient algorithm that is guaranteed to always find the optimal resource allocation. For a subset of tasks in this class, we show that the optimal strategy involves allocating zero resource to some items when the total resource level is low, while for other tasks we show that under the optimal strategy, some resource is allocated to all items regardless of the total resource level. We then illustrate these properties of the optimal solution by using our algorithm on a cued delayed estimation task and a cued detection task. Our work opens the door to normative models of attentional allocation.

## I-16. Surprise-based learning: neuromodulation by surprise in multi-factor-learning-rules

Mohammadjavad Faraji  
Kerstin Preuschoff  
Wulfram Gerstner

MOHAMMADJAVAD.FARAJI@EPFL.CH  
KERSTIN.PREUSCHOFF@EPFL.CH  
WULFRAM.GERSTNER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Surprise is a central concept in sensory processing, learning and attention as well as in the study of the neural basis of behaviour. Even though many successful learning models have been proposed, few of them have explicitly included surprise as a driving factor for learning. More importantly, the role of surprise in synaptic learning rules in neural networks is largely undetermined. Here we study surprise as a potential factor in multi-factor learning rules for neural networks using several known learning problems. We first show that, in classic reinforcement learning, modulating the learning rate by surprise significantly improves performance in both a reversal task as well as in a dynamic decision making task. We then use surprise in a clustering problem to identify whether a new data point should be assigned to one of  $N$  existing clusters or whether it belongs to a new, as of yet unknown cluster. In these analyses, we further show the validity of several mathematical approaches for measuring surprise including information theoretical approaches as well as Bayesian approaches. After having demonstrated that surprise theoretically improves learning, we show that signatures of surprise can be found in (simulated) neural networks during pattern learning as well as in fMRI BOLD response during the dynamic decision making task. Finally, we argue that surprise can have different effects on learning such as modulating the learning rate and regulating the exploration-exploitation tradeoff both of which are biologically plausible within our framework. Importantly, such surprise-driven modulations can enhance performance at both the behavioural and neural network level and are thus suitable as a third factor in multi-factor learning rules.

## I-17. Neural mechanisms of object-based attention

Daniel Baldauf  
Robert Desimone

BALDAUF@MIT.EDU  
DESIMONE@MIT.EDU

Massachusetts Institute of Technology

Attention to a location in space is thought to be mediated by feedback to visual cortex from spatially-mapped control structures such as the frontal eye field, but how we attend to objects and features that cannot be separated by location is not understood. We addressed this question with a combination of MEG and fMRI. We presented two temporally and spatially overlapping streams of objects, faces versus houses, and used a frequency-tagging approach to separate responses to attended and unattended stimuli. Attention to faces versus houses enhanced the sensory signals in the fusiform face area (FFA) and parahippocampal place area (PPA), respectively, as well as in a particular portion of the prefrontal cortex, the inferior frontal junction (IFJ); however, the FFA and PPA frequency-tagged sensory signals were advanced in phase by 25ms compared to IFJ, i.e. the sensory transmission time. By contrast, the top-down information for attention to faces and houses was associated with gamma frequency synchronization between the IFJ and the FFA and PPA respectively, and the IFJ led these areas by an equivalent 25ms, in gamma phase, i.e. the feedback transmission time. With these delays, activity in one area would be optimally timed to impact processing in the connected area. DTI tractography confirmed that the IFJ was connected with both FFA and PPA. Thus, the IFJ seems to be a key source of signals for object-based attention, and it modulates the sensory processing of objects at least in part through phase-advanced gamma-frequency synchronization.

## I-18. Characterizing the impact of category uncertainty on human auditory categorization behavior

Adam Gifford  
Yale Cohen  
Alan Stocker

University of Pennsylvania

MADA@MAIL.MED.UPENN.EDU  
YCOHEN@MAIL.MED.UPENN.EDU  
ASTOCKER@SAS.UPENN.EDU

Categorization is an essential process in auditory perception. However, proper categorization is often challenging because categories often have overlapping boundaries. Here, we explore the degree to which human subjects are able to learn and use category distributions and their prior probabilities, and the computational strategy they employ to make use of this prior information when performing category decisions. In a set of experiments, we asked subjects to classify a tone burst into two categories according to the perceived tone frequency. Tone stimuli were sampled from two uniform category distributions that overlapped in tone frequency. In different blocks of trials, we systematically varied the prior probability of presenting a tone from one versus the other category. We found that subjects consistently learned these changes in prior probabilities and appropriately shifted their decision boundary toward the category with higher probability. To test the computational strategy that governed subjects' decision behavior, we first characterized subjects' individual uncertainty levels in the tone stimuli using a two-tone discrimination experiment. We then generated parameter-free predictions for individual subjects' categorization behavior for a Bayesian model with either an optimal decision strategy or a strategy that selects decisions according to the category posterior probability (i.e., probability matching). Indeed, we found that the model with the probability-matching strategy was a substantially better predictor of behavior than the model with the optimal strategy. Moreover, this finding held even when fitting the model to the data. Our work provides evidence that human observers apply a probability-matching strategy in categorization tasks with ambiguous (overlapping) category distributions. Furthermore, the model fits confirmed that human subjects were able to learn category prior probabilities and approximate forms of the category distributions.

## I-19. Visual categorization reflects second order generative statistics

Feryal Mehraban Pour Behbahani  
Aldo A. Faisal

Imperial College London

FERYAL.MEHRABANPOUR10@IMPERIAL.AC.UK  
A.FAISAL@IMPERIAL.AC.UK

Categorization is a fundamental and innate ability of our brain, however, its underlying implementation is not well understood. Computationally, the task of classification is the same for humans as for machines. In machine learning, classification algorithms fall broadly into Discriminative and Generative (typically Bayesian) algorithms. Previous experimental work on human categorization shows data that is consistent with both discriminative and generative classification. However, the experimental designs used were not able to capture a desired test that unambiguously accepts one method while simultaneously rejecting the other. Therefore, we designed a novel experiment in which subjects are trained to distinguish two classes A and B of visual objects, while exemplars of each class are drawn from Gaussian parameter distributions, with equal variance and different means. During the experiment, we test how the subject's representation of the categories change as a result of being exposed to outliers for only one of the categories, A, far from category B (i.e. increasing category A's variance). Generative classifiers are by necessity sensitive to novel information becoming available during training, which updates beliefs regarding the generating distribution of each class. In contrast discriminative classifiers are sensitive to novel information only if it affects the immediate discrimination of classes. Our results show that, initially, when both categories have equal variance, subjects' classification behavior is consistent with both discriminative and generative algorithms. However, the introduction of the outliers for category A, influences the subject's knowledge of the distribution associated with alternative categories such that objects closer to category B and farthest from category A's outliers will be classified as belonging to category A, only predicted by our simulations of generative classifiers. This confirms and extends our previous work and gives clear evidence that visual categorization is

only consistent with generative and not discriminative classification mechanisms.

## **I-20. Continuous psychometric-neurometric comparison in a perceptual decision making task**

Thiago S Gouvea<sup>1</sup>  
Tiago Monteiro<sup>1</sup>  
Asma Motiwala<sup>1</sup>  
Joe Paton<sup>2</sup>

THIAGO.GOUVEA@NEURO.FCHAMPALIMAUD.ORG  
TIAGO.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG  
ASMA.MOTIWALA@NEURO.FCHAMPALIMAUD.ORG  
JOE.PATON@NEURO.FCHAMPALIMAUD.ORG

<sup>1</sup>Champalimaud Neuroscience Programme

<sup>2</sup>Champalimaud Centre for the Unknown

The involvement of neurons in the formation of a percept has been classically inferred from their choice probability, i.e. the degree to which differences in their firing during repeated presentations of a stimulus are predictive of perceptual judgments. How the brain integrates sensory information to give rise to perceptual judgments is still a mystery. One reason for this may be that information about the unfolding decision is transferred to many regions of the brain and even expressed in behavior, thus decision correlates should be expected in many brain areas. In addition, subjects' choices are often influenced by other biases not originating in the sensory stimuli they are judging. In this way, prior steps in the decision process or preexisting biases can provide new input to decision processes, making it harder to isolate mechanisms by which the brain integrates sensory information to form decisions. Given these complications, many have argued for recording neural activity in multiple brain areas during perceptual decision-making. We propose that this approach can be fundamentally improved by extracting continuous, quantitative behavioral metrics analogous to neuronal choice probability. As a testbed, we present striatal neuron population and behavioral data collected from rats performing a two- alternative forced choice interval discrimination task. We classified both high speed video excerpts and neuronal population data collected from near boundary intervals relative to choice. By testing classifier performance at successive time points within a trial, we determined the degree to which behavior and neural activity were predictive of choice. We found that both measures of choice probability become significant in advance of stimulus offset. We are currently acquiring and analyzing video and neural data recorded simultaneously, thus allowing a continuous psychometric-neurometric comparison. This approach will allow the identification of decision related neural activity not explained by ongoing behavior.

## **I-21. Dissecting the contributions of area MSTd and VIP to heading perception**

Kaushik Lakshminarasimhan<sup>1</sup>  
Sheng Liu<sup>1</sup>  
Eliana M Klier<sup>1</sup>  
Yong Gu<sup>2</sup>  
Gregory C. DeAngelis<sup>3</sup>  
Xaq Pitkow<sup>1,4</sup>  
Dora E. Angelaki<sup>1</sup>

JLAKS@CNS.BCM.EDU  
SLIU@CNS.BCM.EDU  
EKLIER@CNS.BCM.EDU  
GUYONG@ION.AC.CN  
GDEANGELIS@CVS.ROCHESTER.EDU  
XAQ@CNS.BCM.EDU  
ANGELAKI@CABERNET.CNS.BCM.EDU

<sup>1</sup>Baylor College of Medicine

<sup>2</sup>Institute of Neuroscience

<sup>3</sup>University of Rochester

<sup>4</sup>Rice University

In macaque monkeys, both the medial superior temporal (MSTd) and ventral intraparietal (VIP) areas contain neurons that are tuned to heading direction. Previously, we observed that inactivating area MSTd using muscimol injections increases animals' discrimination thresholds [1]. Here, we report that inactivating area VIP fails to

produce significant deficits in heading discrimination. This is surprising because individual neurons have similar sensitivity in the two areas, and VIP neurons exhibit substantially greater correlations with animals' choices than do MSTd neurons. We attempted to resolve this paradox by analyzing neural data recorded from these two brain areas. Using a standard noise model based on empirical pairwise noise correlations, the performance of an optimal linear decoder is inconsistent with the inactivation results. This leads to two alternatives: either (1) decoding is suboptimal or (2) the noise model is incorrect. We evaluate (1) by inferring the decoding weights of individual neurons from their choice probabilities [2]. To reduce uncertainty in the inferred weights, we restrict the weight computation to the modes with the largest noise variances. Although this restriction makes the decoder suboptimal by construction, the resulting predictions are qualitatively consistent with results from the inactivation experiments: the noisiest directions of VIP neural response space contain less information than the corresponding directions of MSTd response space. We explore (2) using an alternative noise model that incorporates additional information-limiting correlations [3,4] and we demonstrate how a robust optimal decoder could account for our experimental observations: greater information-limiting correlations in VIP than in MSTd would imply that less information about heading is available in VIP population activity. The analysis of neural data presented here is guided by theory as well as causal manipulation experiments, and represents a novel approach to understanding how distinct neural populations contribute to perceptual decision making.

## I-22. Computational model-based tDCS selectively enhances guilt-aversion based prosocial behavior

Tsuyoshi Nihonsugi<sup>1</sup>  
Masahiko Haruno<sup>2</sup>

T.NIHONSUGI@GIFU.SHOTOKU.AC.JP  
MHARUNO@NICT.GO.JP

<sup>1</sup>Gifu Shotoku  
<sup>2</sup>NICT Cinet

Understanding neural substrates of prosocial behavior is fundamental in human neuroscience. Although inequity aversion and guilt aversion have been implicated as key driving forces of prosocial behavior: inequity (Fehr & Schmidt, 1999) represents the difference in outcomes between self and the other, while guilt is the difference between the expectation and actual outcome of the other (Charness & Dufwenberg, 2006), it still remains unknown whether these two are dissociable neural processes. Here, we conducted fMRI and trans-cranial direct current stimulation (tDCS) experiments of a trust game, where player A chooses either 'in' or 'out' option. When player A chooses 'in', he also reveals an expectation probability of how likely player B will choose 'roll'. With the knowledge of the probability, player B then decides whether to 'roll' or not. We set inequity and guilt orthogonalized in the task and the utility function for player B can be formalized at time  $t$  as below.  $Utility = B0 + B1 * reward + B2 * Guilt + B3 * Inequity$  49 subjects participated in the fMRI experiment with the role of player B. We conducted model-based event-related analysis of fMRI data with SPM8. More specifically, in association with the presentation of game condition, we included inequity and guilt as parametric modulators. We found correlation with guilt in the right dorsolateral prefrontal cortex (DLPFC  $p < 0.05$ ; FWE), and in the right amygdala and ventral striatum for guilt ( $p < 0.001$ ; uncorrected). 25 different subjects were involved in tDCS experiment, where anodal electrode was put on the right DLPFC (cathodal on right parietal cortex). We compared Bs in the tDCS and sham conditions and found that B2 in the tDCS condition was significantly larger ( $p < 0.05$ ; paired t-test), while B1 and B3 were comparable. These results demonstrate that inequity aversion and guilt aversion are dissociable neural processes and computational model-based tDCS can selectively enhance guilt-aversion based prosocial behavior.

## I-23. Stochastic behavior in rats and its gating by the medial prefrontal cortex

Gowan Tervo<sup>1,2</sup>  
 Mikhail Proskurin<sup>1,2</sup>  
 Maxim Manakov<sup>1,2</sup>  
 Mayank Kabra<sup>1,2</sup>  
 Alison Vollmer<sup>1,2</sup>  
 Kristin Branson<sup>1,2</sup>  
 Alla Karpova<sup>1,2</sup>

GOWANTERVO@GMAIL.COM  
 PROSKURINM@JANELIA.HHMI.ORG  
 MANAKOV@JANELIA.HHMI.ORG  
 MAYANK.KABRA@GMAIL.COM  
 VOLLMERA@JANELIA.HHMI.ORG  
 BRANSONK@JANELIA.HHMI.ORG  
 KARPOVAA@JANELIA.HHMI.ORG

<sup>1</sup>Janelia Farm Research Campus

<sup>2</sup>Howard Hughes Medical Institute

Behavioral variability is often viewed as a reflection of the imperfect nature of decision-making. Purposeful variability in behavior could nevertheless be advantageous when trying to find the appropriate course of action. Action selection through stochastic choice would make behavior maximally variable and completely unpredictable, but would be at odds with one of the brain's central operating principles - to use experience to optimize behavioral choice. In experimental settings where their choices are rewarded only if they escape prediction by electronic competitors, primates, in fact, appear to resort to increasingly complex history-based patterns and even attempt to model the competitor's prediction algorithm, rather than adopting a stochastic strategy. The brain might thus be unable to stop using past experience, implying that there isn't a mechanism for stochastic action selection. We used a newly designed virtual competitive setting for rats to test whether, instead, the failure to observe evidence for stochastic action choices resulted from the inability of previously used virtual competitors to render history- and model-based strategies ineffective. We characterized the rats' behavior in response to electronic competitors that are more powerful at uncovering patterns in a subject's behavior than the ones previously described, and that use prediction algorithms that are progressively harder for the animal to model. We provide three key findings that support the existence of an outcome history- and model-independent — or effectively stochastic — mode of action selection: 1) animals increase the degree of behavioral variability and unpredictability with increased sophistication of the electronic competitor, 2) the behavior against the strongest competitor is maximally variable and is much less sensitive to environmental feedback and 3) in this mode and this mode alone, the influence of medial prefrontal cortex — an element of the circuit that guides behavior based on an experience-derived model of the environmental rules — is suppressed.

## I-24. The role of rat medial prefrontal cortex and superior colliculus in a rapid rule-switching task

Chunyu A Duan<sup>1</sup>  
 Jeffrey Erlich<sup>1</sup>  
 Katherine Fitch<sup>1</sup>  
 Carlos Brody<sup>1,2</sup>

CDUAN@PRINCETON.EDU  
 JERLICH@PRINCETON.EDU  
 KFITCH@PRINCETON.EDU  
 BRODY@PRINCETON.EDU

<sup>1</sup>Princeton University

<sup>2</sup>Howard Hughes Medical Institute

Our ability to inhibit stimulus-driven responses and execute context-appropriate actions is a hallmark of cognitive control. It has been proposed that prefrontal cortex (PFC) exerts cognitive control by inhibiting neural activity that leads to automatic stimulus-driven responses (the 'inhibition model'). This process has been studied extensively in the primate antisaccade task, in which subjects either apply a 'Pro' sensorimotor rule to orient towards a visual target (the 'prepotent' rule), or apply an 'Anti' rule to orient away (the 'non-prepotent' rule). The superior colliculus (SC) has been considered a key structure underlying Pro responses, and according to the inhibition model, needs to be inhibited for correct Anti execution. However, no SC inactivations during the ProAnti task have yet been carried out to test this. To facilitate the use of causal manipulations, we developed a computerized training method for a rodent version of the antisaccade task. This approach allows parallelized training that produces many



trained animals with minimal human intervention. Consistent with the primate antisaccade behavior, we found that trained rats manifest several hallmarks of rapid switching between a prepotent and a non-prepotent rule. To test the inhibition model in the rat ProAnti task, we then performed bilateral inactivations of SC or medial PFC. We found that, contrary to our expectations, bilateral SC inactivations induced a significant behavioral impairment on Anti trials, but not on Pro trials. Bilateral PFC inactivations produced similar results. These data are inconsistent with the inhibition model, and instead suggest that SC and PFC are part of a joint circuit that implements the Anti behavior. To account for these data, we propose a model in which competitive inhibition between task rules occurs within both PFC and SC, with PFC providing an excitatory input to SC that biases responses towards the appropriate rule.

## I-25. Information-theoretic basis of making decisions in the absence of new stimuli

Adam Calhoun<sup>1,2</sup>  
Sreekanth Chalasani<sup>2</sup>  
Tatyana Sharpee<sup>2</sup>

ACALHOUN@SALK.EDU  
SREEKANTH@SALK.EDU  
SHARPEE@SALK.EDU

<sup>1</sup>University of California, San Diego

<sup>2</sup>Salk Institute for Biological Studies

In order to successfully navigate an environment, animals must optimally utilize past information while integrating new information to the greatest extent possible. While new information is typically considered in term of new salient sensory stimuli, in certain conditions the only new information that arrives is the absence of new stimuli. By integrating this information, an animal can decide whether to continue the current task or give up, switching to a new one. We analyze a class of models where behavior involves moving to maximize information about a stimulus with an optimal Bayesian updating rule. Evidence for the use of these strategies in visual and odorant tasks exist, but we examine it from a unique perspective: how to search in the absence of positive input. We find that the model exhibits an emergent switching behavior between local and global search and characterize the switching time. We show a link between switching time and evidence accumulation, and how the Bayesian model is reducible to a drift-diffusion model of decision-making. Finally, we show evidence that the nematode *C. elegans* utilizes this strategy and provide evidence for neural mechanisms.

## I-26. Humans categorize two visual features of a single object sequentially

Yul Kang  
Michael Shadlen  
Columbia University

HK2699@COLUMBIA.EDU  
SHADLEN@COLUMBIA.EDU

When making two decisions, subjects appear to process information sequentially, a phenomenon known as the psychological refractory period (Pashler 1999). This is surprising, because the organization of sensory systems seems capable of parallel processing. We wondered if parallel processing is possible if just one decision is made based on two simultaneous sources of evidence. Subjects viewed a dynamic random dot display that conveyed both a motion direction and a dominant color. The difficulty of the color and motion dimensions was controlled independently. Subjects indicated the color and direction of motion (blue-left, yellow-left, etc.) by choosing one of four targets (saccade or button-press). In the 'reaction time (RT)' experiment, subjects viewed the stimulus until they responded. In the 'short presentation' experiment, the stimulus lasted only 300ms, and subjects responded after the stimulus offset. In 1-feature decisions (e.g., color), ambiguity of the ignored feature (e.g., motion) did not affect speed or accuracy. However, in 2-feature decisions, rendering either feature more ambiguous increased RT additively without improving accuracy with respect to the other feature, consistent with serial processing. From the processing times for each feature inferred from the RT experiments, we predicted that if perceptual

processing is serial, a 300ms stimulus exposure should force the subject to process one feature at the expense of the other. The 'short presentation' experiment contradicted this prediction; accuracy on the two features was not inversely correlated. However, subjects delayed their voluntary responses when either feature was more difficult, suggesting that color and motion information was still combined sequentially to inform the decision. Together, these findings suggest that information bearing on two dimensions of a decision can be acquired in parallel, whereas incorporation of information into a combined decision requires serial access to these parallel streams.

## **I-27. Expected information signals in the posterior cingulate cortex**

David Barack  
Jean-Francois Gariepy  
Michael Platt  
Duke University

DBARACK@YAHOO.COM

Animals must often balance learning about the environment with harvesting rewards. Animals can learn about the environment without reinforcement (Tolman 1948), suggesting independent encoding of information and reward. In contrast to the neural circuit mechanisms of reward and reward cue encoding, the neural mechanisms that acquire information about the environment independently of reward remain mysterious. We focus on expected information, defined as the amount of information learned from an outcome multiplied by the probability of that outcome. While gathering information about the environment by harvesting rewards, animals may anticipate the information to be gleaned from future actions, preferentially selecting those actions that maximize expected information about the current environment. We hypothesize that foraging animals maximize expected information, and encode expected value and expected information separately. To test these ideas, we designed an experiment where fixed rewards were randomly assigned to locations in the environment. On every trial, the amount of information, operationalized as the reduction in the number of possible locations of rewards, about the environment decreased over the course of the trial. Expected information varied only with information gained, while expected value varied with location and size of remaining rewards, thus partially decorrelating expected value and expected information. Response times were correlated with choice number and expected information but not expected value. In our task, neurons in the posterior cingulate cortex (CGp), a posterior midline cortical structure previously implicated in performance monitoring and exploration, preferentially encoded the expected information of a choice (6/11 cells) compared to the expected value of a choice (1/11 cells) or choice number in a trial (2/11 cells). Our sample robustly encoded the interaction of choice number in a sequence and expected information (9/11 cells), as well as the interaction of expected information and expected value (7/11 cells) and the three-way interaction (7/11 cells).

## **I-28. Ventral pallidum neurons signal an information-induced bias in handling reward uncertainty**

Ethan Bromberg-Martin<sup>1,2</sup>  
Okihide Hikosaka<sup>1</sup>

NEUROETHAN@GMAIL.COM  
OH@LSR.NEI.NIH.GOV

<sup>1</sup>National Eye Institute  
<sup>2</sup>Columbia University

Conventional theories of reinforcement learning explain how we choose actions to gain rewards, but we can also choose actions to help us predict rewards. This behavior is known as information seeking (or 'early resolution of uncertainty') in economics and a form of 'observing behavior' in psychology. We recently found that many neurons involved in conventional reward seeking also send signals appropriate for information seeking, suggesting that they share a common neural basis [1,2]. However, little is known about the underlying neural computations. When we make a gamble for uncertain future rewards (such as when we submit a grant application, or an abstract

to Cosyne), how does the brain decide how much we would be willing to pay to get early information about the outcome? To address this question, we developed a task to precisely measure the subjective value a monkey places on getting early information about a gamble's outcome, and how this value depends on features of the gamble including its expected reward magnitude and reward uncertainty. Behaviorally, we found that the subjective value of information scaled in proportion to uncertainty while growing only weakly with expected magnitude. In effect, the availability of information induced a bias in handling reward uncertainty: the subject was more willing to accept uncertain gambles when the subject was promised access to information that could reduce the uncertainty in the future. Neurally, we found a correlate of this behavior in the ventral pallidum (VP). A subset of VP neurons had responses that tracked the degree of reward uncertainty, and these responses were enhanced when the subject was promised early information about a gamble's outcome. This suggests that the VP, a region well-known to motivate seeking of conventional rewards [3,4], also has appropriate signals to motivate information seeking and its enhancement by uncertainty.

## I-29. On the neural dynamics of perceptual decision-making under temporal uncertainty

Federico Carnevale<sup>1</sup>  
 Omri Barak<sup>2</sup>  
 Victor de Lafuente<sup>3</sup>  
 Ranulfo Romo<sup>4</sup>  
 Nestor Parga<sup>1</sup>

FEDERICO.CARNEVALE@UAM.ES  
 OMRI.BARAK@GMAIL.COM  
 LAFUENTE@UNAM.MX  
 RROMO@CORREO.IFC.UNAM.MX  
 NESTOR.PARGA@UAM.ES

<sup>1</sup>Universidad Autonoma de Madrid

<sup>2</sup>Technion

<sup>3</sup>Universidad Nacional Autonoma de Mexico

<sup>4</sup>El Colegio Nacional

Perceptual detection requires the extraction of information about the world from noisy sensory signals. Under uncertainty, the brain uses previous knowledge acquired during learning to transform the sensory inputs into the percepts on which decisions are based. How is this achieved when the uncertainty lies in the timing of the task? We study this issue using a computational model and data recorded in monkeys performing a perceptual detection task in which the stimulation time is variable. According to Signal Detection Theory, this variability leads to a dynamic threshold reflecting the animal's previous knowledge about the task's temporal structure. Here we aim to find neural signatures of this threshold and study possible dynamical mechanisms underlying perceptual detection under temporal uncertainty. Previous knowledge about the task should be reflected in the probability of erroneously reporting a stimulus in stimulus-absent trials (false alarms). We found that premotor cortex neurons' activity during single false alarm trials shows localized fluctuations that resemble their activity under stimulation. We hypothesize that these fluctuations are signatures of false alarms. We devised a procedure to extract the times of these events and estimated the probability of false alarms over time. The resulting probability is not uniform but increases in agreement with the probability of stimulation. What mechanisms support detection under temporal uncertainty? We used a trained recurrent neural network to answer this question. Importantly, the information provided during training was restricted to the behavioral outcome. The resulting network learns to perform the task. Furthermore, it is able to infer the task's temporal structure as shown by the probability of false alarms. Our results show that a model under very simple settings can reproduce our experimental findings. In addition, reverse engineering the model allows us to uncover possible dynamical principles underlying detection under temporal uncertainty.

## I-30. Hierarchical competitions subserving multi-attribute choice.

Laurence Hunt<sup>1</sup>  
Ray Dolan<sup>1</sup>  
Timothy Behrens<sup>2</sup>

LAURENCE.HUNT@UCL.AC.UK  
RAY.DOLAN@UCL.AC.UK  
BEHRENS@FMRIB.OX.AC.UK

<sup>1</sup>University College London

<sup>2</sup>University of Oxford

Standard models of multi-attribute choice propose an 'integrate-then-compare' solution: attributes are integrated to form a unitary option value, followed by comparison of option values. However, previous behavioural findings indicate that this formulation may be misconceived, and emphasise the importance of within-attribute comparison and alternative decision strategies in reward-guided decision making. Here we report evidence for the neural basis of within-attribute comparison in a multi-attribute, three-option choice task. We show that distractor effects, caused by varying the value of a third option, are found in our task to occur selectively within, but not between, attributes. Combined with evidence from subject reaction times, this clearly demonstrated that options were compared within-attribute. We found that both choice behaviour and reaction times could be captured by a leaky, competing accumulator model of decision-making, but one in which evidence accumulation and competition occurred hierarchically. Thus, on any given trial, competition not only occurred on integrated values, but also within-attribute, and over which attribute should be paid more attention. Having fit this model to behaviour, we tested its predictions in human brain activity, indexed using functional MRI. Discrete brain regions carried distinct signals reflecting competition at different levels of the hierarchy. Whereas intraparietal sulcus reflected a competition over which attribute to attend to, dorsal medial frontal cortex reflected a competition over which option to choose. The functional connectivity of intraparietal sulcus to putative regions implementing lower level (within-attribute) comparisons was altered depending upon which attribute was currently more salient in guiding the decision. Taken together, these results suggest that 'integrate-then-compare' formulations of decision making may need revision. We argue that competition via mutual inhibition is a canonical computation operating at multiple hierarchical levels during reward-guided decision making.

## I-31. A Bayesian perspective on flexibly responding to stochastic and non-stationary tasks: a role for str

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. 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, tonically active cholinergic neurons (TANs), may provide a mechanism by which to modulate learning as a function of striatal uncertainty. A previously published neural network model of the basal ganglia was extended with constitutively active TANs and compared to a Bayesian reinforcement learning model. By modulating medium spiny neuron (MSN) activity, TANs were found to improve probabilistic learning but with a tradeoff: long TAN pauses result in better asymptotic performance whereas short TAN pauses facilitate faster learning following reversal. Long TAN pauses also saw greater decreases in entropy across learning, promoting MSN representations able to ignore irrelevant negative feedback at the cost of decreased flexibility following reversal. In the Bayesian reinforcement learner, a multiplicative parameter that decays the effective number of accumulated trials was found to replicate the effects of TAN pause duration. Slower decay resulted in narrower posterior distributions of expected value, facilitating robustness to irrelevant negative feedback at a cost of responsiveness to reversal, mirroring the effects of long TAN pauses in the neural network. Scaling TAN pause duration in the neural network with MSN entropy and decay in the Bayesian learner with action selection entropy mitigated the tradeoff and facilitated superior performance both asymptotically and following reversal. These findings were found to hold across a range

of task-stochasticity. This suggests that controlling of availability of MSNs for reinforcement through changes in TAN behavior a function of MSN uncertainty can promote a flexible and stable learning regime consistent with a Bayesian learner.

### **I-32. Graded representation of evidence accumulation and decision outcome in lateral intraparietal cortex**

Robbe L.T. Goris

Roosbeh Kiani

New York University

ROBBE.GORIS@NYU.EDU

RK97@NYU.EDU

Perceptual decision making requires transforming sensory inputs into motor responses. This transformation is represented in the neural activity of sensory-motor areas in association cortex. For example, responses of lateral intraparietal (LIP) neurons reflect accumulation of evidence when the choice is communicated with an eye movement. Although accumulation of evidence is well established in LIP responses, the full range of computations for the transformation of evidence to choice is less clear. Here we shed light on the diversity of those computations by developing a novel doubly stochastic model. In our model spikes are generated by a point process whose rate at each moment in time has a truncated Gaussian distribution. We show that nearly half of LIP neurons are best explained as accumulators of noisy evidence over time: the variance of the underlying rate grows linearly with time and its autocorrelation drops with the square root of time. Response variability of the other half of neurons is best explained by trial-to-trial modulations of a canonical but stimulus-dependent response buildup. This group of neurons might represent a transformation of the incoming sensory evidence. Alternatively, it may represent the outcome of evidence accumulation across the population. Compatible with the latter hypothesis, these neurons have a delayed representation of choice compared to the accumulator neurons. Moreover, they have a weaker dependence on stimulus strength and a faster pre-saccadic buildup. Our results reveal that the readout mechanism of the decision-making process can be implemented within the same cortical population that represents evidence accumulation.

### **I-33. Adaptive value coding and the optimality of choice**

Kenway Louie

Ryan Webb

Paul Glimcher

New York University

KLOUIE@CNS.NYU.EDU

WEBBR@NYU.EDU

GLIMCHER@CNS.NYU.EDU

Organisms in dynamic environments face constantly changing conditions. Given finite constraints in neural capacity, efficient coding theories require that neural systems adapt to the local statistics of the environment, a well-known phenomenon in sensory circuits. Adaptation also occurs in decision-related brain areas, but the computational benefits and behavioral consequences of adaptive value coding are unknown. Here, we show that adaptive value coding is implemented in parietal cortex via divisive normalization, and characterize the relationship between normalized adaptation and optimal choice behavior. To examine the neurophysiological basis of adaptive value coding, we recorded from lateral intraparietal (LIP) neurons in monkeys performing a temporal context saccade task. We find that value coding in these neurons is significantly dependent on the temporal reward context, despite the fact that such contextual information is irrelevant in individual trials. In contrast to attention or motivation effects, activity is higher (or lower) in smaller (or larger) average reward conditions; notably, this form of adaptive value coding is well-described by a temporal divisive normalization model. Furthermore, we show that this normalization model: 1) explains previously-reported range adaptation effects in orbitofrontal cortex, and 2) predicts adaptive changes in choice behavior specific for likely rewards in a given context. These findings suggest that a canonical gain control algorithm, divisive normalization, may mediate adaptive value coding in decision

circuits and predict specific benefits of such a computation for adaptive choice behavior.

## **I-34. An ideal observer solution for a perceptual decision making task**

Paul Masset<sup>1,2</sup>  
Joshua Sanders<sup>2</sup>  
Justin Kinney<sup>2</sup>  
Adam Kepecs<sup>2</sup>

PMASSET@CSHL.EDU  
JOSHUA21@GMAIL.COM  
JKINNEY@CSHL.EDU  
KEPECS@CSHL.EDU

<sup>1</sup>Watson School of Biological Sciences

<sup>2</sup>Cold Spring Harbor Laboratory

Understanding how perceptual decisions are reached in a noisy environment has been an active area of research in neuroscience. Historically, a large part of the field has been based on a visual discrimination task. Despite the simple decision rule, the perceptual stimulus is too complex to derive an ideal observer solution without making strong assumptions. More recently, similar decision making experiments have been designed using simpler tasks. Here we used a two alternative forced choice auditory discrimination task. Evidence is presented as two auditory click trains generated from two independent Poisson distributions. The subjects have to determine which click train has the higher underlying rate. The two rates always add up to a constant value but vary on each trial. Once the choice is made, the subject report on a scale of 1 to 5 the confidence in the decision made. Here we show that in this task a closed form solution for an ideal observer can be obtained using bayesian inference. Obtaining a full posterior distribution for the probability of the underlying stimulus given the perceived evidence allows us to investigate how high order decision variables, beyond the choice made, such as decision confidence can be represented. We show that the behavior of human subjects in the task follows the qualitative patterns predicted by the model. We show that the accuracy is close to the one predicted by the model. We also show that the inferred probability of being correct correlates with the reported confidence in the task. The model also predicts the pattern observed of reported confidence in which confidence is lowest for high discriminability error trials. In addition to these behavioral predictions, the full posterior distribution will allow a finer analysis of neural data once it is available.

## **I-35. Probing for patterns: Blending reward and information in a non-stationary POMDP**

Jeffrey Cockburn  
Michael Frank  
Brown University

JEFFREY\_COCKBURN@BROWN.EDU  
MICHAEL\_FRANK@BROWN.EDU

Learning and action selection are significantly complicated by the fact that observed events (e.g your flight departing on time) are often determined in part by unobserved processes (e.g the weather conditions en route). The advantage of considering latent processes is evident when events are predictable if considered in concert with their latent influences, but appear random otherwise. Considering latent processes imports additional levels of uncertainty about the true state of the latent process. Recent studies have suggested that ‘exploratory’ behaviour can be conceptualized as targeting the reduction of various forms of uncertainty. However, little work has examined whether organisms actively select actions that would reduce uncertainty about the state of latent processes. We propose a model that defines action selection in terms of both expected value and mutual information shared between action outcomes and latent state value. Critically, the tradeoff between these two components is a function of the belief state’s variance (uncertainty). We investigated this hypothesis using a latent structure reinforcement learning task. Participants were asked to repeatedly pick amongst cards that could be drawn from one of two decks. Participants were blind to the deck in play on each trial, but knew how the payoffs varied across decks. Results reveal an exploitative strategy when inferred belief state uncertainty was low, and an in-

creased probability of foregoing possible reward in favour of cards with more informative outcomes when belief state uncertainty was high. This pattern was observed across a broad range of task parameters, and was even observed when the hidden state value had no bearing on the optimal policy (e.g. the optimal policy was identical across decks). These results suggest that belief state uncertainty predicts the prioritization of information relative to reward, and that the reduction of uncertainty may be valued in its own right.

## I-36. Probing the sources of suboptimality in human Bayesian inference

Luigi Acerbi<sup>1</sup>  
Sethu Vijayakumar<sup>1</sup>  
Daniel Wolpert<sup>2</sup>

L.ACERBI@SMS.ED.AC.UK  
SETHU.VIJAYAKUMAR@ED.AC.UK  
WOLPERT@ENG.CAM.AC.UK

<sup>1</sup>University of Edinburgh

<sup>2</sup>University of Cambridge

When humans are presented with a simple Gaussian distribution of stimuli in an experimental setting, they are often able to combine this prior with noisy sensory information in agreement with the 'optimal' solution of Bayesian Decision Theory (BDT). However, in the presence of more complex experimental distributions (e.g. skewed or bimodal) performance appears suboptimal even after extensive training. Such suboptimality could arise from an inaccurate internal representation of the complex prior and/or from limitations in performing probabilistic inference on a veridical internal representation. We tested between these possibilities by developing a novel estimation task in which subjects were provided with explicit probabilistic information on each trial, thereby removing the need to learn the prior. The task consisted of estimating the location of a hidden target given a noisy cue and a visual representation of the prior probability density over locations, which changed from trial to trial. Priors belonged to different classes of distributions such as Gaussian, unimodal and bimodal. Subjects' performance was in qualitative agreement with the predictions of BDT albeit generally suboptimal. However, the degree of suboptimality was largely independent of both the class of the prior and the level of noise in the cue, suggesting that learning or recalling the prior constitutes more of a challenge to decision making than manipulating the complex probabilistic information. We performed an extensive model comparison across a large set of suboptimal Bayesian observer models. Our analysis rejects many common models of variability in our task, such as probability matching and a sample-averaging strategy. Instead we found that subjects' suboptimality was driven both by a miscalibrated internal representation of the parameters of the likelihood, and by decision noise that can be interpreted as a noisy representation of the posterior.

## I-37. Neural computations underlying dynamic decisions with disparate value types

Vassilios Christopoulos<sup>1</sup>  
Paul Schrater<sup>2</sup>

VCHRISTO@CALTECH.EDU  
SCHRATER@UMN.EDU

<sup>1</sup>California Institute of Technology

<sup>2</sup>University of Minnesota

Decisions involve two fundamental problems, selecting goals and generating actions to pursue those goals. While most research has focused on decisions between static options, humans evolved to survive in hostile dynamic environments where goal availability and value can change with time and previous actions. A predator chasing multiple preys exemplifies how goals can dynamically change and compete during ongoing action. Recent studies provide evidence that the brain generates concurrent action plans when faced with competing goals, and uses information accumulated during the movement to bias the competition until a single goal is pursued [Gold J and Shadlen M (2007), *Annu Rev Neurosci.*]. However, how value information is integrated online during action remains an open question. This information is diverse, involving both the dynamic value of the goal and action

costs, creating a challenging problem in integrating information across these disparate types in real time. We introduce a framework that models the neural computations underlying dynamic decision-making with disparate value types. The model is comprised of a series of control schemes that are attached to individual goals. The key novelty is a relative desirability computation that integrates value information to a single variable that dynamically weighs individual action plans as a function of state and time. Relative desirability weights have the intuitive meaning of the probability of getting the highest pay-off with the least effort by following an action plan. We evaluate the framework in a series of decision tasks and show that it can predict many aspects of decision behavior and its neural underpinnings that have eluded a common explanation.

## **I-38. Visual decisions in the presence of stimulus and measurement correlations**

Manisha Bhardwaj<sup>1</sup>  
Sam Carroll<sup>2,1</sup>  
Wei Ji Ma<sup>3</sup>  
Kresimir Josic<sup>1</sup>

SWEETMANI23@GMAIL.COM  
SRCARROLL314@GMAIL.COM  
WEIJIMA@NYU.EDU  
JOSIC@MATH.UH.EDU

<sup>1</sup>University of Houston

<sup>2</sup>University of Utah

<sup>3</sup>New York University

Our decisions are guided by the information we receive through our senses. Much work has been devoted to modeling the processes by which the brain converts noisy sensory measurements of a set of stimuli into a judgement about a global world state, such as the presence or absence of a target. These models often focus on various decision rules that can be applied to the measurements. By contrast, the measurements themselves are usually modeled in a rather stereotypical fashion, namely as independent and normally distributed. However, experimental data suggests that independence can be questioned. We ask how the statistical structure of the stimulus, together with correlations in measurements impact decision making. We address this question using a simple, analytically tractable model of a target search task where the computations performed by an ideal observer can be interpreted intuitively. We find that measurement correlations can significantly increase performance in the presence of strong external structure, when the stimulus distributions are concentrated on a low-dimensional manifold. If measurement noise is correlated, it can preserve such low dimensional structure, and have little impact on discriminability. Since natural scenes are highly structured, these results may have implications for decision making in natural environments.

## **I-39. Good noise or bad noise? The role of correlated variability in a probabilistic inference framework**

Ralf Haefner<sup>1</sup>  
Jozsef Fiser<sup>2</sup>

RALF.HAEFNER@GMAIL.COM  
FISER@BRANDEIS.EDU

<sup>1</sup>Brandeis University

<sup>2</sup>Central European University

The responses of sensory neurons in cortex are variable, and this variability is often correlated [Cohen and Kohn, 2011]. While correlations were initially seen as primarily detrimental to the ability of neuronal populations to carry information about an external stimulus [Zohary et al., 1994], more recent studies have shown that they need not be information-limiting in populations of neurons with heterogeneous tuning curves [Shamir and Sompolinsky, 2006; Ecker et al., 2011]. In general, information about some variable of interest,  $s$ , is only limited by correlations — sometimes called 'bad correlations' (Pitkow et al., SfN 2013) — that are equivalent to correlations induced by external uncertainty in  $s$ . The greater the magnitude of these correlations, the less information about  $s$  can be



represented by the population. We show in the context of a 2AFC task that correlations of the 'bad noise' structure are induced by feedback connections in a sensory population of neurons involved in probabilistic inference. However, unlike in the traditional encoding/decoding framework, here their presence reflects the fact that the brain has learnt the task. In fact, perceptual learning will increase their magnitude such that stronger 'bad' correlations are simply a side-effect of better psychophysical performance. We further show that increasing 'bad correlations' entails a steeper relationship between choice probabilities (CP) and neurometric performance as has been observed empirically during perceptual learning [Law and Gold, 2008]. Finally, we derive the results of classic reverse correlation techniques as applied to a neural system performing probabilistic inference and relate them to the bottom-up and top-down information flow as predicted by the normative model. Interestingly, we find that despite the fact that CPs are primarily caused by the top-down influence of the decision on sensory responses, they can nonetheless be used to infer the influence of an individual neuron onto the subject's decision.

## I-40. The origin and structure of behavioral variability in perceptual decision-making

Jan Drugowitsch<sup>1</sup>  
Valentin Wyart<sup>2</sup>  
Etienne Koechlin<sup>3</sup>

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

<sup>1</sup>University of Geneva

<sup>2</sup>INSERM, ENS

<sup>3</sup>INSERM, ENS, UPMC Paris

The variability of human behavior in perceptual decision-making is frequently larger than can be explained by variability in the underlying sensory evidence. This additional variability has been hypothesized to have its origin at the sensory, inference, or selection stage of the decision-making process, but the evidence for either alternative is still disputed. To distinguish between these three alternatives, we had humans perform a multi-sample visual categorization task in which they had to decide which of several alternatives was the generative orientation of a sequence of 2 to 16 high-contrast Gabor patterns. A model-based analysis of their behavior strongly supports the inference stage as the main contributor to their choice variability. We ruled out task-independent sensory noise as a viable explanation, as this hypothesis yields absurdly high orientation discrimination thresholds, and did not explain well the subjects' trial-by-trial choices. The selection stage was also an unlikely contributor, since it predicts choice variability independent of the length of the presented sequence, while this variability grows with sequence length. Lastly, we observed an increase in per-category variability when increasing the number of possible choices, which is compatible with noise at the inference stage, but not at the sensory stage. A follow-up experiment, which aimed at identifying the noise structure responsible for the observed variability, revealed that about a fourth of the noise variance could be attributed to deterministic biases in the inference process. We identified these biases as: (i) a skewed mapping of observed orientations into the space of log-likelihood, and (ii) an unequal weighting of the sequentially presented pattern when forming the log-posteriors. Taking these biases into account resulted in largely unbiased residual noise, which is not solely the result of deterministic approximations to Bayesian inference but instead seems to stem from intrinsic noise in the representation of information.

## I-41. Human intracranial electrophysiology supports a heuristic model of perceptual confidence

Yoshiaki Ko<sup>1</sup>  
Megan A. K. Peters<sup>2</sup>  
Hakwan Lau<sup>3</sup>

YOSHIAKI.D.KO@GMAIL.COM  
PETERS.MEGAN@GMAIL.COM  
HAKWAN@GMAIL.COM

<sup>1</sup>Columbia University

<sup>2</sup>University of California, Los Angeles

<sup>3</sup>\_

Human subjects are capable of conscious introspection upon their own perceptual processes, an ability sometimes referred to as metacognition. In sensory psychophysics experiments, this ability is reflected in the fact that subjects' confidence ratings correlate meaningfully with the likelihood of accurate decisions. In traditional psychophysics models, subjects are assumed to rate confidence based on the same internal sensory evidence underlying perceptual decisions. This theoretical view is adopted in recent animal studies, and supported by the finding that there are single neurons whose firing rate reflects both confidence and perceptual decision. Here we present evidence in support of a different view based on direct electrophysiological recording from the human cortex (in surgical epileptics). According to this view, subjects' level of confidence is driven mainly by the 'response-congruent evidence' in support of the chosen alternative, and is largely insensitive to the level of 'response-incongruent evidence' in support of the unchosen alternative. Six patients performed a simple face/house discrimination task, in which stimuli were embedded in noise and titrated to threshold performance. They rated confidence after their perceptual decisions. We analyzed high gamma activity from electrodes around the fusiform gyrus, which distinguished between face and house stimuli. Using reverse correlation analysis, we investigated sources of evidence leading to discrimination and confidence decisions. Whereas the face/house discrimination decision was driven by both house- and face-selective activity, confidence was driven mainly by activity congruent to subjects' discrimination decision. In other words, it is as if when subjects decide a figure resembles a face rather than a house, their confidence is mainly driven by the face-like quality of the figure, but is relatively insensitive to the lack of house-like qualities. While this surprising result may seem suboptimal, we argue that it may reflect a reasonable heuristical computational strategy given constraints faced by observers in everyday situations.

## I-42. Decision-making networks using spiking neurons

Jeff Orchard<sup>1</sup>  
Rafal Bogacz<sup>2</sup>

JORCHARD@UWATERLOO.CA  
RAFAL.BOGACZ@NDCN.OX.AC.UK

<sup>1</sup>University of Waterloo

<sup>2</sup>Oxford University

The brain performs decision-making processes that deal with imperfect, noisy input. We implemented several computational decision-making models using only spiking leaky integrate-and-fire (LIF) neurons. The models included the race model (Vickers, 1970), the inhibitory race model (Usher and McClelland, 2001), the GPR model (Gurney et al., 2001; Stewart et al., 2010), and the log-probability MSPRT model (Bogacz and Gurney, 2007). All spiking implementations successfully performed the decision tasks, even though the log-probability MSPRT network required two nonlinear decodings involving logarithms and exponential functions. For input saliencies of 0.6, 0.5, and 0.1, with added Gaussian noise  $N(0, 1)$ , we monitored the time to threshold for each model over a variety of threshold values. Plotting the resulting decision times versus the decision accuracies, we confirmed that the MSPRT yields shorter decision times than the race models. Interestingly, the GPR model was faster than MSPRT for accuracies below 88%, but slower than MSPRT above that.

### I-43. Activity of the anterior and posterior cingulate cortex during adaptive learning.

Yin Li<sup>1</sup>  
 Matthew Nassar<sup>2</sup>  
 Joshua Gold<sup>1</sup>

YINLI1@UPENN.EDU  
 MATTHEW\_NASSAR@BROWN.EDU  
 JIGOLD@MAIL.MED.UPENN.EDU

<sup>1</sup>University of Pennsylvania

<sup>2</sup>Brown University

In a dynamic environment, a rational agent should adaptively adjust how much it learns from environmental feedback. In environments characterized by periods of stability punctuated by sudden changes, feedback should be weighed more heavily when it is more informative, such as right after a changepoint. The goal of this study is to identify neural correlates of this adaptive learning process in the anterior cingulate cortex (ACC) and the posterior cingulate cortex (PCC), two brain regions known to play roles in reward processing and task control. We recorded from the ACC of two monkeys and the PCC of one monkey while they performed a ten-alternative saccadic-choice task. This task involved static fluctuations (noise) as well as abrupt changes (changepoints) in the identity of the rewarded target. Monkeys exhibited key signatures of adaptive learning: they were more sensitive to feedback in environments with less noisy trial-by-trial feedback and when changepoints had recently occurred. Units in both ACC and PCC exhibited activity that distinguished between feedback indicating reward or error. This feedback encoding was sensitive to the noisiness of the feedback in a subset of units in both areas and sensitive to the time since the last changepoint in a subset of ACC units. In addition, ACC exhibited "mixed" selectivity, with roughly equal numbers of units preferring correct or error feedback, whereas PCC exhibited more unidirectional selectivity, with most feedback-sensitive units preferring error feedback. These results suggest that ACC and PCC may play related but distinct roles in signaling contexts appropriate for adaptive learning.

### I-44. Effects of D2R blockage in the dorsal striatum on perceptual inference and reinforcement learning

Eunjeong Lee<sup>1</sup>  
 Moonsang Seo<sup>1</sup>  
 Olga Dal Monte<sup>1</sup>  
 Bruno Averbeck<sup>2</sup>

SCALLY12@GMAIL.COM  
 MOONSANG\_SEO@GMAIL.COM  
 OLGA.DALMONTE2@NIH.GOV  
 AVERBECKBB@MAIL.NIH.GOV

<sup>1</sup>NIMH/NIH

<sup>2</sup>National Institutes of Health

Selecting an action to achieve the outcome which is best suited to current goals relies on integrating previously acquired knowledge of the environment and immediately available information. If the mapping between actions and outcomes in a context is unpredictable over time, learning does not occur, and decisions must be made on the basis of immediately available information. Alternatively, if action-outcome mappings can be learned by reinforcement, they may be more informative than immediately available information. Emerging research suggests that frontal-striatal circuits are involved in the trade-off between these processes. However, the role of dopamine in frontal-striatal circuits in this has not been examined directly. To investigate this, we trained monkeys on an oculomotor sequential decision making task with two conditions: a fixed condition where animals could use reinforcement to learn action-outcome mappings over trials, and a random condition where animals had to rely on perceptual inference of immediately available information. While the monkeys performed the task we injected locally a dopamine D2R antagonist or saline into the dorsal striatum (dSTR) bilaterally, and recorded LFPs simultaneously in the lateral prefrontal cortex (IPFC) and the dSTR. Behavioral results showed that the D2R antagonist affected accuracy in the fixed condition only. Moreover, under saline, decisions were faster in the fixed than random conditions but after injection this reaction time improvement was decreased. Furthermore, persistent lower band power frequency coupling (beta) was increased within the dStr in the fixed condition when we blocked D2R. Overall, our results suggest the indirect pathway through the basal ganglia, which contains

primarily D2Rs, modulates the transition between making decisions with perceptual inference or reinforcement learning. Furthermore, the effects of antagonist injections in the dStr on the interaction within the dStr in lower frequency power coupling suggests that interaction among dStr neurons facilitates reinforcement learning in the fixed condition.

## **I-45. Effector-specific decision-making in the parietal cortex**

Jan Kubanek  
Lawrence Snyder  
Washington University

KUBANEK@WUSTL.EDU  
LARRY@EYE-HAND.WUSTL.EDU

Parietal cortex has been suggested to be involved in a variety of higher-order cognitive processes. The neural correlates of one such process, decision-making, have been thoroughly studied in the parietal eye fields of the monkey, the lateral intraparietal area (LIP). In this literature, LIP activity is modulated by decision-related variables (DV) in tasks in which animals indicate their decisions using eye movements. However, it has been unclear whether the decision-related neuronal effects in LIP generalize to decisions using response modalities other than eye movements. We recorded from LIP and the parietal reach region (PRR) in a reward-based decision task that elicits Herrnstein's matching behavior. In this task, animals chose one of two targets using an instructed response modality—an eye movement or a reach. We found that the modulation of LIP activity by the decision variable (DV) inferred in this task is substantially stronger during choices made using eye movements compared to choices made using reaches. PRR shows the reverse, encoding the DV exclusively for reach choices. To further determine whether effector-specific decisions are actively made in specific parietal circuits, we reversibly inactivated LIP and PRR using muscimol in a target onset asynchrony task. We found that LIP lesions specifically bias choices made using saccades, not choices made using reaches. PRR showed a reverse trend. This double dissociation revealed using both recordings and inactivations suggests that decisions in parietal cortex are made in circuits specific to particular actions. These findings support the proposed action-based, intentional framework of decision-making in parietal cortex, and support the emerging theories of embodied cognition in which higher-order, decision-related processes can be closely tied to the processes planning particular movements of the body.

## **I-46. Low frequency modulation of the beta rhythm during movement selection in posterior parietal cortex**

David Hawellek  
Yan Wong  
Bijan Pesaran  
New York University

D.HAWELLEK@NYU.EDU  
YAN.WONG@NYU.EDU  
BIJAN@NYU.EDU

The posterior parietal cortex (PPC) plays a key role in the sensory-motor transforms for visually guided movements of the eyes and arms. A central question is how the PPC achieves the selection of appropriate movement plans from multiple competing movement plans. To address this question we made simultaneous recordings of single units and local field potentials (LFP) in the lateral and medial banks of the posterior intraparietal sulcus of the macaque while subjects engaged in a decision making task that temporally dissociated target choice from movement selection. Specifically, the animal had to choose between two spatial targets (choice period) before correctly selecting either a saccade or a reach to express its choice (selection period). We find that during the entire task interval spikes and high frequency field activity (45 -256 Hz) were tightly coupled to the beta rhythm (16-32 Hz). The phase of the beta oscillation predicted spikes and the amplitude of the high-frequency field activity. However, specifically during the selection period frequencies in the theta range (3-8 Hz) modulated the beta rhythm. Strikingly, not the strength of this coupling but the modulation phase dissociated the selection of a reach from a saccade. The beta amplitude was coupled to significantly earlier theta phases for reaches than saccades

(Circular mean reach 256', saccade 289',  $p = 0.0102$ , Watson-Williams test). Theta phases predicted beta frequency amplitudes both within and across cortical areas, suggesting a widespread occurrence of this dynamic pattern across the large-scale circuit. Our findings show that movement related activity within in the PPC evolves on several intertwined time-scales. Ongoing beta band dynamics may reflect processes dedicated to movement planning. The circuit mechanism that flexibly selects specific movements in PPC may involve the packaging of beta-synchronized activity into barrages according to a lower frequency rhythm's phase.

## I-47. Neural mechanism of active learning based on hypothesis testing

Rei Akaishi  
Nils Kolling  
Matthew Rushworth  
University of Oxford

REI.AKAISHI@PSY.OX.AC.UK  
NILS.KOLLING@GMAIL.COM  
MATTHEW.RUSHWORTH@PSY.OX.AC.UK

The ability to learn causal relationships between events is highly developed in primates especially humans. The problem of causal learning has fascinated philosophers and inspired researchers in artificial intelligence. Here we present the results of an investigation of a neural mechanism for learning simple causal structure in a behavioral paradigm in which two alternative events could each cause a single consequence (common effect). The human subjects actively inferred the casual structure between the task events using a mechanism akin to hypothesis testing in science. The hypothesis chosen at a given time can be confirmed or disconfirmed and an alternative hypothesis can be supported. Subjects learned more quickly from decision outcomes that were consistent with the current, main hypothesis. At the neural level, the process involved in confirmation of a current, main hypothesis occurred in hitherto relatively unstudied portions of the medial frontal cortex. It was partly distinguished from the process involved in a switch to an alternative hypothesis in the lateral orbitofrontal cortex. Furthermore, even before the decision outcome was witnessed, we were able to observe neural activity related to both the main and alternative hypotheses. Holding these hypotheses in mind may be critical for enabling the process of hypothesis testing, and credit assignment, when the outcome is finally witnessed.

## I-48. The human sense of confidence obeys properties of a statistical confidence variable

Joshua Sanders  
Balazs Hangya  
Adam Kepecs  
Cold Spring Harbor Laboratory

JOSHUA21@GMAIL.COM  
BALAZS.CSHL@GMAIL.COM  
KEPECS@CSHL.EDU

Decision confidence is an estimate of the probability that a decision maker is correct, given the evidence used to decide. It is a key mental variable, necessary for the brain to optimally exploit imperfect information about the state of its environment. Despite its overt similarity to the concept of confidence in statistical hypothesis testing, the perspective that confidence is computed to solve a statistical problem in the deciding human brain was previously suggested only sparsely as a conjecture, and its implications were not fully explored. Here, we derived testable predictions based on a Bayesian theory of confidence. We start with statistical theory of hypothesis testing and define confidence as a posterior probability that a chosen hypothesis is true given the perceptual evidence used to choose. Based on this general framework we derived and simulated several key testable predictions about the interrelationship between confidence, choice and evidence strength. Next we tested whether human self-reported confidence is characterized by these relationships. First, we appended post-decision confidence reports to a previously characterized decision making task based on acoustic stimuli of graded discriminability. Subjects listened to two streams of Poisson clicks on each trial, indicated which stream was clicking faster on average, and reported confidence in their choice on a 5-division scale. Next we developed a general knowledge task. From

a pair of countries, subjects chose the country with the larger population, and rated their decision confidence. For both tasks we found that the interrelationships between confidence, choice and evidence strength showed patterns intrinsic to statistical confidence. These results suggest that the brain's conscious sense of confidence approximates key properties of a statistical confidence variable, with far reaching implications for understanding economic decision making under risk.

## **I-49. Fronto-parietal correlated discharges vary with performance during cognitive set shifting**

Rodrigo Salazar  
Nicolas Dotson  
Charles Gray  
Montana State University

RSALAZAR@CNS.MONTANA.EDU  
NICHOLAS.DOTSON@LIVE.COM  
CMGRAY@CNS.MONTANA.EDU

Patterns of synchronous activity represent the content of short-term memories (Salazar et al., 2012 in Science). To provide evidence that neuronal interactions are functionally relevant, we investigated the relationship between correlated firing (within and between 8 cortical areas in the prefrontal and posterior parietal cortex) and behavioral performance. The task consisted of an oculomotor delayed match-to-sample task involving a non-instructed switch between an identity- and a location-matching rule. Based the proportion of cells (%) with differences in firing rates for the two rules, we parceled the cortical areas into two categories referred to as rule-selective and non-selective. We found that the incidence of spike count correlations within each network differed strongly. In the selective and in the non-selective network, 17% and 6% were positively correlated, respectively. Surprisingly, the highest incidence was found between parietal area PG and dorsal lateral prefrontal cortex in 25% of the pairs. Furthermore, the magnitude of correlation varied with behavioral performance in 15% and 20% of these pairs for the identity and the location rule, respectively. During the sample presentation and before the match onset, a regression analysis on the correlation coefficients between the magnitude of correlated firing and performance revealed a significant linear fit as a function of rules. The negativity of the slopes represented a positive correlation with performance during one rule and a negative correlation during the other rule. This effect was observed only in the selective network. These findings demonstrate that correlated discharges occur preferentially between cortical areas with rule-selective cells and that their magnitudes vary with behavioral performance. To our knowledge, this is the first report of functional correlated firing between distant cortical areas.

## **I-50. Data-driven algorithms for learning state dynamics in natural decision-making behaviors**

Rich Pang  
Adrienne L Fairhall  
University of Washington

RPANG@UW.EDU  
FAIRHALL@U.WASHINGTON.EDU

Decision-making involves a change of behavior in response to sensory cues. A first step in developing models of sensory-driven decision-making that are relevant to natural behavior is to identify decision points in natural behavioral time-series. This allows the representation of behavior as a sequence of transitions among a discrete set of states. In this project we apply a data-driven method for identifying the number and properties of distinct states in a behavioral time-series. While we emphasize the generality of the method, we demonstrate its utility in identifying the dynamical properties of experimentally measured free- flying mosquito trajectories. Specifically, we represent each trajectory using a variant of the hidden Markov Model (HMM), whose details we learn using the sticky HDP-HMM, an algorithm for parameter inference that allows robust representations of state dynamics as well as agnosticism about the number of distinct states. Using this method, we segment mosquito flight into behavioral motifs with differing dynamical properties. Further, by generalizing the HMM to allow more variable dynamics, we

develop an algorithm with increased predictive capability, and which discriminates with high accuracy between the flight patterns of mosquitos and *Drosophila melanogaster* using little training data. Since this approach requires no domain-specific knowledge, we suggest its utility in analyzing diverse naturalistic time-series.

## I-51. Uncovering stimulus-induced network dynamics during narrative comprehension

Erez Simony  
Chris Honey  
Janice Chen  
Uri Hasson

Princeton University

EREZSIM@GMAIL.COM  
CHRISTOPHER.HONEY@GMAIL.COM  
KANILE@GMAIL.COM  
HASSON@PRINCETON.EDU

Motivation: The goal of functional connectivity (FC) analyses is to map the pattern of large-scale functional interactions in the brain, and how they change over time and across tasks. However, it is difficult to separate the different factors that give rise to FC : these include physiological and instrumental noise, stimulus-induced cognitive processes and stimulus-independent cognitive processes. We therefore introduce an approach (inter-subject functional correlation, ISFC) that isolates the inter-regional correlations that are driven by stimulus-induced processes. Our goal is to use ISFC to map brain networks dynamics that are sensitive to the content of real-world stimuli. We therefore ask: (i) Do the ISFC patterns induced by naturalistic stimuli resemble FC patterns observed in the resting state? (ii) Can we reliably track ISFC network dynamics over time? (iii) How do these network dynamics relate to the semantics of the stimulus? Methods: BOLD responses were measured as thirty-six subjects listened to a full 7-minute auditory narrative, and versions of the story in which different levels of coherence were preserved by reordering of words or of paragraphs. We calculated ISFC in each condition by correlating the response timecourse in one brain area in one individual against the response timecourses in all brain areas of all other individuals. All analyses were replicated in a second set of subjects. Results: Our results demonstrate that (i) stimulus-induced brain networks resemble those observed during the resting state; (ii) we can precisely and replicably track stimulus-induced network states over time; (iii) network dynamics are unique to a story, and the overall reliability of network dynamics in higher-order regions increases with the semantic coherence of the stimulus. Conclusions: Stimulus-induced network dynamics, mapped using ISFC, precisely track large-scale brain states that are sensitive to the semantic content of real-world stimuli.

## I-52. Measuring memory formation in the human brain using cross stimulus suppression

Helen Barron<sup>1</sup>  
Raymond Dolan<sup>1</sup>  
Timothy Behrens<sup>2</sup>

<sup>1</sup>University College London  
<sup>2</sup>University of Oxford

HELEN.C.BARRON@GMAIL.COM  
R.DOLAN@UCL.AC.UK  
BEHRENS@FMRIB.OX.AC.UK

Memory formation is a dynamic process that recruits multiple brain regions. Whilst it has been characterized at the cellular level using spike-timing dependent plasticity, the temporal interactions within and across brain regions are not well understood. Using functional magnetic resonance imaging to measure repetition suppression across two distinct stimuli, we have developed a technique that allows us to quantify the overlap between neural representations in the human brain and thus track the emergence of new associations. Using a range of different stimuli we find that plasticity is functionally localized to those regions of cortex which typically encode features of the paired items. For example, when primary rewards are repeatedly combined together plasticity occurs in pre-frontal cortex, whilst new associations between line drawings and orientation-invariant shapes result in plasticity in

occipital cortex and temporal lobe regions respectively. In addition to these localized cortical associations, a more generic plasticity effect is observed in hippocampus, regardless of the stimulus modality. In contrast to the cortex, the magnitude of this hippocampal plasticity is value dependent. Following associative learning, the hippocampus also exhibits a different temporal profile to the cortex: whilst repetition suppression is observed in the cortex for a limited time period, hippocampal plasticity emerges later, increasing over time. Using repetition suppression to track the associative strength of neural representations, we therefore demonstrate how the dynamics of memory formation can be measured, and the interplay between brain regions assessed.

## I-53. Encoding Certainty in Bump Attractors

Sam Carroll<sup>1,2</sup>  
Kresimir Josic<sup>2</sup>  
Zachary Kilpatrick<sup>2</sup>

SRCARROLL314@GMAIL.COM  
JOSIC@MATH.UH.EDU  
ZPKILPAT@MATH.UH.EDU

<sup>1</sup>University of Utah

<sup>2</sup>University of Houston

A number of experiments have shown that, in various cortical areas, the certainty of a decision can be encoded by the instantaneous firing rates of neurons. Therefore the certainty about a decision, as well as the decision itself can be represented in the activity patterns of a population. Similar observations have been made in tasks involving spatial working memory. Recordings from superior colliculus suggest that an increase in uncertainty in a remembered cue position is represented by reduced neural firing rates during the delay period in an oculomotor delayed response task. Conversely, it has been shown that training in a spatial working memory task that leads to improved performance is accompanied by a rise in firing rates. These observations suggest that certainty about stimulus location in a spatial working memory task may be represented by the level of neural activity during the time the location is in memory. We extend previous work on bump attractors to construct model networks capable of encoding the certainty of a stimulus stored in memory. Such networks support bumps that are not only neutrally stable to perturbations in position, but also perturbations in amplitude. In this model, bump solutions lie on a two-dimensional attractor determined by a continuum of positions and amplitudes. Since the bump's position represents the remembered stimulus location and its amplitude encodes certainty, this attractor provides an analog representation of location and certainty. Such an attractor requires precisely balancing the strength of recurrent synaptic connections. The amplitude of activity bumps is determined by the strength of the initial input to the system, and could thus represent certainty about the stimulus. Moreover, bumps with larger amplitudes are more robust to noise. Generating networks whose bumps can take on a continuum of amplitudes requires that the network's synaptic inhibition precisely cancel excitation.

## I-54. Stable population coding for working memory in prefrontal cortex

John Murray<sup>1</sup>  
Nicholas Roy<sup>2</sup>  
Ranulfo Romo<sup>3</sup>  
Christos Constantinidis<sup>4</sup>  
Alberto Bernacchia<sup>5</sup>  
Xiao-Jing Wang<sup>1</sup>

JOHN.MURRAY@NYU.EDU  
NICHOLAS.ROY@YALE.EDU  
RROMO@IFC.UNAM.MX  
CCONSTAN@WAKEHEALTH.EDU  
A.BERNACCHIA@JACOBS-UNIVERSITY.DE  
XJWANG@NYU.EDU

<sup>1</sup>New York University

<sup>2</sup>Yale University

<sup>3</sup>Universidad Nacional Autonoma de Mexico

<sup>4</sup>Wake Forest University

<sup>5</sup>Jacobs University Bremen



Working memory requires conversion of brief stimulus-driven signals into internal representations that can be maintained across mnemonic delays of several seconds. In primate cortex, electrophysiological studies find stimulus-selective persistent activity in single neurons as neural correlates of working memory. However, recent studies have highlighted temporal variations in delay activity, at single-neuron and population levels. It remains unclear how neuronal populations maintain memory of stimuli despite complex and heterogeneous temporal dynamics. To address this question, we applied population-level analyses to hundreds of single-neuron recordings from primate prefrontal cortex, during two classic working-memory tasks: oculomotor delayed response and vibrotactile delayed discrimination. Both tasks demand working memory of analog stimuli, but they differ in stimulus properties, behavioral response, and neural tuning. We describe the dynamics of neuronal population in terms of its temporal trajectory in a high-dimensional space. Despite strong temporal dynamics and heterogeneity across neurons, we found a low-dimensional subspace, via Principal Components Analysis, in which stimulus coding is stable across the cue and delay. Furthermore, we found that sensory and mnemonic representations are temporally overlapping, rather than sequential. These results, similar for both tasks, suggest working-memory representations are stable and raise the question of which neural mechanisms support working-memory encoding, maintenance, and retrieval. To explore possible mechanisms, we analyzed the activity of several neural circuit models proposed to explain physiological observations in working-memory tasks, including “attractor” models and “random” networks. Each model could explain some, but not all, of the experimental observations: attractor models show low-dimensional, stable memory but cannot account for dynamic trajectories, whereas random networks show complex trajectories but have high-dimensional memory representations. Our work demonstrates the potential of population-level analyses for testing computational models of working memory, and suggests that cortical circuits may operate at an interface between order and disorder that currently lacks an appropriate description.

## I-55. Parametric working memory in rats: sensory vs. Prefrontal cortex

Vahid Esmaeili<sup>1</sup>  
Arash Fassihi<sup>1</sup>  
Athena Akrami<sup>2</sup>  
Mathew Diamond<sup>1</sup>

VAHID.ESMAEILI@SISSA.IT  
FASSIHI@SISSA.IT  
ATHENA.AKRAMI@GMAIL.COM  
DIAMOND@SISSA.IT

<sup>1</sup>International School for Advanced Studies

<sup>2</sup>Princeton University

Parametric Working Memory (pWM), the short-term storage of graded stimuli to guide behavior, has been studied in primates, yet its neuronal mechanisms are not fully elucidated. To address this, we have devised a task where rats compare two sequential stimuli, ‘base’ and ‘comparison’, applied to their whiskers separated by a delay. Stimuli are vibrations, generated as a series of velocity values sampled from a zero-mean normal distribution. Rats should judge which stimulus had greater velocity standard deviation. Previously, we have shown that rats’ ability to hold base stimulus information in working memory for subsequent comparison is similar to that of humans, who receive vibrations on the fingertip. To explore where and how stimulus information is held in the brain we have simultaneously recorded from single neurons in barrel cortex — where whisker information first enters the cortex — and medial prefrontal cortex (mPFC) which has been suggested to be involved in high-level cognitive and emotional processes. We demonstrate that neurons in barrel cortex encode base velocity standard deviation (base sigma) during the presentation of the base stimulus but not during the inter stimulus interval. Barrel cortex neurons increase their discharge rate for increasing values of base sigma (a coding scheme we denote as positive monotonic). In mPFC, we report diverse coding characteristics between and within neurons. Some neurons encode task-related events (such as reward collection) while others encode base sigma. Neurons that encode base sigma do so in a different manner in different stages of the task. They encode base sigma with a positive monotonic function during the presentation of the base stimulus yet some neurons exhibit a negative monotonic relationship to base sigma during the delay. This is, to our knowledge, the first demonstration of neuronal signatures of pWM in rodents.

## I-56. Bayes optimal spike adaptation predicts biophysical characteristics of neurons

Alessandro Ticchi<sup>1,2</sup>

Aldo A. Faisal<sup>1</sup>

<sup>1</sup>Imperial College London

<sup>2</sup>University of Bologna

A.TICCHI12@IMPERIAL.AC.UK

A.FAISAL@IMPERIAL.AC.UK

Sensory input signals can span many orders of magnitude while dynamically changing, and this requires neurons to adapt their bandwidth-limited response behaviour efficiently. We derive from first principles a Bayesian model of neural adaptation to the instantaneous input statistics, and show how existing intracellular data can be explained by its biophysical implementation. We assume that cortical neurons' spiking activity can be described by inhomogeneous Poisson processes, whose instantaneous rate depends both on the received input and on the internal adaptation state. We show how neurons can exploit the information contained in their own spiking history to infer the instantaneous input statistics and to adapt their input-output behaviour accordingly in order to optimise information transfer. Specifically, we prove that the posterior of the instantaneous input rate can be well approximated using a Gamma distribution for various considered temporal dynamics of the input statistics. We derive temporal update equations for the parameters of this distribution. Under the ecologically reasonable assumption that the input stimulus changes on a timescale longer than the average inter spike interval, we find the dynamics of the update equations to mimic the known time course of  $Ca^{++}$  and of some neurotransmitter concentrations. We argue that these quantities can encode the updated belief about the instantaneous input statistics and can regulate neural response accordingly. We also show how the information encoded by a presynaptic Bayesian adapting neuron can be decoded by postsynaptic neurons by weighted temporal integration. Using our first principles approach we link the temporal statistics of the stimulus with the dynamics of neural adaptation. The neural mechanism we propose allows a stochastically spiking neuron to implement temporal predictive coding and, when implemented by both the pre and postsynaptic neurons of a cortical population, produces global adaptation behaviours at the population level which mimic experimental results.

## I-57. Low-dimensional functionality of connectome dynamics: Neuro-sensory integration in *C. elegans*

James Kunert

Eli Shlizerman

J. Nathan Kutz

University of Washington

KUNERT@UW.EDU

SHLIZEE@UW.EDU

KUTZ@UW.EDU

The nematode *Caenorhabditis elegans* is an ideal model organism in the study of neuro-sensory networks, since the connectivity of its entire 302-neuron network (its connectome) has been resolved. Experimental studies have shown through PCA that the network's behavioral responses to stimuli are fundamentally low-dimensional. A description of these responses cannot be drawn from static connectivity data alone, but these studies suggest that computational modeling of network voltage dynamics can assist in relating system inputs to downstream responses. Our study computationally relates, through dynamical full-network simulation of the *C. elegans* connectome, input stimuli to the low-dimensional modal responses known to correspond to behavioral output. Building on recent experimental findings, we posit a full dynamical model of the neuron voltage excitations that allows for a characterization of robustly encoded behavioral responses to input stimuli. Simulations of this model demonstrate, through SVD analysis, the low-dimensional structure of motor neuron responses to given inputs. We then demonstrate through stability analysis that low-dimensional bifurcation structures are ultimately responsible for the appearance of these neural pathways of activity. As a specific example, we observe in our simulation the desired two-mode response to tail-touch. In doing so, we not only show the consistency of our model with observation, but in the process we discover the underlying bifurcation structures which play a critical role in explaining how such responses are encoded within the network. Specifically, key neurons activate a Hopf bifurcation which

induces an oscillatory wave in motor neurons associated with crawling/swimming. Since this study simulates the full network with no assumptions specific to the response in question, its techniques may be extended to the study of any response encoded within similar bifurcation structures.

## I-58. Low-dimensional dynamics underlie auditory cortex responses for different brain states and stimuli

Marius Pachitariu  
Maneesh Sahani  
Nicholas Lesica

MARIUS@GATSBY.UCL.AC.UK  
MANEESH@GATSBY.UCL.AC.UK  
N.LESICA@UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

Many experimental observations point to a significant role of intracortical dynamics in shaping sensory cortical responses—the temporally-extended response to a short stimulus, the spatiotemporal correlations of both spontaneous activity and stimulus-response variability, and the common alternation between periods of widespread activity and quiet (“up” and “down” states). What form best describes the dynamical systems that underlie these phenomena? How do they shape the corresponding neural responses? Here we show that cortical activity under a range of brain states and stimulus conditions is influenced by a common low-dimensional dynamical structure. We analyzed both spontaneous activity and responses to tones and speech from populations of primary auditory cortical neurons in both synchronized and desynchronized brain states. Noise correlations were large and positive (0.1-0.3) in the synchronized state, due to population-wide up/down transitions, and near zero in the desynchronized state (0-0.03), but stereotypical spatiotemporal patterns of spiking were evident in both. We designed a class of dynamical models that captured the statistical structure in the population spiking. These models identified slow dynamics controlling the timing of up and down states, and a faster system controlling the pattern of activity within individual response events. The low-dimensional dynamical structure was largely conserved across states and stimulus conditions, accounting for both spontaneous and driven correlations and the spatiotemporal pattern of both synchronized and desynchronized spiking. Notably, the slow system remained active in the desynchronized state where systematic up/down shifts were few. Rapid, reliable stimulus responses in the desynchronized state clearly rode on top of a dynamically-shaped response. These simple dynamical models thus unify a range of complex cortical phenomena. The similarities across brain states and stimulus conditions in auditory cortex, as well as previous work in motor cortex, suggest that low-dimensional dynamics may underlie much coordinated neural activity throughout the cerebrum.

## I-59. The dimensionality of dendritic computation

Lei Jin<sup>1</sup>  
Bardia Behabadi<sup>2</sup>  
Bartlett W. Mel<sup>1</sup>

LEIJ@USC.EDU  
BARDIAB@QTI.QUALCOMM.COM  
MEL@USC.EDU

<sup>1</sup>University of Southern California

<sup>2</sup>Qualcomm Research

On the path to understanding cortical processing, it is essential to develop good simplified models of the principal cell type, the pyramidal neuron (PN). Both experimental and modeling studies have shown that PN dendrites can act as independent computational subunits with sigmoidal input-output functions, though more recent evidence indicates that a 1-D sigmoid is an inadequate description of a PN dendrite’s spatial processing capabilities. First, the somatic response to a single excitatory input depends on its location, producing a family of different sigmoids increasing in threshold and amplitude with increasing distance from the soma. For two spatially-separated inputs, a dendrite’s input-output function becomes a family of 2-D sigmoidal surfaces, whose peculiar form suggests that PN dendrites could provide a substrate for nonlinear classical-contextual interactions. In this work, we have de-

veloped a dimensionality-reduction approach that allows us to predict a dendrite's time-averaged response to an arbitrary spatial pattern of excitation, and as a side benefit, to determine the maximum dimensionality of spatial processing possible within a PN dendrite. We defined a family of  $N$  spatial basis functions covering the length of a dendrite, each one representing the density of excitation delivered to that region of the dendrite. Using biophysically detailed compartmental simulations, we estimated the number of basis functions needed to accurately predict the dendritic response for arbitrary spatially varying excitation patterns. We found that a PN thin dendrite has, within its  $\sim 200$   $\mu\text{m}$  length, roughly  $N=3$  independent dimensions of spatial analog processing capability, meaning that a branch's input-output function is in effect a 3-D rather than a 1-D sigmoid. This measurement of the true dimensionality of dendritic spatial processing will facilitate the interpretation of experimental results involving active dendritic processing, and help guide the design of future analog neuromorphic hardware systems.

## I-60. A probabilistic modeling approach for uncovering neural population rotational dynamics

Shamim Nemati<sup>1</sup>  
Scott Linderman<sup>1</sup>  
Zhe Chen<sup>2</sup>

SHAMIM@SEAS.HARVARD.EDU  
SCOTT.LINDERMAN@GMAIL.COM  
ZHECHEN@MIT.EDU

<sup>1</sup>Harvard University

<sup>2</sup>Massachusetts Institute of Technology

An important task in computational neuroscience is to understand neural representations of population codes, using either Bayesian approach or dynamical systems approach. We provide a probabilistic interpretation of the jPCA algorithm, recently proposed by Churchland et al. [1] for revealing the rotational patterns and dynamics in neuronal population activity. Using the jPCA technique, the authors showed that neural firing rates during reaching tasks (a non-rhythmic activity) exhibit a quasi-periodic pattern. Here we show that the original jPCA algorithm can be formulated as a linear algebra problem, and therefore be solved using the well-known technique of polar decomposition. Moreover, we show that the technique can be cast into a probabilistic framework which allows for application of Bayesian methods for model identification and parameter fitting. Moreover, missing data also can be handled within this framework using an expectation-maximization approach. Our approach provides an alternative yet more robust approach for dimensionality reduction and visualization in analyzing noisy neural population codes.

## I-61. Orthogonal representation of tasks parameters in higher cortical areas

Wieland Brendel<sup>1</sup>  
Dmitry Kobak<sup>1</sup>  
Ranulfo Romo<sup>2</sup>  
Claudia Feierstein<sup>1</sup>  
Zachary Mainen<sup>1</sup>  
Christian Machens<sup>1</sup>

WIELAND.BRENDEL@NEURO.FCHAMPALIMAUD.ORG  
DMITRY.KOBAK@NEURO.FCHAMPALIMAUD.ORG  
RROMO@IFC.UNAM.MX  
CLAUDIA.FEIERSTEIN@NEURO.FCHAMPALIMAUD.ORG  
ZMAINEN@NEURO.FCHAMPALIMAUD.ORG  
CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

<sup>1</sup>Champalimaud Neuroscience Programme

<sup>2</sup>Universidad Nacional Autonoma de Mexico

Modern experimental techniques allow to record the activity of hundreds or even thousands of neurons simultaneously. Despite these advances, most experimental studies still focus on single-cell responses thereby neglecting much of the distributed information encoded in the population response. This problem is particularly severe in higher order areas such as the prefrontal cortex, where single-cell responses can display a baffling heterogeneity, even if animals are carrying out rather simple tasks. The standard approach to analyzing population activity is by performing principal component analysis (PCA), but its outcome is often hard to interpret. Recent developments

in the field have sought to enhance the sensitivity of dimensionality reduction approaches in order to discover and highlight the most “interesting” (for the researcher) features of the data (e.g. jPCA, Churchland, Cunningham et al., 2012; dPCA, Brendel et al., 2011; Machens, 2010). Here we use demixed PCA (dPCA) to analyze the population activity in several datasets, comprising different animals (rats and monkeys), different higher cortical areas (prefrontal cortex, orbitofrontal cortex, somatosensory cortex) and different experimental tasks (tactile discrimination, odour discrimination, working memory task). Our analysis reveals that the population response in higher cortical areas exhibits a strikingly stereotypical structure. More precisely, we show that the population reliably encodes behaviourally relevant task parameters (such as stimulus, decision, or reward) in completely independent orthogonal subspaces. This observation leads to two important insights. First, from an experimental point of view, it simplifies data analysis and provides a concise way of data visualization, often allowing to summarize the most relevant features of the neural responses in a single plot. Second, from a computational point of view, it serves as a constraint on neural network models of higher cortical areas as we demonstrate on echo state networks and liquid state machines.

## I-62. Co-fluctuation patterns of population activity in macaque primary visual cortex

Benjamin Cowley<sup>1</sup>  
 Matt Smith<sup>2</sup>  
 Adam Kohn<sup>3</sup>  
 Byron Yu<sup>1</sup>

BCOWLEY@ANDREW.CMU.EDU  
 SMITHMA@PITT.EDU  
 ADAM.KOHN@EINSTEIN.YU.EDU  
 BYRONYU@CMU.EDU

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>University of Pittsburgh

<sup>3</sup>Albert Einstein College of Medicine

We seek to understand how populations of neurons in primary visual cortex (V1) process visual stimuli. One way of studying population activity is to extract co-fluctuation patterns, which characterize how neurons increase and decrease their firing rates together. A fundamental question is whether the population response to different visual stimuli use i) the same co-fluctuation patterns, but combine them in different ways, or ii) different co-fluctuation patterns. These two possibilities cannot be distinguished by examining the activity of each neuron individually. To study this question, we simultaneously recorded the response of 61 to 81 neurons to sinusoidal gratings, a natural movie, and white noise, as well as spontaneous activity, using Utah arrays implanted in V1 of two anesthetized macaque monkeys. First, we applied principal component analysis (PCA) to trial-averaged population activity to identify the co-fluctuation patterns for each stimulus. We found that the patterns for gratings and natural movies were dissimilar, but the patterns for both stimuli were subsets of those for noise. Second, we asked how the trial-to-trial variability of the population response to gratings, natural, and noise movies relates to spontaneous activity. We subtracted the trial-averaged population activity from single-trial population activity to obtain single-trial residuals. We found that the co-fluctuation patterns of spontaneous activity closely resembled those of the residuals. Previous work using functional models of single neurons suggest that simple (i.e., gratings and noise) and natural stimuli drive fundamentally different modes of processing in V1. Our results indicate that, while this may be true for gratings, natural stimuli instead drive a subset of modes driven by noise. Furthermore, these results suggest that the trial-to-trial variability of evoked activity arises from similar underlying neural mechanisms that govern spontaneous activity. Our analysis framework provides a new way to elucidate how V1 processes visual stimuli on the population level.

## I-63. Neurons in V4 show quadrature pair selectivity along curved contours

Ryan Rowekamp  
Tatyana Sharpee  
Salk Institute for Biological Studies

RROWEKAMP@SALK.EDU  
SHARPEE@SALK.EDU

As visual signals are processed at successive stages of visual processing, neurons pool information from larger sections of the visual space and also signal the presence of increasing complex visual features. Both of these properties increase the number of stimulus features that modulate a neuron's response. We demonstrate a new algorithm that models a neuron's response as the weighted response of identical subunits shifted across the visual space while allowing for nonlinear computations to take place within each subunit. The subunits incorporate a quadratic filter sensitive to pairwise correlations in the stimulus. By learning the weights associated with the subunits, the model can determine the extent of the invariance of the neuron's response as well as characterize non-invariant computations. We applied this algorithm to the responses of macaque neurons from visual area V4 to natural movie stimuli. We found that the stimulus features associated with excitation of the neuron often formed quadrature pairs with similar positions, orientations, and spatial frequencies but with orthogonal spatial phases. These pairs are positioned such that they form curved contours. This extends Hubel and Wiesel's description of complex cells to curves.

## I-64. Transition to chaos in heterogeneous networks

Johnatan Aljadeff<sup>1,2</sup>  
Merav Stern<sup>3,4</sup>  
Tatyana Sharpee<sup>1</sup>  
<sup>1</sup>Salk Institute for Biological Studies  
<sup>2</sup>University of California, San Diego  
<sup>3</sup>Columbia University  
<sup>4</sup>Hebrew University

ALJADIEFF@UCSD.EDU  
MERAV.STERN@MAIL.HUJI.AC.IL  
SHARPEE@SALK.EDU

There is accumulating evidence that biological neural networks possess optimal computational capacity when they are poised at or near a critical point where the network transitions to a chaotic regime. A majority of theoretical models that predict when a network becomes chaotic consider homogeneous (or nearly homogeneous) networks, but real neural networks are heterogeneous with respect to cell types and connectivity patterns. We derive a mathematical condition for networks with heterogeneous connectivity rules to become chaotic. We apply our results to networks where adult neurogenesis occurs and to clustered networks and show that our theoretical model may provide insight into their structure-function relationship.

## I-65. Stability of neuronal networks with homeostatic regulation

Miha Pelko<sup>1</sup>  
Daniel Harnack<sup>2</sup>  
Antoine Chaillet<sup>3,4</sup>  
Yacine Chitour<sup>3,4</sup>  
Mark van Rossum<sup>1</sup>  
<sup>1</sup>University of Edinburgh  
<sup>2</sup>University of Bremen  
<sup>3</sup>Supélec  
<sup>4</sup>Université Paris Sud

MIHA.PELKO@ED.AC.UK  
DANIEL@NEURO.UNI-BREMEN.DE  
ANTOINE.CHAILLET@SUPELEC.FR  
YACINE.CHITOUR@LSS.SUPELEC.FR  
MVANROSS@INF.ED.AC.UK

Homeostasis is thought to maintain the activity of neurons in a regime which is both biologically healthy and optimal for efficient information processing. Several studies have addressed homeostasis in single neurons using simple feedback cascades. Here we systematically study the effect of homeostatic feedback control of intrinsic excitability on the stability of networks. We first derive the stability criterion for the case of linear rate-based networks and next explore the effects in recurrent excitatory networks of leaky integrate and fire neurons. Our results reveal previously unconsidered constraints on homeostasis in biological networks.

## I-66. Evidence towards surround suppression in perception of complex visual properties

Daniel Leeds<sup>1</sup>  
John Pyles<sup>2</sup>  
Michael Tarr<sup>2</sup>

DLEEDS@FORDHAM.EDU  
JPYLES@CMU.EDU  
MICHAELTARR@CMU.EDU

<sup>1</sup>Fordham University

<sup>2</sup>Carnegie Mellon University

Object perception recruits a cortical network that encodes a hierarchy of increasingly complex visual features. While the earliest stages of human vision have been reasonably well-modeled using local oriented edges, the encoding of visual properties in higher-level cortical regions is less well understood. Prior work has been able to characterize neural selectivity for mid-level features in IT and V4 using simple parts-based, parametric approaches (Hung 2012, Yamane 2008). Although these studies identify individual shapes maximizing neural firing for selected neurons, they provide little insight as to the principles used to distinguish preferred stimuli from anti-preferred stimuli. Here, we explore the grouping of properties salient to cortical perception by defining and using Euclidean feature spaces based on complex visual properties of natural and synthetic object classes. Natural objects are drawn from real-world photographs and are projected into a space reflecting visual grouping based on a SIFT bag-of-words distance metric (Nowak 2006). Synthetic objects are Fribbles (Williams 2000) — animal-like objects composed of geometric shapes — and projected into a space capturing the morphs between Fribble appearances. These visual feature spaces present a new tool for analysis of cortical activity underlying object perception. Using fMRI and pre-selected regions of the ventral visual cortex, we find that object stimulus pairs near one another in feature space can produce responses at the opposite extremes of the measured activity range. Suppression of response to properties slightly altered from preferred properties is observed for different object classes and centered on different locations in a given feature space. These observations may demonstrate extension of surround suppression observed in lower levels of vision — such as the Gabor-like models of V1 receptive fields (Kay 2008) and studies of location-based suppression effects for visual texture perception (Freeman 2011).

## I-67. Structured chaos shapes population noise entropy in driven balanced networks

Guillaume Lajoie<sup>1</sup>  
Jean-Philippe Thivierge<sup>2</sup>  
Eric Shea-Brown<sup>1</sup>

GLAJOIE@UW.EDU  
JEAN-PHILIPPE.THIVIERGE@UOTTAWA.CA  
ETSB@AMATH.WASHINGTON.EDU

<sup>1</sup>University of Washington

<sup>2</sup>University of Ottawa

We study population level neural responses in spiking networks with sparse, random connectivity and balanced excitation and inhibition. Such networks reproduce the irregular firing that typifies cortical activity; moreover, this activity is known to be chaotic, with extremely strong sensitivity of spike outputs on tiny changes in a network's initial conditions [VanVreeswijk and Sompolinsky, 1996, Monteforte and Wolf, 2010, London et al., 2010]. Nev-

ertheless, when these networks receive temporally fluctuating stimuli, our previous work shows that underlying chaotic attractors of limited dimension lead to reduced spiking variability at the single-cell level [Lajoie et al., 2013]. However, recent studies suggest that the impact of noise on network coding cannot be understood by extrapolating from single cell variability alone [Schneidman et al., 2006, Ecker et al., 2011]. Thus, in this work we study the joint activity of entire chaotic networks. We ask: How variable and noisy is this joint spiking activity? What are the network-level sources of this variability, and how might they constrain spiking features relevant for coding? To answer these questions, we extend the work of [Monteforte and Wolf, 2010] to derive a bound for the entropy of joint spike pattern distributions in large spiking networks that receive fluctuating input signals. Our analysis is based on results from random dynamical systems theory and is complimented by detailed numerical simulations. We find that despite very weak conditional correlations between neurons, the resulting joint variability of network responses is surprisingly lower than what would be predicted from measurements of single cells. Moreover, the variability of the network as a whole is strongly controlled by the statistics of the input stimuli.

## I-68. Predicting noise correlations for non-simultaneously measured neuron pairs

Srinivas C Turaga<sup>1</sup>  
 Lars Buesing<sup>1</sup>  
 Adam Packer<sup>2</sup>  
 Henry Dagleish<sup>2</sup>  
 Noah Pettit<sup>2</sup>  
 Michael Hausser<sup>2</sup>  
 Jakob Macke<sup>3,4</sup>

STURAGA@GATSBY.UCL.AC.UK  
 LARS@GATSBY.UCL.AC.UK  
 A.PACKER@UCL.AC.UK  
 HENRY.DALGLEISH.09@UCL.AC.UK  
 NOAH.PETTIT.10@UCL.AC.UK  
 M.HAUSSER@UCL.AC.UK  
 JAKOB.MACKE@GMAIL.COM

<sup>1</sup>Gatsby Computational Neuroscience Unit, UCL

<sup>2</sup>University College London

<sup>3</sup>Max Planck Institute

<sup>4</sup>Bernstein Center for Computational Neuroscience, Tuebingen

Simultaneous recordings of the activity of large neural populations are extremely valuable as they can be used to infer the dynamics and interactions of neurons in a local circuit, shedding light on the computations performed. It is now possible to measure the activity of hundreds of neurons using in-vivo 2-photon calcium imaging. However, many computations are thought to involve circuits consisting of thousands of neurons, such as cortical barrels in rodent somatosensory cortex. Here we contribute a statistical method for “stitching” together sequentially imaged sets of neurons into one model by phrasing the problem as fitting a latent dynamical system with missing observations. This method allows us to substantially expand the population sizes for which population dynamics can be characterized—beyond the number of simultaneously imaged neurons. We describe conditions for successful stitching and use recordings from mouse somatosensory cortex to demonstrate that this method enables accurate predictions of noise correlations between non-simultaneously recorded neuron pairs.

## I-69. Statistical complexity of neural spike trains

Sarah Marzen<sup>1</sup>  
 Michael DeWeese<sup>1</sup>  
 Jim Crutchfield<sup>2</sup>

SMARZEN@BERKELEY.EDU  
 DEWEESE@BERKELEY.EDU  
 CHAOS@CSE.UCDAVIS.EDU

<sup>1</sup>University of California, Berkeley

<sup>2</sup>University of California, Davis

We present closed-form expressions for the entropy rate, statistical complexity, and predictive information for the spike train of a single neuron in terms of the first passage time probability distribution. Our analysis applies to



any one-dimensional neural model where observation of a spike causes the neuron to “reset” to some membrane voltage and in which any noise term is uncorrelated in time. We then use these formulae to study the linear leaky integrate-and-fire and quadratic integrate-and-fire neurons driven by white noise in the naturally spiking regime. The statistical complexity is simply related to the interspike interval’s mean and coefficient of variation. At high coefficients of variation, the neural spike trains of a quadratic integrate-and-fire neuron are more compressible than that of a linear leaky integrate-and-fire neuron. This suggests a new computational benefit for spike generating mechanisms—a decrease in required memory storage of “predictor” neurons.

## I-70. Neuroembeddology: dynamical relation reflects network structures

Satohiro Tajima

Taro Toyoizumi

RIKEN Brain Science Institute

SATOHYRO.TAJIMA@RIKEN.JP

TARO.TOYOIZUMI@RIKEN.JP

Relating the neural dynamics to its anatomical/functional network structure is a central issue in neuroscience. However, this remains a challenging problem for complex systems like the brain. The major difficulty is that conventional methods infer network structure based on a correlational measure. Here, we propose a methodology to infer a causal network structure based on nonlinear forecasting (Sugihara, 2012); the key idea is to explore embedding relationships among network nodes — whether information about a node is extracted by observing another node. Although the embedding is a widely accepted concept in dynamical systems, its use in causal-network estimation is underexplored. Here, we bridged this gap in three steps. First, we provided the information theoretic interpretation of a recently proposed embeddedness measure based on nonlinear forecasting (Sugihara, 2012). We related it to the transfer entropy (TE), an information theoretic analogue of Granger causality (GC). Based on this relation, we demonstrated that embedding measure is able to detect the causality even when TE (and GC) fails to detect it due to severe confounding effects of self-predictability. Second, we extracted the network structure in simulated networks with several different types of dynamical interactions, establishing that cross-embedding relationships effectively capture the true network structures. Third, we applied the embedding analysis to monkey ECoG data. In contrast to a previous study using TE/GC, we found clear asymmetric embedding relationships among cortical regions, which are the hallmark of causal interactions. Moreover, we found a profound difference in such relationships as the brain state changed from an awake to anesthetized condition, providing a structural correlate of consciousness at a whole-brain level. The present theoretical and numerical analyses as well as its application to real data demonstrated a potential of the embedding-based method to be a breakthrough for investigating a wide variety of neurophysiological dynamics.

## I-71. Dynamics of excitatory-inhibitory neuronal networks with exponential synaptic weight distributions

Ramakrishnan Iyer

Nicholas Cain

Stefan Mihalas

Allen Institute for Brain Science

RAMAKRISHNANIYE@GMAIL.COM

NICHOLASC@ALLENINSTITUTE.ORG

STEFANM@ALLENINSTITUTE.ORG

Dynamics of interacting leaky integrate-and-fire (LIF) neuronal networks have been widely studied, both analytically and numerically, with a view to characterizing experimentally observed states of activity in the brain. Analytic approaches typically involve Fokker-Planck methods, based on a continuous stochastic process, that model dynamics of homogeneous neuronal populations with a single partial differential equation. This method abstracts distributions of synaptic weights by their first two moments and works well when higher moments can be neglected. Such an analysis revealed a rich repertoire of states of activity and allowed characterization of phase diagrams of coupled excitatory-inhibitory LIF networks [1] (Brunel 2000). Experiments have observed skewed

non-Gaussian synaptic weight distributions with long tails, such as lognormal [2] (Song et al, 2005). A fast, semi-analytic population statistic approach to study stochastic jump processes with arbitrary distributions [3] (Iyer et al, 2013) showed that equilibrium and transient population firing rates vary according to how heavy-tailed the distributions are. Steady state and transient firing rates in feed-forward networks with exponential synaptic weight distributions have been obtained analytically in [4] (Richardson et al, 2010). We generalize their approach to obtain coupled transcendental equations for stationary rates in coupled excitatory-inhibitory LIF networks, with and without synaptic depression. These provide a quick estimate of the steady state firing rates. We thus have a tool for rapid exploration of parameter space for large scale models, with realistic synaptic weight distributions, in which the observables of interest are the mean firing rates of the populations. Parameter-fitting for interesting network behaviors, which is a hard problem for network models, can be simplified with the use of this method.

## I-72. Predictive coding in active and quiescent states

Veronika Koren<sup>1</sup>  
 Sophie Deneve<sup>1</sup>  
 David G.T. Barrett<sup>1,2</sup>

KOREN.VERONIKA@GMAIL.COM  
 SOPHIE.DENEVE@ENS.FR  
 DAITHIOBAIREAD@GMAIL.COM

<sup>1</sup>Ecole Normale Supérieure

<sup>2</sup>Champalimaud Centre for the Unknown

Animals spend a great deal of time in a quiescent state, either in slow wave sleep or resting. This state is characterised by synchronised bursts of spiking activity (up-states) interspersed by periods of silence (down-states). We ask, what is the function of the quiescent state, and what are the network mechanisms that support its function? Our approach is to treat the quiescent state as a special case of the active state. Recent work has found that the active state supports predictive coding - using a network of integrate-and-fire neurons that are tuned to encode the difference between a signal and a predicted signal. We find that when this predictive coding network is representing a silent signal, or a null-signal, it produces quiescent state activity. Down-state's occur because the best way for a network to represent a silent signal is to produce no spikes at all. Up-states are triggered by neural noise, which randomly causes a few neurons to spike. This activity propagates through the network, causing more neurons to fire, until a full up-state is produced. The escalation of activity corrects for the error introduced by the noise, by recruiting neurons with opposite tuning to the noise driven neurons. Eventually, the prediction error is corrected and the network decays back into a down-state. This similarity, between the spiking activity of a predictive coding network and the spiking activity observed in the quiescent state suggests that the function of this state is to represent a silent null-signal using a noisy predictive coding mechanism. This functional interpretation of quiescent state activity is consistent with the resting behaviour of animals in the quiescent state. Furthermore, the network mechanism underlying this state is particularly natural because it is identical to the mechanism used for complex computations and active state dynamics.

## I-73. A novel point process approach to the stability analysis of spiking recurrent networks

Farzad Farkhooi<sup>1,2</sup>  
 Carl van Vreeswijk<sup>3</sup>

FARZAD.FARKHOOI@FU-BERLIN.DE  
 CORNELIS.VAN-VREESWIJK@PARISDESCARTES.FR

<sup>1</sup>Bernstein Center for Computational Neuroscience, Berlin

<sup>2</sup>National Institutes of Health

<sup>3</sup>Paris Descartes University and CNRS

In this study, we establish a formal link between the classical stability analysis of a recurrent network of spiking neurons with its exact canonical point process equivalent. In this approach the self-consistent solution of inter-event statistics leads to an eigenvalue equation that it determines the stability of asynchronous solution. This

framework allows a straightforward analysis of systems with the short-term synaptic depression (STD). Our result here provides a basic ground for the system phase transition from the asynchronous state to collective dynamics of ensemble synchronization, so-called population spike. We show the stability analysis using inter-event statistics in a self-consistent manner can be generic machinery that determines the conditions that must be satisfied to provide stable computation modes in networks with multiple state variables.

## I-74. Low-dimensional dynamics of somatosensory cortex: experiment and simulation

Cliff C Kerr<sup>1,2</sup>  
 Marius Pachitariu<sup>3</sup>  
 Jordan Iordanou<sup>1</sup>  
 Joseph Francis<sup>1</sup>  
 Maneesh Sahani<sup>3</sup>  
 William Lytton<sup>1</sup>

CLIFFK@NEUROSIM.DOWNSTATE.EDU  
 MARIUS@GATSBY.UCL.AC.UK  
 JORDAN.IORDANOU@DOWNSTATE.EDU  
 JOEY199US@GMAIL.COM  
 MANEESH@GATSBY.UCL.AC.UK  
 BILL@NEUROSIM.DOWNSTATE.EDU

<sup>1</sup>SUNY Downstate Medical Center

<sup>2</sup>University of Sydney

<sup>3</sup>Gatsby Computational Neuroscience Unit, UCL

High-dimensional datasets recorded via multielectrode arrays are often best understood in terms of a low-dimensional representation that highlights those dimensions that account for the greatest amount of variance in the data. In this work, we describe the low-dimensional dynamics of the neural activity in the somatosensory cortex of anesthetized rats that results from tactile stimuli. In addition, we created a simulation of six-layered somatosensory cortex and ventral posterolateral thalamus using a network of 10,000 spiking Izhikevich neurons, with connectivities drawn from empirical data. The low-dimensional representation of the spiking activity was determined via principal component analysis (PCA). In both experiment and simulation, PCA showed that onset and offset responses to the tactile stimuli corresponded to approximately semicircular trajectories in state space; the offset responses were rotated by roughly 90 degrees with respect to the onset responses. To complement the PCA analysis of the signal dynamics, we fit a recurrent linear model (RLM) to quantify the noise dynamics. Most of the noise variance was explained by the first dimension of the RLM, which captured 0.48 and 0.34 bits of information per spike in the experimental data and simulation respectively. Other dimensions of the RLM corresponded to traveling waves, which were observed in both the experiment and simulation with velocities of approximately 0.1 m/s. Notably, these waves were associated with slow dynamics in both the experiment and simulation. Thus, the realistically long time scales of noise correlations in the simulation arose due to intrinsically generated traveling waves of spiking activity, not from long time constants. While low-dimensional dynamics have previously been described in the ensemble activity of motor cortex, to our knowledge this is the first time they have been described in somatosensory cortex.

## I-75. On the brink of instability: network fragility in the epileptic cortex

Sridevi Sarma  
 Duluxan Sritharan  
 Johns Hopkins University

SREE@JHU.EDU  
 DSRITHA1@JHU.EDU

Epilepsy is a network phenomenon characterized by atypical activity at the neuronal and population levels during seizures, including tonic spiking, increased heterogeneity in spiking rates and synchronization. The etiology of epilepsy is unclear but a common theme among proposed mechanisms is that structural connectivity between neurons is altered. It is hypothesized that epilepsy arises not from random changes in connectivity, but from specific structural changes to the most fragile nodes or neurons in the network that cause unstable activity. Here,

the minimum energy perturbation on functional connectivity required to destabilize linear networks is derived. Perturbation results are then applied to a probabilistic nonlinear neural network model that operates at a stable fixed point. That is, if a small stimulus is applied to the network, the activation probabilities of each neuron respond transiently but eventually return to the fixed point. When the perturbed network is destabilized, the activation probabilities shift to larger or smaller values or oscillate when a small stimulus is applied. Finally, the structural modifications in the neural network that achieve the functional perturbation are derived. The findings suggest that the most fragile nodes in the network are excitatory neurons that become more active or inhibitory neurons that become less active. This is consistent with abnormal axonal sprouting of excitatory neurons and loss of inhibitory chandelier cells observed in epileptic cortical tissue. Furthermore, simulation of the unperturbed and perturbed network reect neuronal activity observed in epilepsy patients. The qualitative changes in network dynamics due to destabilizing perturbations, including the emergence of an unstable manifold or a stable limit cycle, may be indicative of neuronal or population dynamics during seizure. That is, the epileptic cortex is always on the brink of instability and that minute changes in synaptic weights can suddenly destabilize the network to cause seizures.

## I-76. Alpha coherence networks and the prerequisites for communication in simulated cortex

Stefan Berteau<sup>1,2</sup>

BERTEAU@BU.EDU

<sup>1</sup>Boston University

<sup>2</sup>Brandeis University

Observations of changes in coherence of oscillations as a function of task conditions and performance in behavioral tasks have led to the suggestion that oscillatory coherence can enhance functional connectivity and communication between regions, or in the case of alpha (8-14 Hz) oscillations, prevent such communication. Interpretations of such studies implicitly assume that oscillatory coherence is an appropriate measure of information transfer. Aimed at investigating this assumption, our modeling work analyzes mutual information in a series of rate-based cortical oscillator models with the potential to both match previous studies and expand upon them. We show that information transfer between brain regions receiving thalamically-driven alpha oscillations is dependent on many factors, not all of which can be easily measured in human subjects. Coherence can be beneficial to communication, but it requires that the two oscillators be maintained with a specific phase difference and not be saturated. The necessary phase difference depends on biophysical factors such as conduction delays and oscillation frequency. In addition, susceptibility to saturation depends on the inter-area connectivity pattern and both intra- and inter-area synaptic weights. Our results suggest that coherence networks are necessary but not sufficient for characterizing alpha-band communication in the brain. This work was supported in part by CELEST, a National Science Foundation Science of Learning Center (NSF OMA-0835976)

## I-77. Contextual influences on gamma oscillations in inhibition stabilized networks

Yashar Ahmadian

Kenneth Miller

Columbia University

YA2005@COLUMBIA.EDU

KEN@NEUROTHEORY.COLUMBIA.EDU

Stimulation enhances LFP power in the gamma frequency range (30-90 Hz) in many cortical areas. Attention typically increases gamma power but lowers it in V1 [1]; surround suppression (SS) raises gamma power in V1 but lowers its peak frequency [2]. As V1 gamma oscillations are not phase coherent [3], we model them as noise that is bandpass-filtered by the intrinsic dynamics of the local cortical network. For strong effective recurrent connectivity, this network is an inhibition-stabilized network (ISN). ISN's exhibit a paradoxical effect: local inhibitory populations lower their mean rates in response to increased mean excitatory external input. We

show that quite generally, an increase in the power of fluctuations in input to either the excitatory or the inhibitory populations raises their gamma power. During SS, both the mean (DC) and the fluctuations of the excitatory input received by a local ("center") inhibitory population from the surround region are raised. The increase in the DC excitatory input suppresses mean inhibitory (and excitatory) responses due to the paradoxical ISN effect, yet the increase in fluctuations raises gamma power. The attention-induced increase of response gains can arise from a change in excitatory feedback from higher areas: either an increase in excitatory input to excitatory populations or a decrease to inhibitory populations (excitatory or disinhibitory mechanisms, respectively). Assuming the mean input and its fluctuations increase or decrease together, a disinhibitory (excitatory) attention mechanism can yield the observed decrease in gamma power in V1 (the increased gamma observed in other areas). When neurons have power-law input/output functions, changes in the DC input lead to changes in the mean network responses, which change the network's filtering properties. We show that in the ISN regime, this effect accounts for the drop in the gamma peak frequency during SS.

## I-78. Dynamical criticality during Induction of anesthesia in human ECoG recordings

Leandro Alonso<sup>1</sup>

Alex Proekt<sup>1,2</sup>

Guillermo Cecchi<sup>3</sup>

Marcelo Magnasco<sup>1</sup>

LALONSO@ROCKEFELLER.EDU

PROEKT@GMAIL.COM

GCECCHI@GMAIL.COM

MGNNSCLB@MAIL.ROCKEFELLER.EDU

<sup>1</sup>Rockefeller University

<sup>2</sup>Cornell University

<sup>3</sup>T.J. Watson IBM Research Center.

In this work we analyze electrocorticography recordings (ECoG) in human subjects as they are anesthetized. We hypothesize that the decrease in responsiveness that defines anesthesia induction is concomitant with the stabilization of neuronal dynamics. To test this hypothesis, we performed a moving vector autoregressive analysis and quantified stability of neuronal dynamics using eigenmode decomposition of the autoregressive matrices, independently fitted to short sliding temporal windows. Consistent with the hypothesis we show that while the subject is awake, oscillations of neuronal activity are predominantly found at the edge of instability, but as the subject becomes anesthetized the dynamics becomes more damped. Analysis of distributions of eigenmodes in the awake and anesthetized brain revealed statistically significant stabilization not present in surrogate data. Thus, stability analysis may suggest a novel way of quantifying changes in neuronal activity that characterize loss of consciousness induced by general anesthetics.

## I-79. Modeling neural responses in the presence of unknown modulatory inputs

Neil C Rabinowitz<sup>1</sup>

Robbe L.T. Goris<sup>1</sup>

Johannes Balle<sup>1</sup>

Eero Simoncelli<sup>1,2</sup>

NEIL.RABINOWITZ@NYU.EDU

ROBBE.GORIS@NYU.EDU

JOHANNES.BALLE@NYU.EDU

EERO.SIMONCELLI@NYU.EDU

<sup>1</sup>New York University

<sup>2</sup>Howard Hughes Medical Institute

Neurons transmit information with spike trains that differ across repeated measurements. The origin of this variability is unknown, but it is common to describe spike count distributions as Poisson, despite the fact that their variance generally exceeds that expected of a Poisson process. This is likely because neurons' firing rates are also at the mercy of numerous uncontrolled and/or unobserved modulatory factors that alter their gain, including

the influence of recently emitted spikes, locally-generated gain control, top-down signals (e.g. attention, arousal, motivation), and physiological conditions (e.g. metabolic resource availability). Regardless of their origin, fluctuations in these signals can confound or bias the inferences that one derives from spiking responses. These effects can be captured by a modulated Poisson model, whose rate is the product of a stimulus-driven response function and an unknown modulatory signal (Goris, Movshon, Simoncelli, 2013). Here, we extend this model, by including modulatory elements that are known (specifically, spike-history dependence, as in previous GLM models, Pillow et al, 2008), and by constraining the remaining latent modulatory signals to be smooth in time. We fit the entire model, including hyperparameters, via evidence optimization (Park & Pillow, 2011), to the responses of ferret auditory midbrain and cortical neurons to complex sounds. Integrating out the latent modulators yields more readily-interpretable receptive field estimates than a standard Poisson model. Conversely, integrating out the stimulus dependence yields estimates of the slowly-varying latent modulators. For example, when applied to array recordings of macaque V1, we find complex spatial patterns of correlation amongst the latent modulators, including clusters of co-modulated units. In sum, use of the modulated Poisson model improves inference, and enables the study of signals underlying non-stationarities in neural responses.

## I-80. The temporal structure of a random network near criticality and human ECoG dynamics

Rishidev Chaudhuri<sup>1</sup>  
 Biyu He<sup>2</sup>  
 Xiao-Jing Wang<sup>1</sup>

RC1520@NYU.EDU  
 BIYU.HE@NIH.GOV  
 XJWANG@NYU.EDU

<sup>1</sup>New York University

<sup>2</sup>National Institutes of Health / NINDS

A widespread but controversial hypothesis is that neuronal networks operate near criticality, poised between a regime where activity fades rapidly and one where activity is amplified. Such systems show long-range spatiotemporal correlations and activity variables can show power-law scaling. Recently we found that power spectra of human ECoG recordings display long-time correlations and a distinctive three-part shape: at low and high frequencies spectra scale as power laws (approximately  $1/f^2$ ), while at intermediate frequencies they are flat. Here we show that both long autocorrelations and tripartite power spectra can emerge from randomly-connected linear networks with excitation and inhibition balanced to place networks near instability. In such networks, long-range interactions create a slow timescale shared across nodes, producing low-frequency power-law scaling. The remaining timescales cluster around the intrinsic timescale of the nodes, producing high-frequency power-law scaling and a characteristic 'knee' in the spectrum. Since the long timescale is more widely shared than fast timescales, correlations in low-frequency activity are stronger and more distributed than those in high-frequency activity, as often observed in data. Moreover, this network mode is preferentially driven by spatially-correlated input and naturally converts spatial correlations in input to long temporal correlations in output. Consequently, decorrelation of network input can explain widespread observations of decrease in low-frequency power upon task initiation. Allowing connection strength to decay with inter-node distance produces timescales intermediate between the slow and fast modes. This decreases the low-frequency scaling exponent, moving it closer to  $1/f$  as connectivity sharpens. Thus variations in effective coupling length might underlie differences in the low-frequency scaling exponent observed between experiments and states. In conclusion, we demonstrate a minimal linear model that explains multiple features of ECoG power spectra and provides insight into factors that can change these spectra upon task initiation, across brain regions and between states like waking and sleep.

## I-81. How brain areas can predict behavior, yet have no influence

Xaq Pitkow<sup>1,2</sup>  
Kaushik Lakshminarasimhan<sup>1</sup>  
Alexandre Pouget<sup>3</sup>

XAQ@CNS.BCM.EDU  
JLAKS@CNS.BCM.EDU  
ALEX@CVS.ROCHESTER.EDU

<sup>1</sup>Baylor College of Medicine

<sup>2</sup>Rice University

<sup>3</sup>University of Geneva

Only one type of correlations ultimately limit the information content of large neural populations, and these are noise correlations that look exactly like the signal. We call this ‘bad noise’. For small noise, these have the structure of ‘differential correlations’. The presence of bad noise in a large neural population generates a redundant code from which the information can be extracted in many different ways with negligible loss. In particular, distinct populations can share bad noise, in which case the brain could read out any or all of them with almost no change in information. One consequence is that every neuron would then have a correlation with behavior (choice correlation  $C_k$ ) equal to the fraction of total information it carries, according to  $C_k = d_k / d'$  where  $d_k$  is the discriminability ( $d$ -prime) for neuron  $k$  and  $d'$  is the discriminability for the whole population. Here we analyze the information content and behavioral correlations of multiple populations that have overlapping sources of information. We derive new results that place constraints on the relevant correlations between areas, and we explain paradoxical experimental observations. In particular, during vestibular discrimination tasks, Ventral IntraParietal cortex (VIP) shows choice correlations that are double those predicted from optimal decoding, yet there is no behavioral change when VIP is inactivated. We quantitatively explain how this is possible. Overall, this framework, and the noise covariance matrix  $E$ , provides a simple way to think about multiple populations and their interactions, but with a massively reduced dimensionality that nonetheless preserves all relevant information.

## I-82. Opening the black box: understanding how recurrent neural networks compute

David Sussillo<sup>1</sup>  
Omri Barak<sup>2</sup>

SUSSILLO@STANFORD.EDU  
OMRI.BARAK@GMAIL.COM

<sup>1</sup>Stanford University

<sup>2</sup>Technion

Recurrent neural networks (RNNs) are important tools in computational and theoretical neuroscience. There has been much recent work on training RNNs to do interesting tasks, for example tasks analogous to those used in experimental settings (e.g. random motion dots task). However, very little is understood about the mechanisms underlying computations in RNNs. On the one hand, they are neural networks. On the other hand, they are nonlinear dynamical systems. In order to understand how RNNs work, we employed fixed point and linearization analyses commonly used in nonlinear dynamical systems theory. For all the examples shown (e.g. 3-bit memory, input-dependent oscillator), the fixed-point and linearization analyses made sense of the RNN mechanism. Further, we found that linearization was appropriate for a larger set of points than just fixed points, including points in state-space exhibiting slow dynamics. Our technique for finding fixed points in the RNN is to define a simple auxiliary function related to the network equations and feed this function to an unconstrained non-convex optimization routine. In summary, we found that the techniques presented here[1] could be used to study and reverse engineer — i.e. to open the black box of — RNNs. The results of such analyses are framed in the language of dynamical systems, and have proven insightful to understanding cortical function, for example in the prefrontal[2] and motor cortices[3]. 1. Sussillo, D.\* & Barak, O\*. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Computation* 25, 626–649 (2013). 2. Mante, V.\*, Sussillo, D.\*, Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84 (2013). 3. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A recurrent neural network that produces EMG from rhythmic dynamics. *Cosyne* 2013, III-67

## I-83. Dynamics of networks of excitatory and inhibitory units with sparse, partially symmetric couplings

Daniel Marti<sup>1,2</sup>  
Nicolas Brunel<sup>3</sup>  
Srdjan Ostojic<sup>4</sup>

DANI.MARTI.ORTEGA@GMAIL.COM  
NBRUNEL@GALTON.UCHICAGO.EDU  
SRDJAN.OSTOJIC@ENS.FR

<sup>1</sup>Group of Neural Theory, ENS

<sup>2</sup>INSERM

<sup>3</sup>University of Chicago

<sup>4</sup>Ecole Normale Supérieure

Over the past ten years, several studies based on multielectrode recordings have revealed that the connectivity patterns among cortical pyramidal neurons have a non-trivial statistical structure. A prominent feature of this structure is the fact that bidirectional connections between neurons appear more often than one would expect if connections were random. How this prevalence of bidirectional connections affects the dynamics of cortical networks, and what possible functional role it plays, is still an open question. In a recent study, Litwin-Kumar and Doiron investigated the effect of clustering induced by reciprocal connections on the network dynamics. Here we study the dynamics of large-scale networks of excitatory and inhibitory units with unclustered, random, sparse, and partially symmetric connections. We investigate networks of both rate units and spiking neurons and show that, in both cases, the nature of the dynamics of the network is strongly determined by the spectrum of eigenvalues of the connectivity matrix. For weak couplings, all the eigenvalues are small and the network is an equilibrium state characterized by constant firing rates. For strong couplings, eigenvalues become large and the network displays an instability that leads to a rate-chaotic regime with heterogeneous and fluctuating firing rates. We show that, in the rate-chaotic regime, adding partial symmetry to the couplings increases the characteristic timescale of the activity fluctuations. This increase is related to the flattening of the eigenspectra of the synaptic matrix that occurs when symmetry increases.

## I-84. Volume transmission as a new homeostatic mechanism

Yann Sweeney<sup>1,2</sup>  
Jeanette Hellgren Kotaleski<sup>2</sup>  
Matthias H Hennig<sup>1</sup>

YANN.SWEENEY@ED.AC.UK  
JEANETTE@NADA.KTH.SE  
M.HENNIG@ED.AC.UK

<sup>1</sup>University of Edinburgh

<sup>2</sup>KTH

Gaseous neurotransmitters such as nitric oxide (NO) provide a unique and often overlooked mechanism for neurons to communicate through diffusion within a volume, regardless of synaptic connectivity. NO is known to provide homeostatic control of intrinsic excitability through regulation of potassium ion channels (Steinert et al., 2011). We conduct a theoretical investigation of the distinguishing roles of nitric oxide volume transmission in comparison with canonical non-diffusive homeostatic mechanisms. Sparsely connected recurrent networks of integrate-and-fire model neurons in a balanced regime are simulated, with modulation of  $K^+$  conductance following activity-dependent NO synthesis implemented as modulation of the firing threshold. These simulations show that diffusive homeostasis provides a robust mechanism for maintaining stable activity in cortical networks across a wide range of perturbations in external input and connectivity. Notably, firing rates of subgroups of neurons receiving distinct input rates remain separated after a global target activity has been achieved, a feature which is not exhibited in networks undergoing non-diffusive homeostasis. Additionally, diffusive homeostasis maintains firing rate heterogeneity and thus provides broad rate distributions without the need for specialised connectivity profiles or neural transfer functions, in agreement with recent experimental evidence (Roxin et al., 2013). This results in networks capable of linearly responding to changes in input, unlike those undergoing non-diffusive modulation which tend to saturate more easily. We show that the differences between diffusive and non-diffusive homeostasis observed in network simulations can be captured by a dynamic mean field analysis with analogous



homeostatic mechanisms.

## I-85. Network dynamics of spiking neurons with adaptation

Moritz Deger<sup>1</sup>  
Tilo Schwalger<sup>1</sup>  
Richard Naud<sup>2</sup>  
Wulfram Gerstner<sup>1</sup>

MORITZ.DEGER@EPFL.CH  
SCHWALGER@GMX.DE  
RNAUD@UOTTAWA.CA  
WULFRAM.GERSTNER@EPFL.CH

<sup>1</sup>Ecole Polytechnique Federale de Lausanne

<sup>2</sup>University of Ottawa

Local neuronal circuits in the neocortex consist of hundreds of neurons. These neurons typically show refractoriness after emitting an action potential (spike), and accumulating refractory effects result in adaptation on multiple time scales. Both finite network size and neuronal adaptation make it difficult to derive population dynamics using traditional, stochastic process approaches. Here we present a novel theory of the population activity of finite-sized, randomly connected networks of spiking neurons with adaptation, which is obtained by approximating neuronal spike emission by a quasi-renewal process. Our theory describes the population activity and its spectral density in coupled networks, exemplified for the typical case of a network of excitatory and inhibitory neurons. Furthermore, we show how correlated noise influences the stationary population activity, and how it can contribute to synchronize or desynchronize neuronal circuits.

## I-86. Connection-type specific biases make random network models consistent with cortical recordings.

Tim Vogels<sup>1,2</sup>  
Christian Tamm<sup>2</sup>  
Michael Avermann<sup>2</sup>  
Carl Petersen<sup>2</sup>  
Wulfram Gerstner<sup>2</sup>

TIM.VOGELS@CNCB.OX.AC.UK  
CTOMM@GMX.DE  
MICHAEL.AVERMANN@GMAIL.COM  
CARL.PETERSEN@EPFL.CH  
WULFRAM.GERSTNER@EPFL.CH

<sup>1</sup>University of Oxford

<sup>2</sup>Ecole Polytechnique Federale de Lausanne

Random sparse network architectures are ubiquitous in computational neuroscience but the underlying assumption that they are a good representation of real neuronal networks has been met with skepticism. Here we used two experimental data sets - a study of triplet connectivity statistics, and a data set measuring neuronal responses to channelrhodopsin stimuli - to evaluate the fidelity of thousands of model networks. Network architectures comprised three neuron types (excitatory, fast spiking, and non-fast spiking inhibitory) and were created from a set of rules that govern the statistics of the 9 resulting connection types. In a high-dimensional parameter scan, we changed the degree distributions (i.e. how many cells each neuron connects with) and the synaptic weight correlations of synapses from or onto the same neuron. These variations converted initially random and homogeneously connected networks, in which every neuron sent and received equal numbers of synapses with equal synaptic strength distributions, to highly heterogeneous networks in which the number of synapses per neuron, as well as average synaptic strength of synapses from or to a neuron were variable. By evaluating the impact of each variable on the networks' structure and dynamics, and their similarity to the experimental data, we could falsify the random sparse connectivity hypothesis for 7 of 36 connectivity parameters, but we also confirmed the hypothesis in 8 cases. 11 parameters had no impact on the results of the test protocols we used, and 10 parameters could not be tested due a lack of appropriate data.

## I-87. Fast sampling in recurrent neural circuits

Mate Lengyel<sup>1</sup>  
 Guillaume Hennequin<sup>1</sup>  
 Laurence Aitchison<sup>2</sup>

M.LENGYEL@ENG.CAM.AC.UK  
 G.JE.HENNEQUIN@GMAIL.COM  
 LAURENCE@GATSBY.UCL.AC.UK

<sup>1</sup>University of Cambridge

<sup>2</sup>Gatsby Computational Neuroscience Unit, UCL

Time is at a premium for recurrent network dynamics, and particularly so when they are stochastic and correlated: the quality of decoding and inference from such dynamics fundamentally depends on how fast the neural circuit generates new samples from its stationary distribution. Indeed, behavioral decisions can occur on fast time scales (~100 msec), but it is unclear what neural circuit dynamics afford sampling at such high rates. We analyzed a stochastic form of rate-based linear neuronal network dynamics with synaptic weight matrix  $W$ , and the dependence of the covariance of the stationary distribution of joint firing rates,  $\Sigma$ , on  $W$ . For decoding,  $\Sigma$  is an inconvenient byproduct of recurrent processing, for sampling-based inference it can be actively used to represent posterior uncertainty under a linear-Gaussian latent variable model — in either case, this variability ultimately needs to be averaged out. The key insight is that the mapping between  $W$  and  $\Sigma$  is degenerate: there are infinitely many  $W$ 's that lead to sampling from the same  $\Sigma$  but differ greatly in the speed at which they sample. We were able to explicitly separate these extra degrees of freedom in a parametric form and thus study their effects on sampling speed. We show that previous proposals for probabilistic sampling in neural circuits correspond to using a symmetric  $W$  which violates Dale's law and results in critically slow sampling, even for moderate stationary correlations. In contrast, optimizing network dynamics for speed consistently yielded asymmetric  $W$ 's and dynamics characterized by fast transients, such that samples of network activity became fully decorrelated over ~10 msec. Importantly, networks with separate excitatory/inhibitory populations proved to be particularly efficient samplers, and were in the balanced regime. Thus, plausible neural circuit dynamics can perform fast sampling for efficient decoding and inference.

## I-88. Relating synaptic efficacies and neuronal separation in a simulated layer 2/3 cortical network

Daniel Miner  
 Jochen Triesch

MINER@FIAS.UNI-FRANKFURT.DE  
 TRIESCH@FIAS.UNI-FRANKFURT.DE

Frankfurt Institute for Advanced Studies

It is well known that local cortical networks exhibit a degree of nonrandom structure and clustering (Perin et al. 2011, Song et al. 2005). Experimental data (Feldmeyer et al. 2006, Perin et al. 2011) suggests that the probability of a lateral synapse existing between two neurons is weakly but significantly dependent on neuronal separation over the scale of a cortical column. We developed a spiking neural network model using parameters drawn from surveys of layer 2/3 of the rodent cortex (Avermann et al. 2012, Feldmeyer et al. 2006, Petersen et al. 2013, Thomson et al. 2007). The excitatory population of the network features multiple plastic and homeostatic mechanisms including short term plasticity, spike timing dependent plasticity, intrinsic homeostatic plasticity, synaptic normalization (Lazar et al. 2009), and structural plasticity. Allowed to evolve spontaneously, the network develops a heavy-tailed distribution of synaptic efficacies and a small population of strong connections which vary relatively little over time, as observed experimentally (Yasumatsu et al. 2008), and as has been previously demonstrated in a simpler, non-topological discrete binary version of a similar network (Zheng et al. 2013). We embed this network in a two dimensional topological sheet to represent a single cortical layer and impose a weak distance-dependent connection probability on the structural growth of the network. The natural spontaneous evolution of the network under this condition leads to the development of increased clustering (compared to a fully random graph with the same sparsity) and a weak but significant relationship between synaptic efficacies and neuronal separations. We study and quantify this relationship and how it varies over a reasonable biological parameter space, and find generally similar results in experimental data (Feldmeyer et al. 2006).

## I-89. Encoding and learning the timing of events in neural networks

Alan Veliz-Cuba<sup>1,2</sup>  
Zachary Kilpatrick<sup>1</sup>  
Kresimir Josic<sup>1</sup>

ALANAVC@MATH.UH.EDU  
ZPKILPAT@MATH.UH.EDU  
JOSIC@MATH.UH.EDU

<sup>1</sup>University of Houston

<sup>2</sup>Rice University

Understanding how the brain learns sequences of events is a fundamental problem in neuroscience. In a learned sequence, the timing is as important as learning the order in which events occur. Recordings in awake monkeys and rats with repeated presentations of a cued sequence of stimuli give some insight into the process. After training, the presentation of the cue by itself can engender temporal patterns of neural activity correlated with those evoked by the stimulus sequence. While many mechanisms of learning the correct ordering have been proposed, how timing is learned remains largely unexplored. Given that the activity of single cells evolves on the timescale of tens of milliseconds, it is likely that such timing is represented in the activity of populations of neurons. However, how the requisite connectivity patterns arise is not understood. We discuss a model neuronal network capable of encoding and learning the timing of a sequence of events. This type of learning relies on two simple but physiologically supported mechanisms: short term facilitation and BCM type synaptic plasticity. We show that short term facilitation allows the neuronal network to encode timing anywhere from milliseconds to seconds. BCM type synaptic plasticity allows the network to learn such timing quickly and accurately. Our model gives predictions that can be experimentally tested. It provides information about the time it takes to learn and how this learning depends on the number of stimulus trials. We show that it can also be used to investigate more complex questions – for instance, how "competing" stimulus sequences affect network structure and activity.

## I-90. Efficient fitting of large-scale neural models

Jascha Sohl-Dickstein<sup>1,2</sup>  
Niru Maheswaranathan<sup>1</sup>  
Benjamin Poole<sup>1</sup>  
Surya Ganguli<sup>1</sup>

JASCHA@STANFORD.EDU  
NIRUM@STANFORD.EDU  
POOLE@CS.STANFORD.EDU  
SGANGULI@STANFORD.EDU

<sup>1</sup>Stanford University

<sup>2</sup>Khan Academy

Efficient fitting of large-scale models is becoming increasingly crucial in systems neuroscience. As the rate of data acquisition grows, neuroscientists are tasked with modeling large-scale recordings of many neurons across long time scales, in response to complex stimuli and behavioral paradigms. The BRAIN initiative promises to drive the production of these datasets orders of magnitude larger. Additionally, as the questions neuroscientists wish to ask become more sophisticated, they also become more computationally demanding. For example, closed-loop or online experiments require models to be fit to neural activity in real time, and could allow breakthroughs in brain-machine interfaces and receptive field characterization if this optimization challenge can be solved. Fitting models to these datasets often boils down to minimizing an objective function summed over the data. Most current approaches fall into two categories: Stochastic gradient descent algorithms process the data in minibatches, and each iteration is computationally cheap. Quasi-Newton methods estimate the curvature of the objective and require fewer iterations, yet each iteration is more expensive to compute. We present an algorithm for optimizing high-dimensional functions over large datasets that combines the computational efficiency of stochastic gradient descent with the second order curvature information accessible by quasi-Newton methods. We unify these disparate approaches by maintaining an independent quadratic approximation for each minibatch, and maintain computational tractability even for high-dimensional problems by storing and manipulating these approximations in a shared, time-evolving low-dimensional subspace. To demonstrate the effectiveness of our approach, we apply this procedure to the problem of fitting generalized linear models over many neurons. We test our algorithm on simulated data consisting of a population of 100 neurons, driven by external noise stimuli, spike-train history, and

connections to other neurons. We release open source Python code for both the optimizer and the GLM.

## I-91. Effects of noise injection in artificial neural networks

Benjamin Poole<sup>1</sup>  
Jascha Sohl-Dickstein<sup>1,2</sup>  
Surya Ganguli<sup>1</sup>

POOLE@CS.STANFORD.EDU  
JASCHA@STANFORD.EDU  
SGANGULI@STANFORD.EDU

<sup>1</sup>Stanford University  
<sup>2</sup>Khan Academy

Intrinsic variability, or noise, is an extremely salient feature of biophysical systems. Many studies have shown that neural systems have evolved to efficiently process information in the face of this noise. Yet the question still persists as to whether this noise could be taken advantage of to solve hard computational problems. Our work builds on a growing body of evidence in machine learning that show the benefits of noise injection on learning representations in neural networks. We first consider the case of injecting noise into the hidden layer of an autoencoder, a simple one hidden layer neural network used for representation learning. By analytically marginalizing out the noise, we derive a set of penalties that can be related to existing explicit regularization strategies for neural networks and autoencoders. In particular, we find that injecting Poisson-like noise into the hidden layer yields a sparsity-inducing penalty. Thus Poisson variability in neural systems may provide a computationally simple mechanism for learning sparse representations. We empirically evaluate the impact and utility of noise in learning by training neural networks for image denoising and digit classification. In both cases, we find that noise injection yields hidden representations that are sparser, more decorrelated, and more robust to random removal of neurons than their noiseless counterparts. Furthermore, the networks trained with noise injection yield improved performance on denoising natural image patches, and achieve state-of-the-art performance on a specific digit recognition task. These results indicate that noise can be beneficial in learning sparse distributed representations of sensory inputs.

## I-92. Unsupervised emergence of continuous tuning curves and sensory maps in spiking neural networks

Robert Guetig  
Max Planck Institute of Experimental Medicine

GUETIG@EM.MPG.DE

Responses of many neurons in mammalian sensory pathways, e.g. simple cells in V1, are selective to specific stimulus features, e.g. edges or bars, and show continuous tuning with respect their target features' characteristic dimensions, e.g. orientation. In many species, sensory neurons form continuous features maps, such that the parameters of their tuning vary continuously across the cortical surface. While the shape and organization of neural tuning is a crucial determinant of cortical sensory representations, our understanding of their formation in spiking neural networks has remained limited. Here we introduce a novel type of self-supervised neural network in which continuous neural tuning and feature maps emerge in an unsupervised fashion. The network consists of an input layer, a plastic processing layer, and a supervisor layer. While the all-to-all connectivity from the input to the processing layer is purely feed-forward, the supervisory layer neurons implement a recurrent teaching signal. Specifically, after each learning episode, each neuron in the processing layer receives a teaching signal from one dedicated supervisor neuron, that signals the desired level of activity and drives multi-spike tempotron learning. Sensory episodes are mimicked by long trials of random Poisson spiking in the input layer. Within this activity, sensory features are embedded as short recurring spike patterns. Driven by a positive feedback instability, the processing neurons efficiently learn to detect the features. Importantly, the organization of the learned feed-forward neural responses is fully determined by the synaptic matrix of the supervisory layer and can realize arbitrary spatial response profiles across the processing layer, without direct lateral interactions. Moreover, when

sensory features span a continuous feature dimension, continuous feature maps emerge. Our results introduce a novel unsupervised network mechanism for spiking neurons that subserves the emergence of continuous neural maps in biologically plausible spiking network models.

## I-93. Learning multi-stability in plastic neural networks

Friedemann Zenke<sup>1</sup>  
Everton J Agnes<sup>2</sup>  
Wulfram Gerstner<sup>1</sup>

FRIEDEMANN.ZENKE@EPFL.CH  
EVERTON.AGNES@UFRGS.BR  
WULFRAM.GERSTNER@EPFL.CH

<sup>1</sup>Ecole Polytechnique Federale de Lausanne

<sup>2</sup>Universidade Federal do Rio Grande do Sul

The connectivity of cortical networks is believed to be structured such that it can store and retrieve memories or perform other vital computations such as decision making, processing of sensory information or inference. Many of these tasks have been demonstrated to be solvable in balanced network models by systematic and mostly manual re-structuring of synaptic connectivity. However, in biological networks the underlying functional structures have to self-organize under conditions where network dynamics and plasticity interact. This is thought to be made possible by a sophisticated interplay of plasticity and homeostatic mechanisms, of which the exact details remain unknown. Here we approach this question in a model of a large balanced network, in which excitatory synapses are plastic and obey a plausible triplet STDP rule. In existing work on associative working memory, recurrent connectivity patterns are written into the synaptic weight structure to form cell assemblies. Here we use correlated external input to trigger the formation of such assemblies. We show that assemblies quickly become unstable when connections are strong enough to maintain persistent activity at elevated firing rates. Building on our previous work we analyze this instability of the network model in a low dimensional dynamical mean field formulation. This analysis shows that the homeostatic triplet STDP learning rule, like many other similar learning rules, has only a single fixed point at the designated target firing rate of homeostasis. Following up on this we then show, how we can construct a plasticity rule that supports multiple stable rates. The so emerging rules contain terms that can be interpreted as heterosynaptic plasticity. Finally we test this idea in network simulations of spiking neurons and explain why inhibitory plasticity can help to maintain network stability, but is not sufficient to do it alone.

## I-94. Interplay between short- and long-term plasticity in cell-assembly formation

Naoki Hiratani  
Tomoki Fukai

HIRATANI@BRAIN.RIKEN.JP  
TFUKAI@RIKEN.JP

RIKEN Brain Science Institute

Some episodes remain in the brain as memory, while others do not. Memory traces are considered to form cell assemblies with the help of synaptic plasticity. However, it is unclear how the various forms of long-term and short-term synaptic plasticity cooperatively create such assemblies. Recent experimental results suggest that, at the circuit level, long-term plasticity is influenced by short-term plasticity. In addition, computational studies have shown that short-term plasticity, particularly short-term depression (STD), enriches neural population dynamics in a timescale relevant to spike-timing-dependent plasticity (STDP). Here, we investigate how the three known forms of synaptic plasticity in cortical circuits (STDP, STD and homeostatic plasticity) contribute cooperatively to the generation, retention and modulation of cell assemblies in a recurrent neuronal network model. We show that multiple cell assemblies generated by external stimuli can survive noisy spontaneous network activity in a particular range of the strength of STD. Furthermore, we show that the selective retention and spatial integration of cell assemblies can be spontaneously performed by the network. These results may have implications for the

understanding of the cortical mechanism of memory deficits.

## **I-95. A Hopfield net trained on images matches retinal spike statistics and encodes images efficiently**

Christopher Hillar<sup>1,2</sup>  
Sarah Marzen<sup>2</sup>  
Urs Koester<sup>1</sup>  
Kilian Koepsell<sup>2</sup>

CHILLAR@MSRI.ORG  
SMARZEN@BERKELEY.EDU  
URS@BERKELEY.EDU  
KILIAN@BERKELEY.EDU

<sup>1</sup>Redwood Center

<sup>2</sup>University of California, Berkeley

We fit Hopfield networks to natural image patches using two retina-inspired schemes: in the first, a network is trained on patches from binarized natural images convolved with a center-surround filter; in the second, two Hopfield neurons are used to code a trinarized pixel intensity. The resulting Hopfield networks find features (memories) that can be used for an image compression scheme which outperforms JPEG in the low-quality regime. In the second scheme, the Hopfield neurons have receptive fields similar to those of ON/OFF retinal ganglion cells, and sparse synaptic weights with coupling similar to the correlation seen for ON/OFF retinal ganglion cells. These results suggest that the Hopfield network trained on natural image patches could be a new functional model of the retina.

## **I-96. One rule to grow them all? Activity-dependent regulation for robust neural circuit specification**

Timothy O'Leary  
Eve Marder  
Brandeis University

TOLEARY@BRANDEIS.EDU  
MARDER@BRANDEIS.EDU

The nervous system comprises many distinct circuits that perform canonical computations such as sensory integration, memory retrieval and motor output. A fundamental question is how these circuits self-organize and maintain their functional properties such that they are robust to environmental perturbations and variability in the circuit components. Central pattern generating circuits are ubiquitous examples of self-organizing circuits that are found in all nervous systems and often have a known canonical function. In this work we use theory and computational modelling to show how a self-regulating and self organizing central pattern generating circuit emerges from a minimal set of activity-dependent plasticity rules that are derived from known biology. Starting with generic assumptions about how individual neurons regulate their intrinsic and synaptic properties, we show how the simplest biologically plausible model derived from these assumptions is equivalent to the engineering principle of integral control. The model is capable of generating diverse cell types in a way that is robust to initial conditions and ongoing activity perturbations, and can explain highly-constrained linear correlations in conductance expression that is observed in-vivo in identified cell types. Furthermore, the model can be used to control synaptic as well as intrinsic conductances to make a self-assembling central pattern generator network. Thus network-level homeostasis is seen to emerge from 'cell autonomous' regulation rules and variability on the single-neuron level is shown to be compatible with reliable, stereotyped circuit function.

## I-97. The dynamics of variability in nonlinear recurrent circuits

Guillaume Hennequin  
Mate Lengyel  
University of Cambridge

GJE.HENNEQUIN@GMAIL.COM  
M.LENGYEL@ENG.CAM.AC.UK

The joint variability in an ensemble of neurons contains important signatures of the circuit connectivity, the ways in which synaptic plasticity is likely to affect it, and - ultimately - the computations the circuit is performing. Exploiting these links quantitatively requires a formal understanding of how correlations arise mechanistically in recurrently connected networks of neurons. Existing theories provide a somewhat blurry picture, as they often connect some high-level statistics of the circuit connectivity to other summary statistics of the pairwise correlation distribution. More elaborate theories can predict pairwise correlations for any specific pair of neurons, but assume either small firing rate variability or weak correlations. Here, we develop a novel theoretical framework to obtain the full correlational structure of a stochastic network of nonlinear neurons described by rate variables. Under the assumption that pairs of membrane potentials are jointly Gaussian - which they tend to be in large networks - we obtain deterministic equations for the temporal evolution of the mean firing rates and the noise covariance matrix that can be solved straightforwardly given the network connectivity. Importantly, our theory requires neither the fluctuations nor the pairwise noise correlations to be weak, and works for several biologically motivated, convex single-neuron gain functions. We apply our formalism to visual area MT, for which data on how particular stimuli affect neuronal variability has been published recently (Ponce-Alvarez et al., 2013). We find that a balanced ring model network with a threshold-quadratic nonlinearity captures the stimulus-dependence of both the Fano factor and the noise correlations. Interestingly, this network operates in a regime very different from previous proposals for variability quenching which all relied on multistable attractor dynamics, but closer to the behavior of the "stabilized supralinear network" of Ahmadian et al. (2013) of which we extend the analysis to the stochastic regime.

## I-98. Specialization within populations of neural feature detectors

Julia Hillmann  
Robert Guetig  
Max Planck Institute of Experimental Medicine

HILLMANN@EM.MPG.DE  
GUETIG@EM.MPG.DE

Many neurons in sensory pathways respond selectively to a narrow class of stimuli such as faces or specific communication calls. However, neural processing is also robust to a large degree of natural variability within such complex stimulus classes. For instance, neural circuits that underly speech processing must tolerate substantial spectral differences between female and male vocalizations. It is commonly believed that such difficult perceptual invariances are realized by populations of "expert" neurons that specialize on different substructures within a sensory object category. However, it is unclear what neural mechanisms of learning can subservise such symmetry breaking within populations of sensory neurons. The emergence and utilization of specialized learners to improve classification performance has also challenged Machine Learning for many decades. One important state of the art ensemble method for neural populations is the Mixture-of-Experts model that induces specialization of expert units via gating neurons that weight the responses of individual experts. However, due to its reliance on complex neural interactions and internal state variables, biologically plausible implementations of this model have remained challenging. Here we show that symmetry breaking and specialization within populations of spiking neurons can emerge through a simple instantiation of a competitive population learning rule. To mimic a sensory classification task we simulate a population of sensory neurons and require the number of responding neurons to discriminate between target and null stimuli. Following error trials, our learning rule uses relative spike-timing within the population and locally available values of depolarization to select only a subset of neurons to undergo learning while the others remain unchanged. We implement this algorithm within populations of integrate-and-fire neurons and demonstrate its high performance in spike pattern classification and speech recognition problems. Weak spike-timing based neural competition during learning is sufficient for neural populations to uncover "hidden"

structure within their input activity.

## **I-99. Role of long-range horizontal connections in visual texture classification**

Felix Bauer<sup>1,2</sup>

Matthias Kaschube<sup>1</sup>

FBAUER@FIAS.UNI-FRANKFURT.DE

KASCHUBE@FIAS.UNI-FRANKFURT.DE

<sup>1</sup>Frankfurt Institute for Advanced Studies

<sup>2</sup>Goethe University Frankfurt

In primary visual cortex, the system of long-range horizontal connections links neurons with non-overlapping receptive fields but similar preferred orientation [1]. While long-range connections have been suggested to play a crucial role in visual perception and during visual cortical development [2], their contribution to 'real world' recognition tasks has only been poorly characterized so far. Here, we pursue a computational study to examine the role of long-range horizontal connections in recognizing and distinguishing among visual textures. Our approach is based on biologically inspired object-classification systems [3]. The computations in the model involve the following key steps: (i) computing the responses of a set of orientation selective neurons to an image; the receptive fields vary in frequency and preferred orientation and their layout in the visual field is constrained by the layout of neurons in an orientation map (neglecting visual field scatter) (ii) pooling the responses across cortical space, using a pooling window of various size. While it is known that long-range connections are excitatory [1], little is known about their target cells. We therefore studied three alternative pooling models: purely excitatory (integrative), subtractive and divisive pooling. The output of our model was then fed into a classifier (linear SVM) that learns to predict a class label. As input images we used textures from the database of [4]. We find that long-range horizontal connections consistently facilitate classification if pooling is excitatory. Subtractive or divisive pooling has a much weaker effect. Moreover, performance continuously increases with the size of the pooling window and saturates at a value, which is close to the range of horizontal connections found by anatomical studies in tree shrew visual cortex [1]. Supported by BFNT 01GQ0840.

## **I-100. The brain as social network: self-organization of a dynamic connectome**

Michael Buice

Nicholas Cain

Ramakrishnan Iyer

Mike Hawrylycz

Stefan Mihalas

MICHAELBU@ALLENINSTITUTE.ORG

NICHOLASC@ALLENINSTITUTE.ORG

RAMAKRISHNANIYE@GMAIL.COM

MIKEH@ALLENINSTITUTE.ORG

STEFANM@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Neuronal systems are vastly more complex than their underlying genetic description. Internal learning rules and environmental stimuli result in a process of self-organization which generates networks responsible for an array of sophisticated computations. Necessarily, the computation performed in any cortical region must be connected deeply to the network architecture. Moreover, the architecture evolves from the interaction of neural activity and the dynamics of plasticity in cortex. From multi-patch measurements at the equilibrium state of this dynamics, there is a demonstrated relation between the probability of direct connections between neurons and the number of shared indirect connections (Perin, et al, PNAS 2011). We consider the outstanding question of connecting plasticity rules in neural networks to the evolving network structure. We analyze the dynamics of a stochastic model of network connectivity in which the state evolution is controlled by learning rules inspired by spike timing dependent plasticity. We derive a mean field representation of this dynamical model. The model admits a state in which the evolved graph matches the statistics of an Erdos-Renyi graph with a homogeneous probability of connection between neurons, implying that the network cannot store information without an explicit



external input. In the biologically realistic regime, there is a phase transition. The homogeneous state becomes unstable, breaking the exchange symmetry between neurons, allowing the network to store information without external input. In this symmetry broken regime, we demonstrate that the network model qualitatively predicts connectivity measures from experimental data. Interestingly, this prediction covers not only neural networks such as the regional connectivity of the mouse brain, but the *C. Elegans* neural network as well as social networks such as Twitter. This work is based on the Allen Mouse Brain Connectivity Atlas (data available via the Allen Brain Atlas portal at <http://www.brain-map.org>).

## II-1. Examining signal and noise correlation in a recurrent network of neurons in culture

Jeremie Barral<sup>1</sup>  
Tatjana Tchumatchenko<sup>2</sup>  
Alexander Reyes<sup>1</sup>

JEREMIE.BARRAL@GMAIL.COM  
TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE  
REYES@CNS.NYU.EDU

<sup>1</sup>New York University

<sup>2</sup>Max Planck Institute for Brain Research

Nearby neurons share common inputs. Therefore their spiking activity is often correlated. Correlated activity can eliminate random fluctuations by averaging out noise when it is not shared between neurons. Yet, strong correlations between excitatory and inhibitory inputs decrease the membrane potential co-variations which leads to active decorrelation of neuronal spiking, as suggested theoretically and observed experimentally in cortical slices. This decorrelation is made possible by cancellation of correlated excitatory and inhibitory inputs at resting potential. Here, we asked whether active decorrelation operates equally on signal and noise correlations and whether it depends on network state and neuronal distance. We studied a recurrent network of cortical neurons in culture and developed an innovative optical device to stimulate optogenetically a large neuronal population with both spatial and temporal precision. Using an ultra-fast video-projector, we stimulated 1 to ~100 neurons with independent trains of light pulses that evoked action potentials with high temporal resolution. We experimentally examined correlation between inhibitory and excitatory inputs in the recurrent network when a sub-population is activated by distinct stimuli which correlation is varied at will. We show that active decorrelation operates in a recurrent network of neurons in culture and that a systematic delay between excitation and inhibition is present. This lag implies that a substantial amount of correlated activity persists and can be modulated by an external input. Although signal and noise are both actively decorrelated, signal correlation remains dominant, allowing information to be encoded by correlated activity of neurons. Additionally, we observe that signal correlation remains high even for distant neurons while noise correlation decays rapidly. These observations suggest that noise correlation results from local shared inputs. Because signal-to-noise ratio increases with distance, a combination of distant neurons would enable to reconstruct the original signal by averaging out noise.

## II-2. Local feed-forward dynamics of scale-invariant neuronal avalanches

Shan Yu  
Andreas Klaus  
Hongdian Yang  
Dietmar Plenz

YUSHAN@MAIL.NIH.GOV  
AKLAUS.MAIL@GMAIL.COM  
YHONGDIAN@GMAIL.COM  
PLENZD@MAIL.NIH.GOV

National Institute of Mental Health/NIH

The identification of cortical dynamics strongly benefits from the simultaneous recording of as many neurons as possible. Yet current technologies provide only incomplete access to the mammalian cortex from which adequate conclusions about dynamics need to be derived. Here, we identify proper constraints introduced by the sub-sampling with a limited number of electrodes, i.e. spatial 'windowing', for well-characterized critical dynamics

–neuronal avalanches–in cortical resting activity. Local field potentials (LFP) were recorded from superficial pre-motor and prefrontal cortex in two awake macaque monkeys during rest using chronically implanted, high-density 96-microelectrode arrays. Negative deflections in the LFP (nLFP) were identified on the full as well as compact sub-regions of the array quantified by the number of electrodes  $N$ , that is, the window size (10–95 electrodes). Spatiotemporal nLFP clusters organized as neuronal avalanches, i.e., the probability in cluster size,  $p(s)$ , invariably followed a power law with exponent  $-1.5$  up to  $N$ , beyond which  $p(s)$  declined more steeply and varied with  $N$  and the LFP filter parameters, producing a 'cut-off'. Importantly, clusters of  $s < N$  emerged from the local feed-forward propagation of nLFPs and, consequently, consisted mainly of neuronal activity from unique, non-repeated cortical sites. In contrast, clusters in the cut-off contained many repeated activations in line with the fact that clusters of  $s > N$  with no repeats cannot be observed within the window. Only clusters with  $s < N$  carried information about spatial cluster organization. Our findings were confirmed in a noise-free neuron-electrode network model. Thus, avalanche identification and statistical analyses need to be constrained to the observation window, in order to reveal the scale-invariant organization of avalanches dominated by locally unfolding feed-forward cascades.

### **II-3. Decorrelation in networks of spiking neurons: from microscopic to macroscopic magnitudes**

Jesus Manrique<sup>1</sup>  
Nestor Parga<sup>2</sup>

JESUSMANRIQUEGOMEZ@GMAIL.COM  
NESTOR.PARGA@UAM.ES

<sup>1</sup>Departamento Fisica Teorica UAM  
<sup>2</sup>Universidad Autonoma de Madrid

Cortical recurrent networks can decorrelate their spiking activity even if the probability of common inputs is large (Renart et al, 2010). An intriguing prediction of this study is that in large networks the cross-covariances of the spiking activity are determined by the auto-covariances of single neurons in the external and recurrent populations. One can wonder whether there is still more statistical structure in the dynamics of cortical networks that could be derived from the activity of the external population and of single neurons in the recurrent network. Most importantly, one would like to know whether macroscopic quantities such as the MUA spectrum could also be predicted. We addressed these issues considering networks of leaky integrate-and-fire inhibitory neurons with slow synaptic kinetics (Moreno-Bote and Parga, 2004; 2006), dense connectivity and strong synapses. We obtained the following results: 1. A self-consistent solution of the first- and second-order statistics of the spiking activity shows that pair spike-count cross-covariances satisfy balance relationships, as in networks of binary neurons. Also, total current correlations are small but those between the components of the currents are on the order of the connection probability. Analytical results are supported by detailed numerical analysis of how correlations scale with the size of the network. Convergence is obtained for networks with about 104 neurons. 2. Cross-covariance functions of the spiking activity also obey balance relationships. They are predicted from the auto- and cross- covariance functions of the external population and the auto-covariance function of neurons in the recurrent network. 3. The power spectrum of the activity of a subpopulation of neurons is determined by the same functions. Analytical results for networks of excitatory and inhibitory neurons with slow inhibitory synaptic kinetics show that the power spectrum of the LFP is predicted from the covariance functions in the external population.

### **II-4. Learning-based crossmodal suppression of ongoing activity in primary cortical areas of the awake rat**

Jozsef Fiser<sup>1</sup>  
Benjamin White<sup>2</sup>

FISER@BRANDEIS.EDU  
BLWHITE@BRANDEIS.EDU

<sup>1</sup>Central European University  
<sup>2</sup>Brandeis University

Ongoing activity is ubiquitous in the cortex and recently has been implied to play a significant functional role in shaping sensory-evoked responses by reflecting the momentary internal state of the brain together with its knowledge about the external world (Berkes et al., Science 2011). A wide variety of studies also reported that ongoing neural signal variability in many cortical areas gets reduced when a stimulus is presented (Churchland et al., Nat Neuro 2010), raising the possibility that this reduction is functionally linked to the decrease of uncertainty due to the stimulus onset (Orban et al., Cosyne 2011). This conjecture would imply that just as stimulus onset reduces uncertainty and hence neural response variability, ongoing signal variability should also decrease due to reduction of uncertainty when the animal can apply acquired long-term knowledge about the environment in a particular situation. To test this hypothesis, we recorded neural activity with multi-electrode arrays simultaneously from the primary visual and gustatory cortices of rats while they learned to associate visual cues with delayed water reinforcement. We found within-modality, cue-dependent suppression of variability of the evoked activity in the visual cortex in line with previous reports. However, we also found that, independent of firing rate changes, spike count variability in the absence of any sensory stimuli was significantly suppressed both during the delay periods in the primary visual cortex, and during the cue period in the ongoing activity of the gustatory cortex. Importantly, this suppression both within and across modalities occurred only in animals that learned the task. These findings demonstrate that not only domain-specific stimulus-evoked, but also experience-based, internally-generated cross-modal signals are capable of suppressing the variability of ongoing activity in the primary cortex supporting the proposal that this activity is not noise but rather it represents information about particular behaviorally-relevant conditions.

## II-5. The sign rule and beyond: boundary effects, flexibility, and noise correlations in population codes

Yu Hu  
Joel Zylberberg  
Eric Shea-Brown

University of Washington

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

Over repeat presentations of the same stimulus, sensory neurons show variable responses. This “noise” is typically correlated between pairs of cells, and a question with rich history in neuroscience is how these noise correlations impact the population’s ability to encode the stimulus. Here we investigate how information varies as a function of noise correlations, with all other aspects of the problem – neural tuning curves, etc. – held fixed. Compared to previous work, we consider a very general setting. We make no assumptions on tuning curves, and do not restrict to particular correlation structures such as the “limited-range” correlations often considered. This approach leads us to derive mathematical theorems, which yield unifying insights for several common (local and global) measures of coding performance. Our main findings are as follows: (1) We generalize previous results to prove a sign rule (SR) – if noise correlations between pairs of neurons have opposite signs vs. their signal correlations, then coding performance will improve compared to the independent case. (2) As also pointed out in the literature, the SR is not a necessary condition for good coding. We show that a diverse set of correlation structures can improve coding, many of which violate the SR. There is structure to this diversity: we prove that the optimal correlation structures must lie on boundaries of the possible set of noise correlations. (3) We provide a novel set of necessary and sufficient conditions, under which the coding performance (in the presence of noise) will be as good as it would be if there were no noise present at all. As experimentally observed noise correlations can have diverse structures including those violating the SR [Cohen:2011], our theory provides a framework to explain such flexibility and suggests a new organizing principle of correlations on the boundary.

## II-6. Noise correlations induced by computation

Ingmar Kanitscheider  
 Ruben Coen-Cagli  
 Alexandre Pouget  
 University of Geneva

INGMAR.KANITSCHIEDER@UNIGE.CH  
 RUBEN.COENCAGLI@UNIGE.CH  
 ALEXANDRE.POUGET@UNIGE.CH

Sensory inputs are usually noisy and uncertain, limiting the information available to subsequent processing by the brain. This implies that information must saturate as the number of neurons increases. This in turn must induce noise correlations because information would not saturate if neurons were independent. Experimental measures have indeed revealed that neurons are correlated and that correlations tend to be positive and proportional to the tuning similarity. This type of correlations has long been thought to be responsible for the information limitation but theoretical work has shown this is not the case. Instead, we have recently found that information is limited by differential correlations: correlations proportional to the product of the derivatives of the tuning curves, which look markedly different from the correlations observed in vivo. This previous work raises two fundamental questions: why do correlations look the way they do in vivo and what is the origin of differential correlations? Here, we investigate these issues by considering the type of correlations that emerge in a network performing a sensory computation, such as extracting the orientation of contours from noisy images. First, we report that the overlap between orientation-tuned filters leads naturally to biologically realistic correlations. Second, we show the latter contain differential correlations due to limited input information, and that all input information can be preserved despite Poisson variability. Third, we find that differential correlations tend to be very small, to the point of being nearly impossible to detect in the correlation pattern. Fourth, we show that these differential correlations explain why the discrimination threshold of single neurons is nearly as small as the discrimination threshold of the animal. Our results suggest that information limiting correlations in the brain might be largely due to the computation performed by neural circuits rather than internal dynamics or intrinsic noise.

## II-7. Biophysically inspired cell assembly detection at multiple scales

Yazan N Billeh<sup>1</sup>  
 Michael T Schaub<sup>2</sup>  
 Costas A Anastassiou<sup>3</sup>  
 Mauricio Barahona<sup>2</sup>  
 Christof Koch<sup>3</sup>

YBILLEH@CALTECH.EDU  
 MICHAEL.SCHAUB09@IMPERIAL.AC.UK  
 COSTASA@ALLENINSTITUTE.ORG  
 M.BARAHONA@IMPERIAL.AC.UK  
 CHRISTOFK@ALLENINSTITUTE.ORG

<sup>1</sup>California Institute of Technology

<sup>2</sup>Imperial College London

<sup>3</sup>Allen Institute for Brain Science

Neuronal monitoring techniques such as calcium imaging and multi-electrode arrays have enabled neuroscientists to monitor activity from hundreds or more neurons simultaneously. Identifying neurons that cooperate in some form within populations is of primary importance for systems neuroscience. Here we introduce a simple biophysically inspired technique for extracting a directed (causal) functional connectivity matrix between both excitatory and inhibitory neurons based on their spiking history. The resulting functional network representation is used with the recently introduced Markov stability method (Delvenne et. al, 2010, PNAS) to identify groups of related neurons (cell assemblies) in the recorded time series at multiple levels of granularity, without prior knowledge of their (hierarchical) relation. This clustering technique uncovers structural features in a network by observing how a Markovian dynamical process such as diffusion evolves over time on the network. We assess our method through synthetic and simulated spike-train data from leaky-integrate-and-fire networks, where we observe that detected spike-train communities correlate well with structural features of the underlying neuronal networks. We further exemplify the utility of the method by analyzing experimental data from retinal ganglion cells of mouse and salamander, in which we identify groups of cells that corresponds to known functional cell types. Finally, we apply our algorithm to the spike trains of hippocampal recordings from rats exploring a linear track, and detect

place cells with high fidelity. This work provides an analysis framework that assists in identifying, aggregating and understanding relevant temporal and spatial information hidden in the data generated from modern systems neuroscience experiments.

## II-8. Decoding thalamic inputs from cortical activity

Audrey Sederberg  
Stephanie E Palmer  
Jason MacLean  
University of Chicago

SEDERBERG@UCHICAGO.EDU  
SEPALMER@UCHICAGO.EDU  
JMACLEAN@UCHICAGO.EDU

The coding scheme employed by the cortical microcircuit remains unclear. The classical approach to characterizing the cortical code has been to focus on the spike activity of single cortical neurons and define such properties as receptive fields, tuning curves, and stimulus selectivity. However, cortical neurons are highly interconnected, with a bias toward proximal connectivity, and it is necessary to consider the response of nearby neurons when considering information representation in the brain. Here we densely sample the spiking activity of 200-600 cortical neurons spanning multiple columns and layers using high-speed two-photon imaging. We electrically stimulate two separate anatomical locations in thalamus, corresponding to separate columns in sensory cortex, and image circuit activity outside of both primary columns. This allows us to emphasize intracortical connectivity and limit the influence of direct thalamic input on our results. We used a decoding approach to reveal in part how information is encoded within the cortical microcircuit. We find that activity at the single-neuron level can decode the identity of the stimulus location. We also identify instances of cooperative decoding in pairs of neurons, finding decoding rates that are higher than either individual cell's performance. Among these cooperative pairs, we find that noise correlations play a strong role in decoding. Additionally, many cell pairs derive their decoding power from a difference in stimulus response lag. However, a cortical cell does not necessarily have access to precise knowledge of when a stimulus was delivered, and the assumption that this information is accessible may be faulty. Thus, we define a decoding model with a 'cortical perspective,' where the onset of activity, rather than absolute time since stimulus delivery, is used to mark time. When we decode using this framework, we find that decoding rates decrease relative to decoders with exact stimulus onset information, but remain above chance.

## II-9. Topological analysis of hippocampal correlations reveals non-random, low-dimensional clique topology

Vladimir Itskov<sup>1</sup>  
Carina Curto<sup>1</sup>  
Eva Pastalkova<sup>2,3</sup>  
Chad Giusti<sup>1</sup>

VLADIMIR.ITSKOV@UNL.EDU  
CCURTO2@UNL.EDU  
PASTAK@JANELIA.HHMI.ORG  
CGIUSTI2@UNL.EDU

<sup>1</sup>University of Nebraska, Lincoln  
<sup>2</sup>Janelia Farm Research Campus  
<sup>3</sup>Howard Hughes Medical Institute

It is often hypothesized that pairwise correlations reflect the information processing mechanisms of the underlying neural circuit. What one can actually learn from the structure of pairwise correlations, however, is often unclear. Correlations can be organized into graphs, with edges reflecting the strengths of correlation between pairs of neurons. In brain areas with convex receptive fields, the structure of the neural code has strong implications for the structure of cliques in these graphs. This can be detected by examining the clique topology of the graph, and is closely tied to the topology of the underlying space. For example, for hippocampal place cell activity during spatial navigation, pairwise correlation graphs are expected to have low-dimensional clique topology, due to the arrangement of place fields in a low-dimensional environment. However, it is not at all clear whether such a low-

dimensional structure should be present in hippocampal activity during sleep or non-spatial behavior. In this work, we are motivated by the question: what can we infer from the intrinsic structure of the neural activity as measured by pairwise correlations? To address this question, we analyzed the structure of cliques in pairwise correlation graphs obtained from population activity in hippocampus under several behavioral conditions: spatial exploration, wheel running, and sleep. We found these graphs to be highly non-random by comparing their clique structure to that of random graphs with matching parameters. The clique topology during open field exploration and wheel running was low-dimensional as compared to random graphs, consistent with what is expected from place cells during spatial navigation. Remarkably, the same low-dimensional structure was also observed in the sleep data. This suggests that the low-dimensional structure of neural activity in hippocampus is due to the structure of the underlying circuit, rather than the structure of external stimuli during navigation.

## II-10. A novel rate-correlation relationship might significantly improve population coding in retina

Joel Zylberberg<sup>1</sup>  
Maxwell Turner<sup>1</sup>  
Jon Cafaro<sup>1,2</sup>  
Fred Rieke<sup>1,3</sup>  
Eric Shea-Brown<sup>1</sup>

JOELZY@UW.EDU  
MHTURNER@UW.EDU  
CAFAROJ@GMAIL.COM  
RIEKE@UW.EDU  
ETSB@UW.EDU

<sup>1</sup>University of Washington

<sup>2</sup>Duke University

<sup>3</sup>Howard Hughes Medical Institute

How do retinal ganglion cells coordinate their activities to convey visual information to the brain? To investigate this, we recorded the co-variability in the responses of pairs of direction selective ganglion cells, over repeat presentations of the same stimuli. Much theoretical and experimental work has investigated how these "noise correlations" impact population coding. Of particular interest are the signs of the noise correlations (positive or negative), and how they relate to the signs of the signal correlations, which describe the co-variation in the mean responses of the cells as the stimulus varies. While noise correlations are defined for each particular stimulus, signal correlations are defined for ensembles of stimuli. Thus, it is important to understand if and how noise correlations vary between stimuli. Previous circuit models – and associated experiments – showed stimulus dependent noise correlation magnitudes, but did not show stimulus-dependent signs of noise correlations. For contrast, we observe that both the magnitudes and the signs of the noise correlations can vary systematically with the stimulus, representing a novel rate-correlation relationship. Through alternating voltage clamp experiments, we traced the origin of these correlations to common fluctuations in the excitatory and inhibitory synaptic inputs to multiple cells, and in the process built a mechanistic model that can reproduce the observed stimulus dependence of firing rates and correlations. Using this model, we ask how population coding performance is affected if the correlations are removed altogether, or if the stimulus-dependent correlations for each cell pair are replaced by their mean value (constant for all stimuli). We find that the stimulus dependent correlations result in better coding performance than when the correlations are removed, or when the noise correlations are constant for all stimuli. Stimulus-dependent correlations, with well-defined origins in circuit mechanisms, might thus be an important feature of the retinal population code.

## II-11. Structure of neuronal correlation during eye movement planning in FEF

Sanjeev Khanna  
Adam Snyder  
Matt Smith  
University of Pittsburgh

SANJEEVBKHANNA@GMAIL.COM  
ADAM@ADAMCSNYDER.COM  
SMITHMA@PITT.EDU

Nearby neurons in visual cortex exhibit correlated spiking activity. Computational analyses suggest there is a strong link between the correlation among neurons and the amount of information that can be represented in a neuronal population, although that link depends on how correlation varies with respect to neuronal tuning curves, distance and other properties. Given its importance in neural coding, the structure of neuronal correlation in higher order cortex is particularly important to understanding perception and behavior. Very little is known, however, about the population activity in these areas. We investigated correlated variability in the frontal eye fields (FEF), a part of prefrontal cortex that contains neurons responsive to visual stimuli and near the time of an eye movement. We predicted that the structure of neuronal correlation in FEF would be consistent with that found in visual cortex, reflecting a conserved architecture across neocortex. We used a multi-contact probe to record from FEF neurons in alert rhesus macaque monkeys. The animals performed a conventional memory guided saccade task and we measured neuronal correlation on both short and long time scales (spike count correlation and synchrony). We found that both spike count correlation and synchrony decreased with increasing neuronal distance, consistent with previous observations in visual cortex. Additionally, correlation was highest among pairs of neurons with strong visual responses, and weakest among pairs of neurons with strong motor responses. We conclude that the structure of neuronal correlation appears to be based on circuitry that is conserved across much of sensory cortex, while neuronal populations closer to the motor output operate under different constraints.

## II-12. Spike count correlation relates to the phase and amplitude of EEG oscillations measured at the scalp

Adam Snyder<sup>1</sup>  
Cory Willis<sup>2</sup>  
Matt Smith<sup>1</sup>

ADAM@ADAMCSNYDER.COM  
CWILLIS53@GMAIL.COM  
SMITHMA@PITT.EDU

<sup>1</sup>University of Pittsburgh

<sup>2</sup>University of Pennsylvania

An understanding of neuronal communication, both within and between regions of the brain, is fundamental to explaining perception and behavior. Communication among neurons has been studied at a variety of scales, from the relatively coarse resolution afforded by scalp electroencephalogram (EEG) to the fine resolution achieved in extracellular microelectrode recordings of individual neurons. Although the knowledge gained to date by these parallel efforts has been invaluable, a detailed understanding of neuronal communication will ultimately need to bridge these scales so that the relationships between individual neurons can be understood in the context of the interactions of large-scale brain regions, and vice versa. Here we take the first steps toward this goal with the use of simultaneous EEG and single unit recordings in alert macaque monkeys. Since EEG signals are generally considered to be manifestations of coherent synaptic input to a region of cortex, we reasoned that EEG oscillations—which have been linked to a variety of perceptual and cognitive processes, but whose computational mechanisms remain unknown—would be related to the correlation in spiking activity between pairs of neurons in the underlying cortical area. To test this prediction, we recorded population spiking activity from microelectrode arrays chronically implanted in visual area V4 of awake, behaving macaques simultaneously with signals from EEG electrodes on the scalps of the animals. Consistent with our predictions, we found that the amplitude and phase of EEG oscillations indexed the correlation of the underlying spiking activity. We found this relationship between EEG oscillations and spiking correlation to be frequency-specific, which may be related to important functional distinctions between different EEG frequency bands. This is a crucial development that directly links EEG signals to physiological processes with clear computational consequences, and opens an avenue for understanding

spiking correlations in terms of the interaction among large-scale brain networks.

### II-13. Efficient coding theory in nonlinear networks with noise

Jerome Tubiana<sup>1</sup>  
Zhuo Wang<sup>2</sup>  
Daniel Lee<sup>2</sup>  
Haim Sompolinsky<sup>3</sup>

JEROME.TUBIANA@ENS.FR  
WANGZHUO@SAS.UPENN.EDU  
DDLEE@SEAS.UPENN.EDU  
HAIM@FIZ.HUJI.AC.IL

<sup>1</sup>Ecole Normale Supérieure

<sup>2</sup>University of Pennsylvania

<sup>3</sup>Hebrew University

Theories for efficient coding rely on maximizing mutual information between high-dimensional input patterns and the output activations of neurons. The inputs are assumed to be linearly mixed by a set of weights and passed through saturating scalar nonlinearities. The weights along with the nonlinearities are optimized to determine the most efficient code that maximizes mutual information. Unfortunately, few analytic solutions exist for this optimization problem. Conventional infomax assumes zero input noise and vanishing output noise, resulting in a solution where the nonlinearities equalize the output distributions and the weights render the outputs as statistically independent as possible. For Gaussian inputs, this implies the optimal weights decorrelate the inputs, supporting the principle of redundancy reduction stated by Attneave and Barlow. But what about the presence of input noise? We consider a nonlinear network model where both input and output noise are present. We parameterize the noise terms using the input to output noise variance ratio, and derive a novel analytic solution when this ratio approaches infinity and the overall noise scale vanishes. In contrast to the infomax solution, in this limit the analytic solution results in a one-third power dependence for the optimal nonlinearity that does not equalize the output distributions. Additionally, we analytically compute the optimal weights for Gaussian inputs that results in imperfect decorrelation of the outputs. We also compare this analytic solution with infomax for dimensionality reduction, where the number of outputs is less than the number of inputs. We use perturbative expansions to explain the most significant terms in both limits and their roles in determining the optimal weights and nonlinearities for efficient codes. Finally, we discuss the implications our analysis has for modeling neural representations in various layers of sensory processing.

### II-14. Determining the optimal distributions of neural population codes

Wentao Huang  
Kechen Zhang  
Johns Hopkins University

WNTHUANG@GMAIL.COM  
KZHANG4@JHMI.EDU

It has long been suggested that the sensory systems might approach an information-theoretic optimum for neural coding and computation. Despite many progresses over the years, it remains unclear how well the optimality hypothesis holds for a large population of neurons with realistic distribution of response properties. One major difficulty is that in order to accurately evaluate the mutual information between the stimulus and the population response, one has to average over all possible response patterns of the whole population, which quickly leads to combinatorial explosion as the population size increases. The computational intractability makes it virtually impossible to derive theoretical optimal distributions for many realistic neural populations. Here we present an efficient asymptotic formula for mutual information between a large neural population and the stimulus, which in general is a high dimensional vector. The original information optimization problem is converted into a convex optimization equation about a population density distribution of tuning function parameters, and the problem can be solved efficiently by numerical optimization of a convex objective equation. This algorithm allows us to obtain and compare the optimal neural populations under many conditions and constraints. For example, the optimal



distribution of the preferred stimuli across the population can become either continuous or discrete, and can be clustered either around those stimulus features that are the most prevalent in the input, or around features that are rare, all depending on the shapes of the tuning functions. The emergent optimal population distributions may account for various types of sensory neuron distributions known to exist in the sensory systems, providing additional support to the optimal coding hypothesis for large neural populations.

## II-15. Transformation of optimal motion encoding in the mouse early visual system

Alfred Kaye

Edward M Callaway

Tatyana Sharpee

Salk Institute for Biological Studies

ALFRED.KAYE@GMAIL.COM

CALLAWAY@SALK.EDU

SHARPEE@SALK.EDU

In mice, direction selective retinal ganglion cells (DSRGCs) form a complete representation of visual motion. It has recently become clear that this information reaches the lateral geniculate nucleus (LGN) and visual cortex of mice, yet its role in cortically driven behavior remains unclear. In this work, we show that DSRGC tuning curves are optimal for carrying information about errors between preferred and observed directions of motion. This result contrasts with Fisher Information optimal theories that predict that neurons should be more sharply tuned than is observed, do not match the observed direction preferences of retinal ganglion cells, and do not naturally take into account synergies when multiple neurons jointly encode the same variable. An information-theoretic theory of optimal direction selectivity for encoding natural optic flow distributions reconciles experimentally observed features of retinal motion processing with the efficient coding hypothesis. The theory shows that synergistic encoding of motion accounts for the observation that the LGN both combines subtypes of DSRGC inputs and sharpens direction selectivity, and predicts that anisotropies in the natural distribution of motion (e.g., horizontal axis predominance) would give rise to anisotropies in direction tuning. Finally, we extend this theory to consider schemes in which cortical direction selectivity is derived from the DSRGC motion pathway rather than from the classical model in which direction is computed de novo in cortex.

## II-16. Statistical inference for directed phase coupling in neural oscillators

HaDi MaBouDi<sup>1</sup>

Hideaki Shimazaki<sup>2</sup>

Mehdi Abouzari

Shun-ichi Amari<sup>2</sup>

Hamid Soltanian-Zadeh<sup>1</sup>

MABOUDI@IPM.IR

SHIMAZAKI@BRAIN.RIKEN.JP

ABOUZARI@IPM.IR

AMARI@BRAIN.RIKEN.JP

HAMIDS@RAD.HFH.EDU

<sup>1</sup>Institute for Research in Fundamental Sciences

<sup>2</sup>RIKEN Brain Science Institute

Phase coupling and synchronization have been an important concept in the study of oscillations of neural systems for the past few decades. The phase coupling of various neuronal measurements was observed within local activity of cortical and subcortical areas and even between activities in distant brain regions [1]. Nevertheless, previous methods for estimating phase couplings are limited. Both standard phase correlation and phase coherence compute static pairwise dependency of two oscillators, therefore are limited to uncover underlying couplings of multiple oscillators, and are blind to their underlying dynamics. Recently, Cadieu and Koepsell [2] provided a model-based method that allows us to directly and simultaneously estimate underlying undirected couplings of multiple oscillators using the Kuramoto model of an oscillator network. However, this and other previous methods cannot reveal causal relations between the oscillators. Here we propose a method to estimate the causality in the phase space from measurements of neural oscillations. A new probabilistic method was developed to estimate

directed coupling parameters of the phase oscillators from noisy multivariate circular time-series data. In this method, we constructed a probabilistic description of the Kuramoto model, and developed an algorithm to estimate the coupling parameters under the maximum likelihood principle. We demonstrate that the proposed method recovers the underlying direction and weights of couplings between noisy dynamic oscillators while previous approaches by the phase correlations and the Granger causality were unable to recover these values correctly. We stress that the method performs well on a relatively large number of oscillators even in the presence of noise. The estimated directed graphs provide us useful tools to uncover causal relations of cortical networks, such as casual interactions of local micro-circuitries, and extract their network topology. [1] Buzsaki, G. and et al. (2013) Neuron. [2] Cadieu and Koespell (2010) Neural Comput.

## **II-17. Nonlinear transfer of signal and noise correlations in cortical networks**

Dmitry Lyamzin<sup>1,2</sup>  
Samuel Barnes<sup>3</sup>  
Tara Keck<sup>3</sup>  
Nicholas Lesica<sup>1</sup>

D.LYAMZIN@UCL.AC.UK  
SAMUEL.BARNES@KCL.AC.UK  
TARAKECK@GMAIL.COM  
N.LESICA@UCL.AC.UK

<sup>1</sup>University College London

<sup>2</sup>Ludwig Maximilian University of Munich

<sup>3</sup>King's College London

Correlated spiking is a prominent feature of cortical activity. During sensory processing, information transmission is controlled by the balance of signal and noise correlations, and their structure across a network is an important signature of its organization and function. Studies of the relationship between signal and noise correlations typically report a positive dependency, with variations due to changes in stimulus properties and behavioral context, but the factors that establish and modify correlations are not well understood. Here we show that the correlation structure in cortical networks is controlled by correlation transfer nonlinearities with complex dependencies on the properties of signal and noise inputs. We find that a change in the strength or correlation of either signal or noise inputs can affect both spike train signal and noise correlations. We first demonstrate these effects directly using whole-cell recordings of responses to injected currents and show that they are reproduced within an analytical framework based on a simple threshold nonlinearity. We then analyze the correlation structure of spike trains recorded in vivo and show that correlation transfer nonlinearities impose a variety of dependencies between signal and noise correlations across populations of neurons in auditory cortex, and that these effects underlie changes in correlation structure that result from changes in sound location. Our results demonstrate that correlation transfer nonlinearities can shape the correlation structure of cortical networks, even imposing a negative dependency between signal and noise correlations under certain conditions. It is likely that the correlation transfer nonlinearities that we observe in auditory cortex also operate in other sensory systems and, thus, our results are a critical step toward a general understanding of how correlated spiking relates to the function of cortical networks.

## **II-18. Existence of multiple bursting oscillators in cultured neural network and transitions among them**

June Hoan Kim  
Ryoun Heo  
Joon Ho Choi  
Kyoung Lee  
Korea University

LOVEPHY@GMAIL.COM  
NEUROCOMPLEX@EMPAS.COM  
FRAMEBUILDER@GMAIL.COM  
KYOUNG@KOREA.AC.KR

Synchronized neural bursts (SBs) are a salient dynamic feature of biological neural networks, having important roles in brain functions. Some earlier studies demonstrated traces of determinism embedded in the sequence

of SB events. For example, inter-burst-intervals (IBIs) followed a long-tailed distribution, an indication that a history-dependent memory was involved in the generation of SBs. Moreover, some motifs of SB were observed to recur over time. Very recently, we also obtained a set of experimental results suggesting that the seemingly random sequence of spontaneous SBs could be generated by a single sub-threshold burst oscillator working in the presence of strong noise [Choi et al., PRL 2012]. The existence of a single oscillator could be inferred only from the responses of the system to extrinsic electrical stimulation. In this study, we report that similarly prepared cultured neural networks can support not just one, but several oscillators, having different periods that cover three orders of magnitude, ranging approximately from tens of milliseconds to ten seconds. These oscillators are very noisy with fluctuating periods. Interestingly, they implement time-sharing and take turns. That time-sharing strategy for modulations in neural network is good to perform multi-modality or multi-functionality. The transitions among them are sudden and the dwelling times are random. Nevertheless, the transitions clearly support several preferred passages. In other words, the cultured neural networks seemed to support several attractors that may have a non-trivially structured basin of attraction. Considering that the system under consideration was a nonlinear system with an almost infinite degree of freedom, the very existence of low-dimensional attractors and preferred transition passages was quite surprising. Through a set of bicuculline (antagonist of inhibitory neural connections) applications, we demonstrate that the dominance of one oscillator over the others can be controlled by the overall inhibitory cell-to-cell connectivity.

## II-19. Extracting a signal in noise: Effect of bursts and tonic spikes on detection and discriminability

Clarissa Shephard<sup>1,2</sup>  
 Christian Waiblinger<sup>3</sup>  
 Cornelius Schwarz<sup>3</sup>  
 Garrett B Stanley<sup>1,2</sup>

CSHEPHARD3@GATECH.EDU  
 CHRISTIAN.WAIBLINGER@UNI-TUEBINGEN.DE  
 CORNELIUS.SCHWARZ@UNI-TUEBINGEN.DE  
 GARRETT.STANLEY@BME.GATECH.EDU

<sup>1</sup>Georgia Institute of Technology

<sup>2</sup>Emory University

<sup>3</sup>University of Tuebingen

Sensory information is transformed from physical attributes of a stimulus at the periphery to complex representations at the level of cortex that ultimately give rise to perception. Despite the complexity of the sensory information, neurons can have strongly nonlinear characteristics that give sensory systems the remarkable capability to extract relevant information (features) from noisy environments (context conditions). In this work, we explored the effect of stimulus context on the encoding of features in the whisker thalamocortical circuit of the fentanyl-cocktail anesthetized rat using a classic signal-in-noise paradigm. The sensitivity of thalamic single and multi units to detect features (signal) in varying contexts (noise) was estimated using signal detection theory. Neurons' sensitivity increase monotonically with the detectability of the stimulus input, however an increase in the noise amplitude decreases the response to the embedded features in absolute terms. Ultimately this leads to a loss of detectability of embedded features with increasing context amplitude. Interestingly, the loss of detectability is not due to increasing neural responses to noise, but rather decreasing neural responses to features. We hypothesize that increasing the context (noise) amplitude is depolarizing the thalamus, leading to a more tonic encoding scheme whereas presenting the features in isolation, while the thalamus is extremely hyperpolarized, leads to a burst encoding scheme. Preliminary results suggest that cells ( $n=12$ ) fire more 'burst' spikes in response to features in isolation than in noise and that this leads to a higher detectability, but a lower discriminability. This is supported by previous results in the vibrissa thalamocortical circuit that suggest discrimination is enhanced in response to a periodic adapting stimuli.

## II-20. Decoding analog multiunit activity in a brain-machine interface for individuals with tetraplegia

Wasim Malik<sup>1</sup>  
John Donoghue<sup>2</sup>  
Leigh Hochberg<sup>2</sup>

WQM@MIT.EDU  
JOHN\_DONOGHUE@BROWN.EDU  
LEIGH\_HOCHBERG@BROWN.EDU

<sup>1</sup>Harvard Medical School  
<sup>2</sup>Brown University

Brain-machine interfaces (BMIs) offer the potential to restore communication and movement ability in people with paralysis following spinal cord injury, brainstem stroke, ALS, or limb amputation. An intracortical BMI records motor-related neural activity, decodes it to estimate the intended movement, and controls assistive devices accordingly. Single-unit activity (SUA), obtained from motor cortical neuronal ensembles, is well established as a useful signal in BMIs due to its extensively studied relationship to voluntary movement. Despite its information richness, the SUA's temporal instability poses a practical challenge for chronic BMIs, since SUA decoding involves frequent and time-consuming spike-sorting and filter recalibration. To address this issue, we consider a signal that samples from broader collections of neural activity, namely the analog wideband multiunit activity (MUA). Besides improving decoder longterm stability and robustness, MUA can help avoid the time-consuming process of manual spike-sorting. Following recent animal studies showing that MUA encodes movement intention and its use in an early offline control of a robot wrist, we investigate the potential of MUA for motor decoding in a clinical BMI. We analyzed MUA decoding performance as part of the ongoing BrainGate pilot clinical trial (IDE) with participants S3 and T2 with brainstem stroke. We recorded the raw wideband neural activity sampled at 30 kHz using an intracortical array of 96 microelectrodes implanted in the precentral gyrus in the region of arm representation. The raw signal at each electrode was bandpass filtered (300 Hz – 6 kHz), integrated (50 ms non-overlapping windows), and downsampled to obtain MUA. For comparison, we computed the SUA as the spike rate within identical timebins from all available sorted units. The participant imagined moving a computer-controlled cursor on a two-dimensional screen during a 4-target center-out-back task under open-loop motor imagery without feedback control. Neural activity and cursor kinematics recorded during this period were used to calibrate a Kalman filter assuming a velocity-based neural tuning model. During subsequent closed-loop control assessment with visual feedback, the filter was used to decode neural activity in real time and move the cursor in an 8-target center-out-back task and a step-tracking pursuit task optimized to acquire performance metrics using the Fitts law. We conducted leave-one-out cross-validated offline open-loop decoding analysis from 5 different days with one participant. MUA yielded higher Pearson's correlation coefficient between actual (computer-controlled) and decoded cursor velocity than SUA (0.80 vs. 0.73;  $p < 0.001$ ). In closed-loop online real-time control sessions on 10 different days with both participants, the MUA decoder provided a higher target acquisition rate (52% vs. 39%), higher path efficiency (0.58 vs. 0.42;  $p < 0.01$ ), and statistically equivalent target acquisition time compared to the SUA-based decoder. Thus our open- and closed-loop decoding analysis confirmed that MUA consistently performed at least as well as SUA in standard motor tasks. Our analysis demonstrates the feasibility of decoding MUA in a BMI to estimate intended movement kinematics. As MUA decoding can reduce system complexity and improve longterm stability of a BMI without sacrificing decoding accuracy, it may enable autonomous and robust BMIs for clinical function restoration and rehabilitation of people with tetraplegia.

## II-21. Maximum Variance Differentiation (MVD) explains the transformation from IT to Perirhinal cortex

Marino Pagan<sup>1</sup>  
Eero Simoncelli<sup>2,3</sup>  
Nicole Rust<sup>1</sup>

MARINOPAGAN@GMAIL.COM  
EERO.SIMONCELLI@NYU.EDU  
NRUST@PSYCH.UPENN.EDU

<sup>1</sup>University of Pennsylvania  
<sup>2</sup>New York University  
<sup>3</sup>Howard Hughes Medical Institute

Neural processing of signals for object recognition and target search has been shown to implement an "untangling" transformation, whereby initial population representations are converted into a format that is linearly separable. Despite the appeal of this description, understanding the precise nature of these computations has proven difficult. Here we propose that a transformation analogous to the well-known Energy model for V1 complex cells can explain the untangling of target-match signals flowing from IT to Perirhinal cortex, collected as monkeys performed a delayed-match-to-sample task. For any N-way classification problem, "untangling" amounts to increasing the separation between the class means of the population response. For the IT to Perirhinal transformation, we hypothesized that untangling is achieved by transforming variance differences into mean differences through the use of a squaring operation. We optimized a linear-nonlinear-linear (LNL) response model to achieve this goal. The first linear stage transforms a population of IT inputs to maximize the variance differences between the classes in the output population. These linear responses are squared, and followed by a final orthogonal linear transformation. We find that a linear decoder operating on the responses of this Maximum Variance Differentiation (MVD) model attains target match vs. distractor performance close to that of an ideal observer operating on the IT population, and far better than a linear decoder operating directly on the input IT population, suggesting that most of the target match information embedded in IT population responses lies in the class variances. Furthermore, the MVD population matched Perirhinal linear decoder performance, suggesting that an MVD transformation within Perirhinal cortex may act on input arriving from IT. These results provide evidence that within the family of LNL models, a generalization of the Energy model is sufficient to explain the untangling of visual target match information.

## II-22. Predictive coding with hodgkin-huxley-type neurons

Michael Schwemmer<sup>1</sup>

Sophie Deneve<sup>2</sup>

Eric Shea-Brown<sup>3</sup>

Adrienne L Fairhall<sup>3</sup>

SCHWEMMER.2@MBI.OSU.EDU

SOPHIE.DENEVE@ENS.FR

ETSB@WASHINGTON.EDU

FAIRHALL@U.WASHINGTON.EDU

<sup>1</sup>Ohio State University

<sup>2</sup>Ecole Normale Supérieure

<sup>3</sup>University of Washington

Recently, Boerlin et al. [Boerlin and Deneve, 2011, Boerlin et al., 2013] derived a neural network that can perform a variety of computations. The derivation is based upon the premise that the membrane potentials of neurons in the network track a prediction error between a desired output and the network estimate — and that the neurons spike only if that prediction error exceeds a certain value. This leads to a recurrent network of leaky-integrate-and-fire (LIF) neurons with a mixture of instantaneous and slower synaptic dynamics. Although this framework is flexible in that the derived network can perform any linear computation, it is important to ask whether the framework generalizes to more realistic synaptic and spike-generating dynamics. Here, we show that the predictive coding framework of [Boerlin et al., 2013] can be extended to include slow, alpha-function synaptic dynamics and Hodgkin-Huxley-like (HH) spike-generating currents. First, we explore the performance of an LIF network with alpha-function synapses. We find that increasing the magnitude of the leak current can significantly degrade performance, but that this degradation can be countered by increasing the gain of the synapses. However, this also acts to increase the firing rate of the network to unrealistic levels. Secondly, we find that the network functions but with a significant decrease in performance when one naively replaces the LIF dynamics with that of the HH-type dynamics. Moreover, increasing the overall gain of the synapses to regain accuracy in this instance quickly leads to unrealistic firing rates. We then show that one can significantly improve performance and maintain realistic firing rates by including a novel modification to the network dynamics. Importantly, we show that this adjustment allows highly accurate network performance to be attained for a wide range of network level parameters, as well as a variety of HH-type neuronal models.

## II-23. Zipf's law and criticality in multivariate data without fine-tuning

David Schwab<sup>1</sup>  
Ilya Nemenman<sup>2</sup>  
Pankaj Mehta<sup>3</sup>

DSCHWAB@PRINCETON.EDU  
ILYA.NEMENMAN@EMORY.EDU  
PMEHTA314@GMAIL.COM

<sup>1</sup>Princeton University

<sup>2</sup>Emory University

<sup>3</sup>Boston University

Recently it has become possible to measure simultaneous collective states of many biological components, such as neural activities or antibody sequences. A striking result has been the observation that the underlying probability distributions of the collective states of these systems exhibit a feature known as Zipf's law. That is, the frequency of a state is inversely proportional to its rank in a frequency table. Translating this into statistical mechanics reveals that these systems appear poised near a unique critical point, where the extensive parts of the entropy and energy are exactly equal. Here we present analytical arguments and numerical simulations showing that such behavior naturally arises in systems with an unobserved random variable, such as common input stimulus to a neural system, that affects the observed degrees of freedom. The mechanism does not require fine tuning and may help explain the ubiquity of Zipf's law in disparate systems. Our work nonetheless suggests that biological systems operate in a special regime. The system size required to exhibit Zipf's law depends on the sensitivity of the neurons to the hidden variable fluctuations. The system must be well adapted for Zipf-like behavior to emerge at moderate system sizes. In particular, we consider a directed graphical model where a hidden variable couples to the observed degrees of freedom. The coupling to the observed data may be heterogeneous and not conditionally independent given the hidden variable. We show analytically that if the distribution of the hidden variable is not too sharply peaked, under general conditions the resulting data will exhibit Zipf's law in the thermodynamic limit. This prediction is tested via numerical simulations of multiple model systems. Finally, we demonstrate this mechanism at work in data using the spiking output of a blowfly motion-sensitive H1 neuron in response to a time-varying input stimulus.

## II-24. Heaps' law as a tuning principle for retinal population codes

Bas van Opheusden<sup>1</sup>  
Gasper Tkacik<sup>2</sup>  
Olivier Marre<sup>3</sup>  
Michael Berry<sup>4</sup>

SVO213@NYU.EDU  
GASPER.TKACIK@IST.AC.AT  
OLIVIER.MARRE@GMAIL.COM  
BERRY@PRINCETON.EDU

<sup>1</sup>New York University

<sup>2</sup>Institute of Science and Technology, Austria

<sup>3</sup>Institut de la Vision - INSERM

<sup>4</sup>Princeton University

It is a long-standing belief that neurons encode relevant information in an optimally efficient way (Barlow 1961). However, it has been difficult to extend the efficient coding hypothesis to populations of correlated neurons. Recently, it has been suggested that retinal population codes are tuned to a critical point; the point where the distribution of code words starts to be dominated by collective correlations (Bialek 2007, Tkacik 2013). A direct implication of criticality is Zipf's law: after sorting code words from most to least likely, the probability of any word is related to its rank by a power law. Moreover, in a critical population, the Zipf exponent must be equal to -1. Thus, Zipf's law is often viewed as an important signature of criticality in physical systems. However, we show that non-critical population codes can also obey Zipf's law. Next, we introduce an alternative explanation for why retinal population codes should obey Zipf's law with exponent -1. We show that Zipf's law implies the so-called Heaps' law: the number of distinct words the retina uses in a particular time window scales with its duration as a power law. Moreover, if the Zipf exponent is -1, then the Heaps exponent is 1. Remarkably, even though such a system repeats some code words quite often, the number of completely novel words per unit time is asymptotically

constant. We propose that retinal populations have Zipf exponent -1 so that Heaps' law is satisfied, with exponent close to 1. This might be a valuable property to have, because the retina can encode more stimuli by using more distinct words. However, repetition of code words is useful because the brain can learn their probabilities and perform inference. We argue that Zipf's law with exponent -1 balances these two requirements optimally.

## II-25. Fast population coding in recurrent networks: single cell vs. network mechanisms

Rainer Engelken<sup>1,2</sup>  
Michael Monteforte<sup>3</sup>  
Fred Wolf<sup>1</sup>

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

<sup>1</sup>Max Planck Institute for Dynamics and Self-Organization

<sup>2</sup>Bernstein Center for Computational Neuroscience

<sup>3</sup>Georg August University

Quickly tracking changes of time-varying input is crucial for many neural computations such as coincidence detection, spike-timing dependent plasticity as well as for the processing speed of neuronal networks. Recent experiments show a surprisingly high cut-off frequency of ensembles of unconnected neurons in a fluctuation-driven regime [1-3]. Theoretical models can explain this by the fast-onset dynamics of action potentials (AP) [4-6]. Recent evidence for this hypothesis comes from whole-cell recordings, where the ability to encode high frequencies was dramatically reduced when AP onset was slowed by experimental manipulations or was slow due to immature AP generation mechanisms [7]. In contrast, recurrent networks in the balanced state can track fast-varying stimuli, as long as the rate of change of the input is sufficiently small so that the network stays balanced [8]. Until now, these two mechanisms on the single neuron and network-level have been studied separately. Here we develop a scenario where we can directly compare the two and address how the speeds at these two levels interact. We run numerically exact network simulations using an analytically solvable extension of the theta model with adjustable AP onset rapidness [12]. We drive populations of recurrently connected neurons with a sinusoidally varying input mean or variance embedded in Poissonian noise, whose net input variance is kept independent of the number of synapses so as to keep the overall activity constant. We study the frequency response of the network as a function of the AP onset rapidness, the number of synapses and the strength of the recurrent synapses, while the input statistics remain fixed. Our results indicate that the high frequency response depends strongly on the AP onset rapidness, both for recurrent and uncoupled networks. For an intermediate range, the recurrent input can facilitate tracking of high-frequency input even at low AP onset rapidness.

## II-26. Neural population coding of multiple stimuli

Emin Orhan  
Wei Ji Ma  
New York University

EORHAN@CNS.NYU.EDU  
WEIJIMA@NYU.EDU

Neural population coding has traditionally been studied for the encoding of a single stimulus, albeit a multidimensional one. However, in natural vision, stimuli rarely appear in isolation. Rather, they generally appear within the context of other stimuli. It is thus of interest to understand how neural populations can simultaneously represent multiple stimuli. Here, we extend the neural population coding framework to the simultaneous encoding of multiple stimuli. We incorporate two biologically motivated properties in our encoding model: (i) the mean response of a neuron to the simultaneous presentation of multiple stimuli is assumed to be a weighted average of its mean responses to the individual stimuli; (ii) noise correlations between neurons are assumed to decay exponentially as a function of the absolute difference between their preferred stimuli (limited-range correlations). We study the asymptotic behavior of the maximum-likelihood estimator (Fisher information matrix), as well as the behavior of

the optimal linear estimator reading out the neural responses to simultaneously estimate multiple stimuli. We report four novel results: (i) stimulus-dependent correlations between estimates of different stimuli are predicted, the precise form of which depends on the parameters of the encoding model and on the decoding method; (ii) encoding accuracy is also found to be stimulus-dependent in a way that depends on the parameters of the encoding model and the decoding method; (iii) making the mean responses of neurons depend on multiple stimuli (stimulus mixing) leads to substantial reductions in encoding accuracy; (iv) under conditions of strong stimulus mixing, increasing the noise correlations between neurons significantly improves encoding accuracy even when the tuning properties of the neurons are homogeneous and even when a suboptimal decoding method is used, thus broadening the range of conditions under which noise correlations are beneficial for encoding accuracy.

## **II-27. How much signal is there in the noise?**

Leon Gatys<sup>1,2</sup>  
Alexander S Ecker<sup>1</sup>  
Tatjana Tchumatchenko<sup>3</sup>  
Matthias Bethge<sup>1,4</sup>

LEON.GATYS@BETHGELAB.ORG  
ALEXANDER.ECKER@UNI-TUEBINGEN.DE  
TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE  
MATTHIAS@BETHGELAB.ORG

<sup>1</sup>University of Tuebingen

<sup>2</sup>Bernstein Center for Computational Neuroscience, Tuebingen

<sup>3</sup>Max Planck Institute for Brain Research

<sup>4</sup>Bernstein Center for Computational Neuroscience

In 1995, Shadlen and Newsome raised the question 'Is there a signal in the noise?' when considering the high input regime in which cortical neurons operate. They promoted the idea that excitation and inhibition are tightly balanced, so that changes in upstream firing rate are encoded primarily in the variance of the postsynaptic current (PSC) distribution while leaving its mean mostly unaffected. Although there is a growing body of experimental evidence supporting the balanced state hypothesis (e.g. Okun and Lampl 2008), relatively little work has been carried out to study the information-theoretic properties of such a neural information channel. Here, we take an information-theoretic approach to develop a general understanding of the encoding process and perform a systematic analysis to estimate the channel capacity. In our model the PSC takes the form of a scale mixture distribution and the upstream population firing rates are encoded in the scale variable of this mixture distribution. As a consequence, the resulting information channel does not allow for the conventional distinction between signal and noise, as it is known for an additive noise channel. We estimate the channel capacity of a balanced population of 10,000 neurons to be 1.15 bits/ms. The encoding performance scales logarithmically with population size, as it is the case for any intensity code. We find, however, that in the balanced state channel different sources of variability can lead to surprising and somewhat unintuitive effects: In particular, additional noise resembling the unreliable synaptic transmission of spikes robustly enhances the channel capacity by up to 10% for realistic population sizes. The underlying mechanism suggests a new beneficial effect of noise for signal transmission that is genuinely different from stochastic resonance.

## **II-28. Neuronal population decoding can account for perceptual lightness illusions**

Douglas A Ruff<sup>1</sup>  
David H H Brainard<sup>2</sup>  
Marlene Cohen<sup>1</sup>

RUFFD@PITT.EDU  
BRAINARD@PSYCH.UPENN.EDU  
COHENM@PITT.EDU

<sup>1</sup>University of Pittsburgh

<sup>2</sup>University of Pennsylvania

The relationship between the intensity of the light reflected from an achromatic object and how light the object



appears depends on visual context. It has been difficult to relate this complexity to the activity of individual neurons, because neurons respond in varied ways to stimulus intensity. Here we ask whether a population decoding approach clarifies the neural underpinnings of perceived lightness. We used stimuli derived from Adelson's checker-shadow illusion, such that probe disks presented on checkerboard images were seen to lie either within a shadowed region ('shadow' condition) or within a luminance-matched region without a shadow ('no shadow' condition). We evaluated whether similar changes in the lightness of the probe disk across the two conditions were revealed by human psychophysics and by lightness estimates obtained through decoding the responses of several dozen neurons in either area V4 or V1 of rhesus monkeys. Our psychophysical experiments confirmed and quantified the checker-shadow illusion for our stimuli: probe disks with the same intensity were perceived to be lighter in the 'shadow' condition than in the 'no shadow' condition. When we decoded the intensity of the probe disks using population responses of V4 neurons to the same stimuli, we found that the decoded intensities for the 'shadow' condition were consistently higher than those for the 'no shadow' condition. Moreover, there was a quantitative match between the psychophysical and neural effects. In contrast, although disk intensity could be decoded from the population of V1 neurons as accurately as from the V4 neurons, the V1 decoding differences were not in agreement with the perceptual illusion. This result suggests that the checker-shadow illusion arises at least in part from cortical computations. More generally, our data support the notion that decoding the responses of neural populations can shed light on the neural correlates of complex perceptual phenomena.

## II-29. Adaptation of intrinsic circuit correlations and synergy in retinal synchrony

Ann Hermundstad<sup>1</sup>

Olivier Marre<sup>2</sup>

Stephanie E Palmer<sup>3</sup>

Michael Berry<sup>4</sup>

Gaspar Tkacik<sup>5</sup>

Vijay Balasubramanian<sup>1</sup>

ANNHERM@PHYSICS.UPENN.EDU

OLIVIER.MARRE@GMAIL.COM

SEPALMER@UCHICAGO.EDU

BERRY@PRINCETON.EDU

GASPER.TKACIK@IST.AC.AT

VIJAY@PHYSICS.UPENN.EDU

<sup>1</sup>University of Pennsylvania

<sup>2</sup>Institut de la Vision - INSERM

<sup>3</sup>University of Chicago

<sup>4</sup>Princeton University

<sup>5</sup>Institute of Science and Technology, Austria

While adaptation has been documented at the single-neuron level, less is known about adaptation in networks. Recent work predicts that network-level adaptation should occur in response to stimulus variations. We test this prediction using simultaneous recordings from large populations of retinal ganglion cells and a new extension to the maximum entropy framework, called time-dependent maximum entropy (TDME). With TDME, the probability distribution of population activity patterns is specified by the neuron's time-dependent propensity to fire, and a set of pairwise couplings that reflect the intrinsic tendency of neurons to emphasize or avoid coincident firing. In addition to reproducing the distributions of firing patterns observed in data, TDME provides a generative model to study the occurrence and structure of rare firing events. Unlike traditional static maximum entropy models, intrinsic circuit couplings in TDME have a clear interpretation: they are the probabilistic (generative) counterpart of so-called "noise correlations". To extract the contribution of these intrinsic circuit couplings, we construct two TDME models of population activity: one in which circuit correlations are preserved, and another in which circuit correlations are destroyed by shuffling stimulus presentations. We use these techniques to explore how intrinsic circuit couplings shape population activity in response to switches between different naturalistic stimuli in the same retina. We find that intrinsic circuit correlations change during natural stimulus presentations to: (a) increase synergistic information conveyed by synchronous states when many neurons fire simultaneously, (b) increase the average information per firing event, and (c) convey more information per firing event than random couplings constrained to reproduce single-cell firing rates. Together, these results suggest that adaptation occurs not only on the single-neuron level, but also on the population level, to adapt the vocabulary of activity patterns to stimulus variations.

## II-30. A data-driven model-free population receptive field estimation framework

Francesca Pizzorni-Ferrarese  
Jonas Larsson  
University of London

FRANCESCA.PIZZORNIFERRARESE@RHUL.AC.UK  
JONAS.LARSSON@RHUL.AC.UK

Most studies characterizing human visual area functions by functional magnetic resonance imaging (fMRI) use phase encoded retinotopic mapping techniques to identify visual area boundaries (Engel et al., 1994; Sereno et al., 1996; Larsson & Heeger, 2006). Because these methods ignore much of the information in the time-series, they have increasingly been replaced by methods to model population receptive fields (pRF) from fMRI data (Larsson & Heeger, 2006; Dumoulin & Wandell, 2008; Zuiderbaan et al., 2012; Lee et al., 2013). pRF modelling techniques involve fitting a Gaussian (or Difference of Gaussians) model of the pRF to individual voxel time series, allowing reconstructing a visual field map and yielding estimates of average receptive field size and suppressive surrounds. However, a limitation of pRF methods is that they require specific assumptions about the pRF structure. To overcome this limitation, we propose a new model-free data-driven approach that estimates the structure of the pRF and the hemodynamic response function (HRF). By making no a priori assumptions about the specific pRF shape our method provides a way to map pRF structure at different spatial locations in an unbiased way. The method is based on the idea of reverse correlation (Ringach & Shapley, 2004) applied to BOLD signals. A limitation of reverse correlation is that it ordinarily requires large numbers of spatially independent stimuli to avoid biased estimates. By making specific assumptions about the temporal structure of the fMRI response, our approach relaxes these requirements, providing rapid and reliable estimates of pRFs with fewer stimuli than conventional reverse correlation using standard retinotopic stimuli. We demonstrate the application of the method to investigate the effects of contrast and visual field position on surround suppression in human visual cortex.

## II-31. Robustness of criticality under different stimulus conditions

Mark Ioffe<sup>1</sup>  
Gasper Tkacik<sup>2</sup>  
William Bialek<sup>1</sup>  
Michael Berry<sup>1</sup>

MIOFFE@PRINCETON.EDU  
GASPER.TKACIK@IST.AC.AT  
WBIALEK@PRINCETON.EDU  
BERRY@PRINCETON.EDU

<sup>1</sup>Princeton University

<sup>2</sup>Institute of Science and Technology, Austria

Previous results have suggested that the maximum entropy models inferred from neural data are tuned to a critical point [Tkacik 2008, Mora 2011]. Specifically, a maximum entropy model for the codewords that incorporates constraints on mean firing rates and pairwise correlations is mathematically identical to the Ising model. This model is one of the simplest models that can have a phase transition, marked by the divergence of the heat capacity. By first fitting an Ising-like model to the data, and then introducing into it a parameter equivalent to the temperature, one can explore whether the retinal codeword distributions are tuned close to the critical point. Being close to a critical point maximizes the dynamic range of codewords, and such an anomalously high dynamic range has been reported in [Tkacik 2013]. We tested these ideas with a 252-channel MEA [Marre 2012] that allowed us to record from 128 retinal ganglion cells. We presented the same retina with the same natural movie, displayed at two widely differing mean luminance levels. We found significant changes in average firing rates, pairwise correlations between neurons, and the distribution of codewords between the two conditions, demonstrating significant changes in the microscopic details of the distributions. Nonetheless, the heat capacity per neuron showed a peak as a function of effective temperature that grew and shifted toward the real operating point of the network ( $T=1$ ) with increasing network size (Fig 1). While these heat capacity curves differed between the two conditions for small subnetworks, there was a visible convergence with system size, suggesting a robustness in this aspect of neural coding. In sum, we find that the neural code retains key signatures of a critical state despite dramatic changes in its microscopic details.

## II-32. Mapping large scale retinal population activity with high density multi-electrode arrays

Sahar Pirmoradian<sup>1</sup>

Gerrit Hilgen<sup>2</sup>

Oliver Muthman<sup>3</sup>

Alessandro Maccione<sup>4</sup>

Upinder Bhalla<sup>5</sup>

Luca Berdonini<sup>4</sup>

Evelyne Sernagor<sup>6</sup>

Matthias H Hennig<sup>1</sup>

SPIRMORA@INF.ED.AC.UK

GERRIT.HILGEN@NCL.AC.UK

JENSOLIVER@NCBS.RES.IN

ALESSANDRO.MACCIONE@IIT.IT

BHALLA@NCBS.RES.IN

LUCA.BERDONINI@IIT.IT

EVELYNE.SERNAGOR@NCL.AC.UK

M.HENNIG@ED.AC.UK

<sup>1</sup>University of Edinburgh

<sup>2</sup>University of Newcastle upon Tyne

<sup>3</sup>Tata Institute of Fundamental Research

<sup>4</sup>Italian Institute of Technology

<sup>5</sup>National Centre for Biological Sciences

<sup>6</sup>Newcastle University

Retinal ganglion cells can be categorized according to a range of physiological criteria, some of which correspond to distinct anatomical features such as dendritic stratification depth in the inner plexiform layer. So far, however, this classification largely relied on pooling data from multiple preparations due to limitations in the number of neurons that could be simultaneously recorded. Here, we present a novel approach for localizing ganglion cells recorded in the adult mouse retina during light stimulation with high density 4096 channel multielectrode arrays. Recording channels were arranged on a 64x64 lattice and separated by 42  $\mu$ m, enabling simultaneous sampling of activity at near cellular resolution of >500 units on a large patch of the retina. Typically, signals from the same cell were detectable on multiple neighboring channels. We exploited this to substantially improve the detection and isolation of spike events, to estimate the current source location, and to cluster spiking events according to spatial proximity, thus identifying the activity of single units. A comparison with standard spike sorting techniques showed this method could substantially improve spike detection and correctly cluster spikes that would otherwise be assigned to different sources. An analysis of responses to full field flashes (a sequence of dark/light stimuli) led to the following findings: response latencies were broadly distributed from 50-950 ms, with faster responses in OFF cells; distribution of the ratio of ON versus OFF responses were broad and bi-modal, without clear peak at equal ratios (ON-OFF cells) and higher abundance of OFF cells compared to ON cells; the spatial distribution of ON and OFF cells in the dorsal retina was random. In summary, our results are a first step towards a comprehensive and statistically sound characterization of the activity of a large population of neurons without variability caused by combining data from multiple preparations.

## II-33. Latent variable models (almost) always display signatures of criticality

Laurence Aitchison<sup>1</sup>

Nicola Corradi<sup>2</sup>

Peter Latham<sup>1</sup>

LAURENCE@GATSBY.UCL.AC.UK

NICOLA.CORRADI@GMAIL.COM

PEL@GATSBY.UCL.AC.UK

<sup>1</sup>Gatsby Computational Neuroscience Unit, UCL

<sup>2</sup>Cornell University

There is a large body of experimental work suggesting that the brain is, in some sense, critical - a notion from statistical physics implying that the system is fine-tuned to a point with remarkable properties. Observations of criticality in the brain are thought to point to a deep, underlying organizational principle. However, here we show that some of these observations arise trivially from simple, well-understood processes, so cannot tell us anything new about the brain. We show that two classic signatures of criticality can arise from any process that involves time-varying latent variables (e.g. firing rates) that can be thought of as generating observations (e.g. spike times).

In particular, we look at two properties of the distribution over energies, where the energy of an observation is simply its negative log-probability. The first property is that the variance of the distribution over energies is large for systems at the critical point. Secondly, the distribution over energies is flat at the critical point, which also gives rise to Zipf's law. Furthermore, one may want to ask whether a proposed latent variable can account for some, all or none of the observed criticality. To address this question, we developed a measure, the percentage of criticality explained. Applying this measure to retinal data told us that 94% of the criticality in our spiking dataset could be explained simply by the underlying firing rate. We believe that many other observations of criticality arise from well-understood processes. It is important to discover which these are, so that we can focus our attention on only those signatures of criticality that tell us something non-trivial about computation in the brain.

## II-34. Generation of sequences through reconfiguration of ongoing activity in neural networks: A model of c

Kanaka Rajan<sup>1</sup>  
Christopher Harvey<sup>2</sup>  
David Tank<sup>1</sup>

KRAJAN@PRINCETON.EDU  
CHRISTOPHER\_HARVEY@HMS.HARVARD.EDU  
DWTANK@EXCHANGE.PRINCETON.EDU

<sup>1</sup>Princeton University  
<sup>2</sup>Harvard Medical School

Complex timing tasks are the basis for experiments identifying the neural correlates of behaviors like memory-based decision making in brain areas like the posterior parietal cortex (PPC). Recently, cellular-resolution imaging of neural activity in PPC during a virtual memory-guided 2-alternative forced choice task [Harvey, Coen & Tank, 2012] showed that individual neurons had transient activation staggered relative to one another in time, forming a sequence spanning the entire duration of the task. Motivated by these results, our goal here is to develop a computational framework that reconciles the emergence of biologically realistic assemblies or trajectories of activity states, with the ability of the same neural population to translate sensory information into long time-scale behaviors. We build an echo state network to test our hypothesis that during memory-based decision making, sensory cues set up an initial network state that follows the intrinsic dynamics of the brain area to generate activity underlying a behavioral response. We start with a firing rate network which exhibits rich ongoing dynamics correlated over numerous time scales. This network acts as a dynamic reservoir whose modes can be tapped to perform the task through minimal reconfiguration or partial in-network (PIN) training, and not complete rewiring. 1) Only weights carrying the inputs to a subset of neurons are subject to change. 2) There is no external unit that feeds back network output, nor learning of readout weights. 3) The learning rule targets as many units as required for the task, the fraction varying with task demands. We show that the PINned network performs a timing task involving a sensory cue, its storage in working memory during a delay period, and response to its retrieved trace. We change the fraction of trained units until the duration and shape of the sequential activation pattern in the network is comparable with PPC neurons imaged during the task. Further, we show that like the PPC, the network's activity is specific to cue/outcome pairing, temporally confined to a task epoch, and contains similar levels of extra-sequential noise. Notably, since no sequence-specific wiring diagram is embedded a priori units remain spatially intermixed, like the PPC where there is no topological organization of active neurons. Finally, we study the properties and functional consequences of the synaptic connectivity matrix which is initially random but acquires non-normal features because of PINning. We are currently exploring whether such partially trained networks can be extended to simulate our ability to generalize across different task conditions.

## II-35. A method for fitting linear neural models to extracellular electrophysiology data

Corinne Teeter  
Christof Koch  
Stefan Mihalas

CTEETER@GMAIL.COM  
CHRISTOFK@ALLENINSTITUTE.ORG  
STEFANM@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

With the recent explosion in computer power and experimental techniques to collect data, it is now possible to begin to make biologically based, large scale network models in order to make predictions, guide experimental investigations, and study perturbative effects of neural systems. Realistically behaving, fast, neuron models are needed to populate biologically realistic network models. Investigators have been quite successful quickly fitting linear models to subthreshold electrophysiology data (Mensi et al. 2012). However, subthreshold data is not available in the vast majority of the experimental paradigms that would be most useful in the generation of large neural network models including experiments recording large neural populations or in vivo experiments. Maximum likelihood estimation (MLE) can be utilized to fit simple linear models using solely the spike-times available in extracellular data (Paninski 2004). However, MLE is computationally intensive and is impractical when variables that are needed to add more biological realism, such as a voltage dependent threshold, are built into the model. (Dong, et al 2011). Here we present a method of estimating the parameters of a generalized linear integrate and fire neuron (GLIF) (Mihalas & Niebur, 2009) that requires only knowledge of the input and the timing of spikes. Using somatic electrophysiological patch clamp pipeline data from mouse visual cortex, we show that GLIF models with several after spike currents, and a voltage adaptive threshold variable are comparable to current subthreshold models in predicting the spiking behavior of neurons. A large database of GLIF models will be constructed and made publically available.

## II-36. Direct and indirect pathways of the basal ganglia exert opposing but not identical control of motor

Ian Oldenburg  
Bernardo Sabatini  
Harvard University

IAN.OLDENBURG@GMAIL.COM  
BSABATINI@HMS.HARVARD.EDU

The Basal Ganglia are an interconnected group of nuclei that are critical to the control of voluntary movements. Anatomically the Basal Ganglia are arranged in a loop, from cortex to Basal Ganglia to thalamus and back to cortex. Within the Basal Ganglia there are two major pathways, the Direct and Indirect pathways, with opposite effects on movements. Despite anatomical and behavioral evidence, functional evidence of a loop from basal ganglia to cortex does not exist. Indeed, such a study has been very difficult with classical approaches due to the intermingled distribution of direct and indirect pathway striatal neurons, and the extensive interconnections of the basal ganglia. Using channelrhodopsin in either pathway we show strong bidirectional modulation of cortical activity in the awake behaving animal during quiescent periods, but only inhibition via the indirect pathway during movements and cues. Furthermore the relationship between pathway modulation and cortical unity identity is different between the two pathways, suggesting different mechanisms or targets of Basal Ganglio-thalamo-cortical projections. Thus, we show that the canonical looped architecture is only playing a dominant role while the animal is not moving, with each pathway differentially affecting the cortex.

## II-37. Spatial learning by *Drosophila melanogaster* in virtual reality

TJ Florence<sup>1,2</sup>  
 Charles Zuker<sup>1</sup>  
 Michael Reiser<sup>3,4</sup>

FLORENCET@JANELIA.HHMI.ORG  
 CZ2195@COLUMBIA.EDU  
 REISERM@JANELIA.HHMI.ORG

<sup>1</sup>Columbia University

<sup>2</sup>Janelia Farm Visitor Program

<sup>3</sup>Janelia Farm Research Campus

<sup>4</sup>Howard Hughes Medical Institute

Recently we showed that *Drosophila* excel at visual place learning, using a paradigm analogous to the Morris water maze but with heat in place of water. This demonstration enables the use of the powerful *Drosophila* genetic toolkit towards the study of long-standing questions in insect navigation. Experiments on tethered, behaving insects have provided insight into odor tracking, motion vision, and feature recognition, among other stimulus-evoked behaviors. Taking inspiration from decades of previous efforts, as well as recent work in mammals, we have developed a virtual-reality spatial learning task for tethered walking flies. We simulate a thermal Y-maze. Warm (35°C) lanes encourage the fly to walk through the maze, while an aversive (40°C) surround minimizes walking outside the maze. One maze arm ends in a 'rewarding' cool zone, while the other ends in an aversive hot zone. Arms are associated with distinguishing visual patterns. We challenge each fly to learn which arm of the maze leads to the cool zone. Our standard experiment consists of ten training trials and four test trials (repeated twice). Test trials restart as soon as an animal 'chooses' a maze arm. Over the course of the experiment, flies improve in locating the cool zone. After training, flies demonstrate a significant preference for the rewarded side. Control flies trained to a condition in which the cool zone location is randomized, while the visual scene remains constant, develop no side preference. In ongoing experiments, we are using this platform to investigate: 1) the consequence of targeted perturbations of the sensory experience during virtual navigation, and 2) the neural correlates of visual place learning, using calcium imaging to optically record from targeted populations of neurons in the fly brain during behavior.

## II-38. Redundant hippocampal spatial coding offsets competition between interference and generalization

Alexander Keinath<sup>1</sup>  
 Joshua T Dudman<sup>2,3</sup>  
 Isabel Muzzio<sup>1</sup>

ATKEINATH@GMAIL.COM  
 DUDMANJ@JANELIA.HHMI.ORG  
 IMUZZIO@SAS.UPENN.EDU

<sup>1</sup>University of Pennsylvania

<sup>2</sup>Janelia Farm Research Campus

<sup>3</sup>Howard Hughes Medical Institute

Recent work highlighted neuroanatomical and neurochemical differences along the longitudinal hippocampal axis (dorsal-ventral in rodents, anterior-posterior in humans). These differences are paralleled by a gradient in the scale of spatial representations. At the single cell level, place field size of pyramidal cells increases when progressing from dorsal to ventral regions. At the population level, activity in the dorsal region becomes decorrelated over short distances, while activity in ventral regions stays correlated over larger distances, echoing single cell spatial scaling. This has been taken to suggest that space is represented at a fine resolution in the dorsal hippocampus, but this representation deteriorates in more ventral regions. However, a shift in spatial scale does not necessarily imply a loss of spatial information. Rather, if the fidelity of spatial information is preserved along the longitudinal axis, this gradient may instead signal a shift from sparse to distributed coding. To test this possibility, we recorded from dorsal or ventral CA1 in freely moving mice exposed to multiple visuospatial and olfactory cues. Using techniques derived from PCA, as well as linear and nonlinear location reconstruction methods, we found that, despite the broad spatial tuning of ventral single cells, populations of equal size in both regions faithfully represented space regardless of condition. To explain this apparent redundancy in spatial coding, we modeled

via neural networks the impact that a representational gradient would have on memory. Our results indicate that, among other advantages, such a gradient offsets the interference-generalization tradeoff inherent in network memory. Specifically, dorsal representations robustly guard detailed information against interference, at the expense of generalization of common features across contexts. Conversely, ventral representations successfully extract and generalize commonalities, at the cost of high interference for details.

## II-39. Phase comparison of functionally distinct inputs in CA1 pyramidal cells during locomotion

Katie Bittner<sup>1,2</sup>  
Jeffrey Magee<sup>1,2</sup>

BITTNERK@JANELIA.HHMI.ORG  
MAGEEJ@JANELIA.HHMI.ORG

<sup>1</sup>Janelia Farm Research Campus  
<sup>2</sup>Howard Hughes Medical Institute

Most neuronal circuits receive at least two functionally distinct forms of input (intrinsic vs. extrinsic; sensory vs. motor; etc). In many pyramidal neuron based microcircuits integration of these two input signals can proceed nonlinearly through the production of active dendritic voltage signals. For example, appropriately timed, perisomatically located, hippocampal (SC) and distal dendrite targeting entorhinal (EC3) input produces a distal dendritic Ca<sup>2+</sup> plateau potential that drives burst firing output from CA1 pyramidal neurons in vitro. Related signals have been observed in neocortical pyramidal neurons. Until recently it was unknown whether these events occurred in vivo and, if so, during what behavioral states. Here we used simultaneous whole-cell patch and field potential recordings in head-fixed mice running on a linear track treadmill to study this dendritic plateau driven burst firing (plateaus) in CA1 neurons. Here we find that during locomotion, plateaus occur within the neuron's place field with initiation probability peaking near the peak of the firing field. Plateaus produce a large (32.4mV; n=12), slow (duration; 51.7ms) somatic depolarization that appears similar to that measured in vitro. Interestingly, plateaus exhibit a dramatic level of theta-phase modulation (~97%) that peaks late in the theta cycle (~350°). This late phase peak in plateau potential initiation is near the theta-phase preference of EC3 inputs, suggesting a theta-phase dependent interaction of SC and EC3 inputs. We tested this idea by manipulating the phase of SC inputs by injecting phase adjusted theta frequency currents into CA1 somas. Biasing AP firing earlier in the theta phase decreased plateau probability to ~48% of control whereas biasing AP firing later in phase increased plateau probability 276%. These data indicate that CA1 pyramidal cells compare theta-phase relations of SC and EC3 inputs with backpropagating action potentials acting as a link between the two isolated subcellular input integration compartments.

## II-40. Representation and generalization in Drosophila olfactory learning

Ann Kennedy  
Peter Wang  
Daisuke Hattori  
Richard Axel  
L.F. Abbott

AK3024@COLUMBIA.EDU  
PETERWANG724@GMAIL.COM  
DH2453@COLUMBIA.EDU  
RA27@COLUMBIA.EDU  
LFA2103@COLUMBIA.EDU

Columbia University

In the *Drosophila* olfactory system, the mushroom body translates sensory representations of odors into learned behavioral responses [1]. Sensory input from the olfactory periphery drives sparse patterns of activation among Kenyon cells, the intrinsic cells of the mushroom body. Kenyon cell activity is read out by a small set of output neurons, which project to reward and behavior-associated regions of the brain and are broadly tuned to odor identity in the naive fly. Pairing of an odor with reward or shock produces an odor-specific modification of the output neuron's response to the trained odor; this change in output neuron response is required for subsequent recall of the conditioned behavior [2-4]. This suggests that environmentally-modulated plasticity at the Kenyon cell-output

neuron synapse drives learning of novel responses to trained odors [5]. Kenyon cell have several properties that are favorable for associative odor learning, such as sparse [6, 7] heterogeneous [8] representations and normalization of odor-evoked responses [9]. We examine associative learning in a spiking model of the mushroom body built according to the known organization of the fly olfactory system. We find that due to overlapping Kenyon cell representations, output neurons are susceptible to over-generalization when trained to respond to multiple odors. We devise two forms of synaptic plasticity that overcome this problem by driving output neuron responses to separate fixed points for trained and untrained odors. These distinct plasticity mechanisms lead to different experimentally testable predictions. To improve the robustness of the system to overlapping odor representations, we also reconsidered the role of APL [10], a large inhibitory interneuron hypothesized to improve odor separability by providing recurrent inhibitory normalization to the Kenyon cell population. We construct a model of APL recurrence that fits experimentally-observed patterns of Kenyon cell recruitment, including increased Kenyon cell activation at higher odor concentrations and sublinear responses to mixtures of odors. Using this revised model, we study generalization and extrapolation of associative learning across concentrations of a learned odor and in odor mixtures.

## II-41. The functional role of randomness in olfactory processing

Kamesh Krishnamurthy<sup>1</sup>  
 Ann Hermundstad<sup>1</sup>  
 Thierry Mora<sup>2</sup>  
 Aleksandra Walczak<sup>2</sup>  
 Venkatesh N Murthy<sup>3</sup>  
 Charles Stevens<sup>4</sup>  
 Vijay Balasubramanian<sup>1</sup>

KAMESH@MAIL.MED.UPENN.EDU  
 ANNHERM@PHYSICS.UPENN.EDU  
 TMORA@LPS.ENS.FR  
 AWALCZAK@LPT.ENS.FR  
 VNMURTHY@MCB.HARVARD.EDU  
 STEVENS@SALK.EDU  
 VIJAY@PHYSICS.UPENN.EDU

<sup>1</sup>University of Pennsylvania

<sup>2</sup>Ecole Normale Supérieure

<sup>3</sup>Harvard University

<sup>4</sup>Salk Institute for Biological Studies

Recent experimental studies in the insect and mammalian olfactory systems observe diffuse and apparently random projections between the second stage (antennal lobe/olfactory bulb) and third stage (mushroom body/piriform cortex) of processing. What functional role does randomness serve? We start with two key characteristics of olfactory stimuli: (a) typical odors contain small numbers of the possible volatile molecules, and (b) organisms continually encounter relevant novel odorants. Results from compressive sensing prescribe random measurements as an optimal way of representing such signals. Thus we propose that the olfactory system exploits the inherent sparseness of odors to compactly represent both new and familiar odors by increasingly randomizing the sensory signal at each processing stage. This hypothesis also predicts that the responses of small, random subsets of third order olfactory neurons will be sufficient to robustly decode odors. Olfactory Receptor Neurons are unlikely to provide random measurements of odors because the bindings of odorants to receptors are statistically dependent. Despite these correlations, which we characterize in *Drosophila*, the responses are sufficiently unstructured to permit decoding of odor mixtures from random ORN subsets. A computational model then shows that inhibitory nonlinearities at the second stage decorrelate ORNs — allowing the decoding of larger mixtures from smaller subsets of neurons. Finally, random projections to the third stage, combined with nonlinearities, increase the separation between all pairs of odorants and produce larger separations than those produced by structured connectivity. As predicted, small subsets of neurons in the third stage are sufficient to achieve the same separation performance as the entire population. Ongoing work seeks to characterise experimentally-observed responses to mixtures and sequential presentation of odors, using the computational model.



## II-42. Two types of representation of sensory information in cortical feedback axons in awake mice

Gonzalo Otazu  
Hong Goo Chae  
Dinu Albeanu

OTAZU@CSHL.EDU  
HGCHAE@CSHL.EDU  
ALBEANU@CSHL.EDU

Cold Spring Harbor Laboratory

Odorant information is transmitted from the olfactory epithelium by way of the olfactory bulb into the piriform cortex. Similar to other sensory modalities, the bulb also receives massive amounts of cortical feedback whose role is still unknown. Computational models of sensory processing have speculated that cortical feedback conveys prior information that is used to disambiguate the incoming sensory input. However, there are no physiological data on the activity of feedback axons in awake animals. As a first step in understanding the role of piriform cortex feedback in olfactory processing, we have directly imaged the activity of corticobulbar axons using two photon calcium imaging in awake mice. We found that cortical feedback boutons provided a continuous level of activation even in the absence of odor input. Odor presentation enhanced activity as well as suppressed spontaneous activity with roughly equal numbers of enhanced and suppressed responses. Notably, we found that cortical feedback axons could be classified according to the type of response to odors. One set of axons responded to odors by enhancing their activity, while a different set of axons responded by suppressing its spontaneous activity, with only a small fraction of boutons (<5%) showing enhancement for some odors and suppression for other odors. Furthermore, enhanced and suppressed axons to a particular odor were spatially clustered in spatial domains. Interestingly, the majority of odor evoked responses lasted for several seconds following odor termination, indicating that feedback activity does not only reflect the current odor environment but also carry information about the recent past. We are currently investigating the feedback response dynamics during odor presentation in more naturalistic conditions (i.e using odor pulses). Furthermore, we are probing circuit substrates underlying the distinct enhanced vs. suppressed responses, and long lasting vs. transient feedback responses by specific manipulations of activity in the piriform cortex.

## II-43. Long-distance inhibition by short axon cells gate the output of the olfactory bulb.

Arkarup Bandyopadhyay  
Fred Marbach  
Francesca Anselmi  
Matthew Koh  
Martin B. Davis  
Hassana K. Oyibo  
Priyanka Gupta  
Dinu Albeanu

ABANDYOP@CSHL.EDU  
FREDIMARBACH@GMAIL.COM  
FANSELM@CSHL.EDU  
MKOH@CSHL.EDU  
DAVISM@CSHL.EDU  
OYIBO@CSHL.EDU  
PGUPTA@CSHL.EDU  
ALBEANU@CSHL.EDU

Cold Spring Harbor Laboratory

Odors activate a distributed set of glomeruli in the input layer of the rodent olfactory bulb (OB). Lateral interactions in the OB sculpt this incoming information, before being routed to the cortex by the Mitral/Tufted (M/T) cells. Surprisingly little is known about how lateral interactions within the OB shape the output of individual M/T cells *in vivo*. Most studies have focused either on computation within one glomerulus or on the granule cells in the output layer. Here we describe long-range lateral cross talk in the glomerular layer *in vivo*, demonstrating that substantial computation happens even before the odor information reaches the M/T cells. Short Axon (SA) cells, a genetically identifiable population of interneurons in the OB, are thought to mediate interglomerular inhibition largely based on anatomy and some recent direct evidence in OB slices. We set out to directly assess the influence of interglomerular cross talk mediated by SA cells on the bulb output *in vivo*. We find that SA cells carry odor concentration signals and broadcasts it throughout the glomerular layer. Optogenetic activation of SA

cells strongly inhibits both spontaneous and odor-evoked firing of M/T cells, even at large distances from the recording site. These observations suggest that SA cells mediate inhibitory gain-control in the glomerular layer. Mechanistically, SA cells directly inhibit the External Tufted (ET) cells, which thereby reduces the excitatory drive of M/T cells. Contrary to the conventional model of glomerular connectivity, ET cells have been suggested to relay odor inputs to M/T cells from the OSNs. We show that indirectly inhibiting the ET cells drastically reduces the responsiveness of M/T cells, providing *in vivo* evidence for this model. In summary, SA cells acting via ET cells are crucial in determining the firing rate of M/T cells *in vivo*, thereby gating the OB output.

## **II-44. Learning hierarchical representations of natural auditory scenes**

Wiktor Młynarski

MLYNAR@MIS.MPG.DE

Max Planck Institute for Mathematics in the Sciences

The biological auditory system is capable of identifying objects in the environment and reconstructing their spatial configuration. This complex inference task is performed upon two one-dimensional time series of air-pressure entering each ear. Current theories suggest that auditory neurons achieve it by responding to coherent streams i.e. patterns revealing strong correlations across frequency, pitch, or location channels. Such streams hypothetically correspond to separate sound sources. It is however not known, how this computation can be performed in a natural auditory environment and which coherence patterns can be observed in ecological conditions. This work approaches the auditory scene analysis problem by proposing a hierarchical, generative model of binaural sounds recorded in a fully natural environment. The model consists of three layers. To separate temporal information (phase) from amplitude, a sparse coding of sound epochs from left and right ear (first layer) is formed using complex valued basis functions in the second layer. Representation of phase as a temporal shift requires placing continuity priors on learned basis functions, since the unconstrained model captures cross-frequency and bandwidth invariants. The third layer explains intra- and inter-aural correlations of phases and amplitudes. The fundamental assumption made, is that an auditory scene at a short time instant can be explained as a combination of only a few sparse causes. High level structures captured by the model represent spectral composition of sound sources and their spatial relationships. They form a more invariant representation of "what" and "where" aspects of the auditory scene than second layer coefficients. Obtained results suggest which kinds of auditory streams useful in source separation can be extracted from a natural binaural input. Additionally they allow one to form experimental predictions of how spatial sensitivity of auditory neurons is formed by combining phase and amplitude cues.

## **II-45. Invariant selectivity of auditory neurons due to predictive coding**

Izzet Yildiz<sup>1</sup>

Brian Fischer<sup>2</sup>

Sophie Deneve<sup>1</sup>

IZZETBURAKYILDIZ@GMAIL.COM

FISCHER9@SEATTLEU.EDU

SOPHIE.DENEVE@ENS.FR

<sup>1</sup>Ecole Normale Supérieure

<sup>2</sup>Seattle University

Predictive coding theory proposes to view the brain as an interactive mechanism that constantly adapts and reshapes its responses to better fit to the environment. It has been highly influential for the neuroscience community since it can account for many aspects of neuronal structure and function such as the classical and non-classical receptive field properties and nonlinear responses of sensory neurons in the visual system. However, the consequences of this theory have been rarely exploited in the auditory domain. The goal of this work is to close this gap by building a model of optimal auditory perception, matching it to a spiking neural network structure and exploiting its dynamics to acquire novel interpretations of experimentally observed neuronal and behavioral responses. We propose that auditory neurons are predictors rather than filters of their input and we investigate their "true selec-

tivity" which is invariant under adaptive and contextual modulations in contrast to their spectrotemporal receptive fields that depend on the context of the auditory scene. We train our model on a large database of natural speech to recover this invariant selectivity, i.e. predictive fields, and show that it can account for nonlinear contextual effects such as combination sensitivity, two-tone suppression and forward suppression. Moreover, we show that the model neurons adapt rapidly to new input statistics (such as behaviorally relevant tones) and can explain the experimentally observed task-dependent receptive field changes in the primary auditory cortex neurons. Furthermore, we employ the model to gain novel insights into the underlying neuronal mechanisms of auditory stream segregation, i.e. behaviorally grouping of tones into one or more separate sound sources. Finally, we discuss how the measurable quantities produced by the model can be used to analyze electrophysiological data and predict behavioral and neuronal responses in the auditory system.

## II-46. Bayesian perception of the pitch of non-stationary natural sounds

Vincent Adam  
Maneesh Sahani

VINCENT.ADAM87@GMAIL.COM  
MANEESH@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

Many natural sounds evoke a percept of pitch; but despite being a fundamental perceptual attribute, pitch is difficult to define. The pitch of periodic sounds can generally be related to their fundamental frequencies. However, small deviations from periodicity often do not destroy the initial periodicity-related pitch percept—possibly leading to ambiguous situations in which a given sound might have arisen by deviation from one of many different base sounds that evoke different pitches. Such ambiguity reveals the ill-posed nature of pitch estimation. One principled way to deal with both ambiguity and the stochasticity of sound generation and transduction is to frame pitch perception as inference in a probabilistic model. Previous work has used such an ideal-observer approach to model human pitch estimation, reproducing qualitatively the human percept of pitch for a set of standard psychophysical stimuli [1] and providing a closer account of human pitch perception in an ambiguous setting than alternative models [1][2]. Natural sounds, however, are not stationary—the timbre, amplitude and pitch may all vary continuously—and not all of them evoke a pitch. Thus pitch perception depends first on a determination of whether the sound is pitched at all, and only then on a pitch estimate that must account for the non-stationarity. Here, we extend the previous approach to capture local timbre and amplitude variations in a way that mimics human voiced speech production, and we solve the pitch detection problem by introducing a competing generative model for non-pitched sounds. These extensions lead to a principled approach to pitch detection and pitch-strength estimation that matches human psychophysics. With an additional smooth prior on pitch variation it allows efficient pitch-tracking of speech sounds. [1] P. Hehrmann and M. Sahani, Bayesian Pitch, Cosyne 2010 [2] P. Hehrmann, Pitch Perception as Bayesian Inference, PhD thesis, Gatsby Unit, UCL, 2011

## II-47. Learning noise-invariant auditory receptive fields for speech recognition

Phillip Schafer  
Dezhe Jin

PBS929@GMAIL.COM  
DJINPSU@GMAIL.COM

Pennsylvania State University

Humans and animals routinely communicate in noisy acoustic environments, but the neural mechanisms enabling this ability are largely unknown. One possibility is that the receptive fields of auditory cortical neurons are tuned to produce noise-invariant representations of sound. Here we propose two noise-invariance objectives for learning spectro-temporal receptive fields (STRFs) for neural populations and evaluate their performance on a speech recognition task. We first compute STRFs that maximize the correlation between the firing rate responses to clean and noisy speech. The resulting representations give improved recognition performance in noise over other

methods, but also highlight limitations of the linear STRF method—due to orthogonality constraints, improved noise invariance in some neurons comes at the expense of reduced invariance in others. To bypass this limitation, we use a different, discriminative criterion in which each neuron is trained to identify speech sounds drawn from a brief temporal window. The neural responses are thresholded to produce a sparse spike code that gives further speech recognition performance gains. We compare the STRFs from both training schemes to physiological data and propose a novel decoding paradigm for the spike code that further improves the recognition performance.

## **II-48. Sensitive periods for functional connectivity in auditory cortical circuits**

Liberty Hamilton<sup>1</sup>  
Jascha Sohl-Dickstein<sup>2,3</sup>  
Shaowen Bao<sup>1</sup>

LIBERTYHAMILTON@GMAIL.COM  
JASCHA@STANFORD.EDU  
SBAO@BERKELEY.EDU

<sup>1</sup>University of California, Berkeley

<sup>2</sup>Stanford University

<sup>3</sup>Khan Academy

The auditory cortex undergoes rapid maturation during early postnatal development, as manifested by the emergence and refinement of cortical sound representations. This process is shaped by acoustic experience in a "critical period" of heightened plasticity. Within the critical period, sensory experience in a series of time windows changes the representation of different acoustic features, including characteristic frequency (CF), tuning bandwidth (BW), and frequency modulation (FM) selectivity. The existence of multiple sensitive periods for different acoustic features suggests that the CF, BW, and FM selectivity of cortical neurons are determined by different circuits that undergo experience-dependent modifications in distinct time windows. However, the relative contributions of different cortical microcircuits in mediating single-site changes in these sensitive periods are still poorly understood. To address this, we acquired simultaneous multi-unit recordings in putative layers 2-4 of the mouse primary auditory cortex (AI), anterior auditory field (AAF), and secondary auditory cortex (AII) after pulsed noise exposure during different sensitive windows of the auditory critical period. We then used Ising models to describe functional connectivity within auditory cortical circuits (between pairs of simultaneously recorded sites) as well as functional connectivity between cortical sites and sound stimuli. Noise exposure in early windows reduced stimulus-related functional connections in all auditory fields, with slight differences between AI, AAF, and AII. These changes may be related to alterations in thalamic input. Late noise exposure, in contrast, was associated with increased intrinsic firing rate and reduced spread of functional connectivity between deep and superficial layer sites in AI, indicating that late changes were primarily corticocortical. We compared these functional connectivity results with receptive field response properties to determine how single-site feature selectivity is related to circuit properties. Our results suggest that changes to distinct auditory circuits underlie plastic changes observed at the single site level.

## **II-49. Population coding of complex sounds in the avian auditory system**

Mike Schachter  
Tyler Lee  
Julie Elie  
Friedrich Sommer  
Frederic Theunissen

MIKE.SCHACHTER@GMAIL.COM  
TYLERLEE@BERKELEY.EDU  
JULIE.ELIE@BERKELEY.EDU  
FSOMMER@BERKELEY.EDU  
THEUNISSEN@BERKELEY.EDU

University of California, Berkeley

Auditory neuroscientists have yet to describe how complex sounds are represented in the ensemble response. More generally, the role of correlations, synchrony, oscillations, traveling waves or other properties of population activity in sensory coding remain an open question. Here we attempted to characterize how populations of

neurons throughout the Zebra Finch auditory system interact to process behaviorally relevant communication signals by investigating the pairwise coupling structure amongst local field potentials recorded simultaneously from multiple brain regions. We probed the avian auditory system with acoustic stimuli that span the entire vocalization repertoire of the zebra finch (alarm calls, distance calls, distress call, songs, etc) and also include synthetic noise-like stimuli. To study the population activity, we then developed a novel analysis that involves the estimation of a time varying pair-wise coherence across all simultaneously recorded sites. This approach allows us to investigate both average correlations as well as transient correlations, including those produced by waves of activities or elicited by specific stimuli. On average, we found that, as expected, the pairwise coherence declined quickly as a function of recording distance. We have, however, observed transient stimulus-driven long-range coherences including inter-hemispheric coupling. We also observed particular sites that appear to be sources and sinks of waves of correlated activity. Finally, we have quantified stimulus-dependent deviations from the average coupling, and showed that it depends on the class of vocalization used. These analyses are not only revealing to understand the functional connectivity of the system but also demonstrate the important role that a population code plays in the representation of sensory information.

## II-50. Changes in encoding of communication signals by populations of neurons in the auditory cortex

Isaac Carruthers<sup>1</sup>  
 Ryan Natan<sup>1</sup>  
 Andrew Jaegle<sup>1</sup>  
 Laetitia Mwilambwe-Tshilobo<sup>1</sup>  
 Diego Laplagne<sup>2</sup>  
 Maria Geffen<sup>1</sup>

CIS@SAS.UPENN.EDU  
 RNATAN@GMAIL.COM  
 AJAEGLE@MAIL.MED.UPENN.EDU  
 LMWILAMBWE10@GMAIL.COM  
 DLAPLAGNE@MAIL.ROCKEFELLER.EDU  
 MGEFFEN@MAIL.MED.UPENN.EDU

<sup>1</sup>University of Pennsylvania  
<sup>2</sup>Rockefeller University

Communication signals undergo acoustic transformations when presented in varying environments or when produced by different speakers. In order to create a representation of a complex acoustic signal, the auditory system needs to exhibit invariance to acoustic transformations that do not change the identity of auditory objects. We tested the degree to which populations of neurons in the primary auditory cortex (A1) and a putatively higher auditory area, the supra-rhinal auditory field (SRAF), recorded in awake rats, exhibit invariance to transformations in the basic acoustical features of rat ultra-sonic vocalizations (USVs). We computed the specificity index — a measure of how well a classifier, trained on neuronal responses to original vocalizations, performed when tested on transformed vocalizations. The specificity index was lower for ensembles of neurons in SRAF as compared to ensembles of neurons in A1. This suggests that populations of neurons in SRAF exhibit greater invariance to acoustic transformations than populations of neurons in A1. We further tested whether the population code was transformed between A1 and SRAF. The relative temporal spiking activity within ensembles of neurons is reflected by the so-called 'noise- correlations' — the correlated activity of multiple neurons. In order to evaluate whether noise- correlations improve the quality of the population code for our stimuli, we measured the performance of a model classifier on groups of simultaneously recorded neurons when noise-correlations were left intact, and when noise-correlations were removed. We found that the performance of our classifier was significantly higher with noise-correlations intact, both for groups of neurons recorded in A1 and for groups recorded in SRAF. The relative improvement was similar for both areas. This indicates that temporally correlated neuronal activity within population of neurons carries information about communication signals in addition to that encoded by distinct neurons, in both A1 and SRAF.

## II-51. Clustering of cortical responses to natural sounds reveals organization for pitch and speech

Samuel Norman-Haignere  
Joshua McDermott  
Nancy Kanwisher

SVNH@MIT.EDU  
JHM@MIT.EDU  
NGK@MIT.EDU

Massachusetts Institute of Technology

Compared with visual cortex, the structure of auditory cortex is poorly understood. We applied a hypothesis-neutral, data-driven approach to investigate auditory functional organization in humans. Using fMRI, we measured responses to individual sounds throughout auditory cortex, using a diverse collection of 165 everyday sounds. We then applied a clustering algorithm to the response profiles of each voxel (i.e., the response to each sound). The clustering algorithm fit the collection of response profiles with a small number of exponentially-distributed mixture components, but was otherwise unconstrained — in particular, it had no information about voxel locations, and no prior on response profiles. Once clusters of voxels were determined from their functional profiles, the voxels in each cluster were projected back into the brain to test whether functional clusters were located in consistent anatomical locations. We also examined whether the response profiles of each cluster could be explained by a set of acoustic measures. The analysis revealed 3 clusters of voxels that were consistent across subjects, each of which was located in a distinct anatomical region: 1) an early auditory region near Heschl's Gyrus that responded strongly to all 165 sounds, and whose response was best predicted by a measure of amplitude modulation 2) a region in the anterior superior temporal plane that responded to sounds with a pitch, and whose response was most strongly correlated with a measure of pitch salience and 3) a region in the superior temporal gyrus that was highly selective for speech sounds and whose response profile could not be explained by acoustic measures. These results suggest that pitch and speech are central to the organization of high-level auditory cortex, possibly analogous to the importance of visual motion and faces in the organization of high-level visual cortex.

## II-52. Characterizing central auditory neurons in *Drosophila*

Allison Baker<sup>1</sup>  
Alex Vaughan<sup>2</sup>  
Bruce Baker<sup>3,4</sup>  
Rachel Wilson<sup>1</sup>

ALLISONEVABAKER@GMAIL.COM  
VAUGHANA@JANELIA.HHMI.ORG  
BAKERB@JANELIA.HHMI.ORG  
RACHEL\_WILSON@HMS.HARVARD.EDU

<sup>1</sup>Harvard Medical School

<sup>2</sup>Cold Spring Harbor Laboratory

<sup>3</sup>Janelia Farm Research Campus

<sup>4</sup>Howard Hughes Medical Institute

*Drosophila* use acoustic communication signals during courtship, but little is known about the representation of sound in central auditory pathways. Previous studies suggest that flies process auditory information on different timescales in parallel channels within the central nervous system. We have identified novel central auditory neurons whose processes overlap with the axon terminals of a subset of primary auditory neurons. We believe this population represents the primary projection neurons in one of these parallel processing channels. We are investigating the synaptic inputs and receptive field properties of these neurons using *in vivo* whole-cell patch clamp recordings, a combination of acoustic and piezoelectric stimuli, and genetic inactivation of afferent auditory neurons. Our preliminary findings indicate that these central auditory neurons encode low-frequency amplitude modulations and that the preferred modulation frequency varies across the population. Pharmacological experiments suggest that GABAergic inhibition plays a key role in shaping these diverse responses. Based on these neurons' receptive field properties, we hypothesize that they receive convergent inputs from a functionally heterogeneous set of primary auditory neurons. These findings shed light on how auditory neurons in any system extract the relevant features of acoustic stimuli.

## II-53. Synaptic computation of interaural level differences in rat auditory cortex

Michael Kyweriga  
Whitney Stewart  
Carolyn Cahill  
Michael Wehr  
University of Oregon

MKYWERIGA@GMAIL.COM  
WHITNSTEWART@ME.COM  
CCAHILL26@GMAIL.COM  
WEHR@UOREGON.EDU

Cortical sensory neurons are tuned to specific stimuli. The synaptic mechanisms underlying this tuning typically follow either of two different strategies: inheritance or local processing. In the inheritance strategy, the tuning of excitation matches that of the spiking output of the cell. In the local processing strategy, excitation is more broadly tuned than spiking output, with inhibition suppressing spiking for non-preferred stimuli. In this strategy, inhibition refines the inherited tuning by suppressing spiking that would otherwise be driven by excitation. Interaural Level Differences (ILDs) are sound localization cues that are extensively processed in the auditory brainstem and mid-brain. Here we directly compared spiking responses with synaptic excitation and inhibition in the same neurons across a range of ILDs. We found that cells preferring contralateral ILD sounds (CP cells), followed the inheritance strategy. CP cells received proportional excitatory and inhibitory inputs across midline and contralateral ILDs. These cells spiked when they received enough excitation and inhibition did not shape the spiking output. In contrast, cells preferring midline ILD sounds (MP cells), showed evidence of the local processing strategy. MP cells received excitatory inputs across midline and contralateral ILDs. However, sounds from contralateral ILDs elicited strong inhibition that quenched the spiking output, only allowing spiking to occur at midline ILDs. These results suggest that in the rat auditory cortex, CP cells do not utilize inhibition to shape ILD sensitivity, whereas MP cells do. Accordingly, a conductance-based integrate and fire model demonstrated that simulated removal of inhibition converted MP cells into CP cells, whereas CP cells were unchanged. We conclude that an auditory cortical circuit computes sensitivity for midline regions of auditory space.

## II-54. Near-optimallinear decoding for marginalization: application to heading estimation

HyungGoo R Kim<sup>1</sup>  
Xaq Pitkow<sup>2,3</sup>  
Ryo Sasaki<sup>1</sup>  
Dora E. Angelaki<sup>2</sup>  
Gregory C. DeAngelis<sup>1</sup>

HKIM@BCS.ROCHESTER.EDU  
XAQ@CNS.BCM.EDU  
RSASAKI@CVS.ROCHESTER.EDU  
ANGELAKI@CABERNET.CNS.BCM.EDU  
GDEANGELIS@CVS.ROCHESTER.EDU

<sup>1</sup>University of Rochester

<sup>2</sup>Baylor College of Medicine

<sup>3</sup>Rice University

Neural responses of many sensory neurons depend on multiple stimulus parameters, but often only one stimulus parameter needs to be estimated to govern behavior. To estimate one stimulus feature and ignore other (nuisance) parameters, the brain can marginalize the joint posterior probability function. If nuisance parameters simply change the gain of responses, then optimal marginalization is done automatically with a simple linear readout (Ma et al 2006). However, nuisance parameters may not simply scale sensory responses, making linear decoding suboptimal. In area MSTd, multisensory neurons that respond to both optic flow and physical translation of the head are involved in heading estimation. When large moving objects distort the optic flow field, they exert complex effects on the heading tuning of MSTd neurons and may also bias perception. Because these effects violate the assumptions of Ma et al, it is not clear how effectively heading can be estimated from MSTd activity in the presence of moving objects. We sought the best approximation to the marginal posterior that can be obtained using a linear probabilistic population code (PPC). We show that multinomial logistic regression (MLR) can be used to estimate decoding weights for each neuron that minimize the KL divergence between the true marginal posterior and the

linear PPC approximation. Using simulations of MSTd responses to self-motion and object motion, we show that MLR can be used to obtain an excellent approximation to the marginal posterior over heading, even when object motion is assumed to shift the heading preferences of model neurons. We applied the MLR method to responses of MSTd neurons tested with many combinations of heading and object directions. Heading was recovered accurately from multisensory responses of MSTd neurons using decoding weights derived from MLR, whereas decoding by visual or vestibular heading tuning yielded large errors.

## II-55. Integration of visual and tactile signals in behaving rats

Nader Nikbakht<sup>1</sup>  
Azadeh Tafreshiha<sup>1</sup>  
Rodrigo Q. Quiroga<sup>2</sup>  
Davide Zoccolan<sup>1</sup>  
Mathew Diamond<sup>1</sup>

NADER.NIKBAKHT@GMAIL.COM  
ASEYEDTA@SISSA.IT  
RQQG1@LEICESTER.AC.UK  
ZOCCOLAN@SISSA.IT  
DIAMOND@SISSA.IT

<sup>1</sup>International School for Advanced Studies

<sup>2</sup>University of Leicester

Our experience of the world depends on integrating cues from multiple senses to form unified percepts. How the brain merges information across sensory modalities has been the object of debate. To measure how rats bring together information across sensory modalities, we devised an orientation categorization task that combines vision and touch. Rats encounter an object—comprised of alternating black and white raised bars—that looks and feels like a grating and can be explored by vision (V), touch (T), or both (VT). The grating is rotated to assume one orientation on each trial, spanning a range of 180 degrees. Rats learn to lick one spout for orientations of 0°45 degrees ('horizontal') and the opposite spout for orientations of 90°45' ('vertical'). Though training was in VT condition, rats could recognize the object and apply the rules of the task on first exposure to V and T conditions, suggesting that the multimodal percept corresponds to that of the single modalities. To quantify performance, we fit psychometric curves to the choice data of each rat in each modality. Taking the slopes of the curves at their inflection point as a proxy for sensory acuity, we found that rats have good orientation acuity using their whiskers and snout (T condition); however typically performance is superior by vision (V condition). Performance is always highest in the VT condition, indicating multisensory enhancement. Is the enhancement optimal with respect to the best linear combination? To answer this, we computed the performance expected by optimal integration in the framework of Bayesian decision theory and found that most rats combine visual and tactile information better than predicted by the standard ideal—observer model (maximum likelihood prediction). To study the neuronal correlates of visuo-tactile integration, we are setting up electrophysiological recordings in somatosensory, visual and parietal association areas.

## II-56. 3D surface tilt estimation in natural scenes from image cue gradients

Johannes Burge  
Brian McCann  
Wilson Geisler  
University of Texas at Austin

JBURGE@MAIL.CPS.UTEXAS.EDU  
BRIAN.C.MCCANN@GMAIL.COM  
GEISLER@PSY.UTEXAS.EDU

Estimating the 3D shape of objects is a critical task for sighted organisms. Thus, it is important to understand how different image cues should be combined for optimal shape estimation. Here, we examine how gradients of various image cues—disparity, luminance, texture—should be combined to estimate surface tilt in natural scenes. Surface tilt is the direction in which a surface is receding most rapidly; for example, a ground plane straight-ahead has a surface tilt of 90 deg. Estimating surface tilt is necessary for recovering surface orientation and 3D shape. To determine how image cues to surface tilt should be optimally combined, we collected a database of



stereoscopic natural images with precisely registered range images, using a robotically positioned DSLR camera and laser range scanner. For each pixel in each registered image (~109 samples) we computed the gradients of range, disparity, luminance, and texture within a local area (0.6 deg). Finally, we computed the conditional mean of the range-gradient orientation (the ground-truth surface tilt), given the orientations of the image gradients. These conditional means are the Bayes optimal (MMSE) estimates of the surface tilt given the image cues, and are free of assumptions about the underlying joint probability distributions. A rich set of results emerges. First, the prior probably distribution over surface tilts in natural scenes exhibits a strong cardinal bias. Second, the likelihood distributions for disparity, luminance, and texture are each somewhat biased estimators of surface tilt. Third, the optimal estimates of surface tilt are more biased than the likelihoods, indicating a strong influence of the prior. Fourth, when all three image cues agree, the optimal estimates become nearly unbiased. Fifth, when the luminance and texture cues agree they often override disparity in the estimate of surface tilt, but when they disagree, they have little effect. Word limit reached.

## II-57. A new database for natural motion: movie data and analysis of natural motion statistics

Stephanie E Palmer<sup>1</sup>  
 Jared Salisbury<sup>1</sup>  
 Heather Yee<sup>1</sup>  
 Hanna Torrence<sup>1</sup>  
 Don Ho<sup>1</sup>  
 David Schwab<sup>2</sup>

SEPALMER@UCHICAGO.EDU  
 JARED.SALISBURY@GMAIL.COM  
 HKYEE@UCHICAGO.EDU  
 JTORRENCE@UCHICAGO.EDU  
 DONHO58@GMAIL.COM  
 DSCHWAB@PRINCETON.EDU

<sup>1</sup>University of Chicago

<sup>2</sup>Princeton University

While the characteristics of static images have been explored in large image repositories, much less is known or measured in the temporal domain. We describe a new database for natural motion, which is comprised of a large collection of fixed-camera, high resolution grayscale and color video from field recordings of moving objects. Videos include motion of both animate and inanimate objects from a variety of camera-to-subject distances. In contrast to other motion databases (CCTV surveillance footage, various head-attached camera recordings, commercial movies), this set carefully selects natural subjects to maximize the quantity, quality and continuity of motion in each clip. These data are ideal for use as stimuli in visual neuroscience experiments because of their variety of motion and the fine temporal and spatial information available in each recording. We also report on preliminary analysis of the motion content of these scenes. Static images strikingly exhibit a power-law distribution of spatial frequencies. Power-law behavior in the frequency distribution of temporal fluctuations in total scene luminance has also been observed in natural contexts, and we find that our scenes similarly display power-law behavior with a scaling exponent that depends on the particular motion content. To quantify the dynamics of motion in our scenes, we fit translation fields to pairs of stimuli and analyze the statistics of the resulting flow. These flow fields are then used to guide the trajectories of "tracer" particles released at various points in the frame and subjected to the inferred flow from the natural movie. One important aspect understanding of predictive computation in the brain is characterizing the structure of predictable events in the world itself. These tracer particles allow us to define the predictive information content of naturalistic trajectories.

## II-58. Higher-order stimulus correlations drive a sparse representation in mouse V1 well-suited for discrim

Emmanouil Froudarakis<sup>1</sup>  
 Philipp Berens<sup>2</sup>  
 R. James Cotton<sup>1</sup>  
 Alexander S Ecker<sup>2</sup>  
 Fabian H. Sinz<sup>2</sup>  
 Dimitri Yatsenko<sup>1</sup>  
 Peter Saggau<sup>1</sup>  
 Matthias Bethge<sup>2,3</sup>  
 Andreas S. Tolias<sup>1</sup>

EFROUD@CNS.BCM.EDU  
 PHILIPP.BERENS@UNI-TUEBINGEN.DE  
 RCOTTON@CNS.BCM.EDU  
 ALEXANDER.ECKER@UNI-TUEBINGEN.DE  
 FABIAN.SINZ@UNI-TUEBINGEN.DE  
 YATSENKO@CNS.BCM.EDU  
 PETER@CNS.BCM.EDU  
 MATTHIAS@BETHGELAB.ORG  
 ATOLIAS@CNS.BCM.EDU

<sup>1</sup>Baylor College of Medicine

<sup>2</sup>University of Tuebingen

<sup>3</sup>Bernstein Center for Computational Neuroscience

The neural code is believed to have adapted to the statistical properties of the organism's natural environment. However, the principles that govern the organization of ensemble activity in the visual cortex during natural visual input are unknown. We used a novel 3D high-speed imaging method to image populations of up to 500 neurons in the mouse primary visual cortex and characterize the structure of their activity. We found that higher-order correlations in natural scenes, carrying their perceptually informative statistical structure, induce a sparser code in which information is encoded by reliable activation of a smaller set of neurons when compared to control stimuli. Ensemble activity patterns spread out in the high-dimensional neural space, making the read-out of natural stimuli easier. This functional benefit of sparse coding for discrimination could be a general principle governing the structure of the population activity throughout cortical microcircuits.

## II-59. Representation and reconstruction of natural scenes in the primate retina

Alexander Heitman<sup>1</sup>  
 Martin Greschner<sup>2</sup>  
 Greg Field<sup>3</sup>  
 Peter Li<sup>4</sup>  
 Daniel Ahn<sup>1</sup>  
 Alexander Sher<sup>5</sup>  
 Alan Litke<sup>5</sup>  
 E. J. Chichilnisky<sup>6</sup>

ALEXKENHEITMAN@GMAIL.COM  
 GRESCHNER@SALK.EDU  
 GDFIELD@USC.EDU  
 PETERLI@SALK.EDU  
 DANNYSJAHN@GMAIL.COM  
 SASHAKE3@UCSC.EDU  
 ALAN.LITKE@CERN.CH  
 EJ@STANFORD.EDU

<sup>1</sup>University of California, San Diego

<sup>2</sup>University of Oldenburg

<sup>3</sup>University of Southern California

<sup>4</sup>Salk Institute for Biological Studies

<sup>5</sup>Santa Cruz Institute for Particle Physics

<sup>6</sup>Stanford University

Artificial stimuli such as sinusoidal gratings and white noise are useful for measuring and modeling the light responses of visual neurons. However, it remains unclear how results obtained with such stimuli generalize to understanding visual function under natural conditions. We have begun testing this generalizability in the collective visual responses of retinal ganglion cells (RGCs): (1) How accurately can models of RGC response replicate spike trains elicited by natural stimuli? (2) How accurately can natural stimuli be reconstructed from the activity of RGC populations? We performed 512-electrode recordings from identified RGC populations in isolated peripheral macaque retina. Functional cell type classification revealed complete mosaics of ON-parasol and

OFF-parasol RGCs, which provide the dominant input to the magnocellular layers of the LGN. Naturalistic stimuli consisted of selected images from the van Hateren database, spatio-temporally displaced to match previously recorded fixational eye movements in awake macaques. (1) The replicability of RGC responses to natural scenes was examined by model fitting and simulation. A generalized linear model was fitted to responses to white noise and natural scenes, and tested for cross-validated accuracy. The replication of responses to naturalistic stimuli was markedly inferior to the replication of responses to white noise stimuli. Additional work will focus on variants of the model that more accurately replicate RGC responses. (2) The reconstruction of naturalistic stimuli was assessed with optimal linear decoding from recorded RGC spike trains. Filters used for decoding converged rapidly, exhibiting a spatial structure similar to receptive fields measured with white noise. These filters permitted reconstruction of features of the images at spatial scales similar to the receptive field size. Additional work will focus on the importance of fixational eye movements in reconstruction.

## II-60. A selective color-opponent pathway in the mouse retina

Maximilian Joesch<sup>1</sup>  
Markus Meister<sup>2</sup>

JOESCHKROTKI@FAS.HARVARD.EDU  
MEISTER@CALTECH.EDU

<sup>1</sup>Harvard University

<sup>2</sup>California Institute of Technology

Like most mammals, mice are dichromatic: They perceive a two-dimensional color space. Their cone photoreceptors produce two pigments: the S- and M-pigments that maximally absorb at 370nm and 510nm, respectively. The two pigments form countervailing gradients in the retina, ranging from pure S at the ventral edge to mostly M at the dorsal margin, with a rare pure-S cone uniformly distributed over the entire retina. How RGCs combine these inputs to detect color remains obscure. Color vision relies on a comparison of light detected by different pigments, and this would seem difficult given the spatial segregation of pigments and their frequent co-expression in the same cone. Here we report for the first time a genetically defined retinal ganglion cell with stereotyped color-opponency: It shows an OFF-response to UV light and an ON-response to green light. The receptive field is arranged antagonistically, with a UV-OFF center and Green-ON surround. A transgenic mouse line labels these cells with a fluorescent marker, which enabled a concerted study to test the underlying retinal pathways. The one cone-selective pathway described in mammals, the S-ON bipolar cell, seems to play no role. As a further puzzle, the color-opponent response is most prevalent in ventral retina, which contains virtually no M pigment, and should be devoid of Green-sensitive cones. Recent work suggests that the Green-ON surround may derive instead from rod photoreceptors, which absorb maximally at 500nm. Rods and cones are active together over a wide range of light levels, including the crepuscular intensities relevant to mice. Ongoing work investigates how those signals are segregated into color-opponent pathways leading to the ganglion cell. This study promises a deeper understanding of retinal color-opponency. Furthermore, the genetic control of color-opponent RGCs offers an entry into downstream circuits and color processing in the brain.

## II-61. Early representations of whisker self-motion revealed using closed-loop control in anesthetized rats

Avner Wallach  
Knarik Bagdasarian  
Ehud Ahissar

AVNER.WALLACH@WEIZMANN.AC.IL  
BAGDASARIAN.KNARIK@WEIZMANN.AC.IL  
EHUD.AHISSAR@WEIZMANN.AC.IL

Weizmann Institute of Science

Sensation is an active process, in which the activity of the sensory organ affects and selects the nature of the sensed data that are derived from the external environment. The vibrissal somatosensory system is currently a popular animal model in the study of mammalian active perception. Rats, mice and several other rodent species

use the long hairs (vibrissae or whiskers) on either side of their snout to interrogate their proximal surroundings. Exploratory rhythmic movements of the whisker array are used to actively acquire tactile information on both the location and nature of nearby objects. This information is conveyed from the whisker follicles via the trigeminal nerve towards the brainstem, and from there to motor nuclei and to higher level structures. The sensory information streamed into the vibrissal system during active sensation and the processing it undergoes are still largely unknown. One of the main factors hindering a systematic characterization of this system is that the self-generated movements are shut-down under anesthesia. Here we present a novel paradigm in which a computer-based, real-time motion control algorithm is used to generate awake-like whisker-pad movements in anesthetized rats. Thus, the computer artificially closes the sensory-motor loop, substituting for the neural pathways which are affected by anesthesia. Using this hybrid brain-computer system we were able to characterize the response of individual neurons to a wide range of natural-like movements. Moreover, the ability to precisely repeat the very same movements numerous times enabled us to reconstruct typical population responses to such movements. Our analysis show that the secondary afferents in the brainstem's interpolar trigeminal nucleus preferentially respond to the protraction phase of the whisking cycle, while primary afferents of the trigeminal ganglion exhibit more diverse response patterns. These results reveal early neuronal transformations of the vibrissal system.

## **II-62. An optimal object-tracking model provides a unifying account of motion and position perception**

Oh-sang Kwon  
Duje Tadin  
David Knill

University of Rochester

OSKWON@CVS.ROCHESTER.EDU  
DUJE@CVS.ROCHESTER.EDU  
KNILL@CVS.ROCHESTER.EDU

A central biological function of vision is to track objects over time. Vision uses sensory signals about object position and pattern motion to iteratively update internal state estimates of object velocity and position. The problem becomes complicated when the generative process for the motion signals is ambiguous. For example, the pattern motion generated by the texture on an object provides a strong but ambiguous motion signal that may indicate object motion (e.g. running tiger), pattern motion within an object (e.g. barber pole), or both (e.g. flying soccer ball that is rotating). An optimal motion tracking system resolves this source attribution problem using internal models of different kinds of motion in the world. We show that a number of illusory influences of pattern motion on perceived position arise naturally from the behavior of such a tracking system. For example, such a system accounts for the perceived position shift of objects in the direction of pattern motion and the apparent curvature of a moving object with embedded pattern motion. Most importantly, the quantitative pattern of these illusory effects is well-modeled by a Kalman filter that optimally tracks position and velocity of an object through a process of statistical inference in which position and motion estimates are derived by integrating noisy sensory inputs with the prediction of a forward model that reflects natural dynamics. The model further predicts previously unreported influences of the uncertainty of spatial coding on perceived motion that are born out in psychophysical experiments. In sum, the model provides a unifying theory of perceptual interactions between motion and position signals. The results also highlight the fact that our percepts of motion and position in dynamic displays necessarily co-vary with one another and that, for example, it is near impossible to completely isolate motion processing from the influence of position signals.

## **II-63. Binding by synchrony in complex-valued deep neural networks**

David P Reichert  
Thomas Serre  
Brown University

DAVID\_REICHERT@BROWN.EDU  
THOMAS\_SERRE@BROWN.EDU

Deep learning approaches have achieved state-of-the-art performance in tasks such as object recognition (Krizhevsky et al., 2012). As hierarchical neural networks that learn representations from sensory data, these approaches are also relevant to computational neuroscience (Cadieu et al., 2013): for example, deep convolutional nets are closely related to biological models like HMAX (Serre et al., 2007), and deep Boltzmann machines have been applied as models of generative cortical processing (Reichert et al., 2013). However, deep networks are still outmatched by the power and versatility of the cortex, perhaps in part due to the richer neuronal mechanisms available to real cortical circuits. Here, we argue that timing of neuronal firing could be such a key mechanism and qualitatively transform how deep networks represent the sensory world. Synchrony in particular has long been hypothesized to play a crucial role in the brain, for instance by binding (Singer, 2007) distributed neuronal populations dynamically to represent coherent entities in the sensory world, such as visual objects. To make aspects of spike timing like neuronal synchrony amenable to the theoretical language of deep networks, we present a formulation based on complex-valued neuronal units that have both a firing rate and a phase, the latter indicating properties of spike timing, e.g. activity relative to a global network rhythm. We apply this approach to deep networks such as deep Boltzmann machines, and demonstrate in simulations what functional roles synchrony could play, gating information transmission and binding distributed representations. By combining a hypothetical key element of biological processing with state-of-the-art machine learning, we aim to lay the groundwork for what could lead both to improved machine learning approaches and novel powerful models for computational neuroscience.

## II-64. Interplay between sensory statistics and neural circuit constraints in visual motion estimation

James E Fitzgerald<sup>1</sup>  
Damon Clark<sup>2</sup>

JAMESFITZGERALD@FAS.HARVARD.EDU  
DAMON.CLARK@YALE.EDU

<sup>1</sup>Harvard University

<sup>2</sup>Yale University

Many animals use visual signals to estimate motion. The Hassenstein-Reichardt correlator and motion energy models each suppose that animals estimate motion by computing pair correlations between spatiotemporally separated visual signals. In contrast, the Bayes optimal motion estimator utilizes prior information regarding expected stimuli to compute motion from a broader set of correlations. Statistical light-dark asymmetries in natural environments imply that both even-ordered and odd-ordered correlations contribute to the optimal estimator. Interestingly, recent experiments indicate that humans and flies perceive motion from triple correlations that signify naturalistic motion. It remains unclear whether animals directly compute stimulus correlations or optimally account for the statistics of natural motion. Here we study how nonlinear visual processing enables the computation of higher-order stimulus correlations that improve estimation accuracy. More specifically, we frame visual motion estimation as a problem of constrained statistical inference and consider four candidate constraints. The brain can either: (i) only compute pair correlations in the stimulus; (ii) apply a static nonlinearity to visual inputs before computing pair correlations; (iii) construct higher-order estimators using only filtered signals already computed for its pairwise estimator; or (iv) linearly combine ON/OFF motion processing channels that underlie pairwise motion estimation. We find that beyond the first constraint, each optimal estimator incorporates pair and triple correlations into its motion estimate. Nevertheless, each constraint requires the use of a distinct subclass of these correlation types. Optimal static nonlinearities and ON/OFF processing incorporate many higher-order correlations. Interestingly, our results reveal two ways that efficient visual coding could improve motion estimation. First, contrast equalization transforms visual input streams in a manner that makes their cross-correlation more predictive of motion. Second, segregating motion signals into ON/OFF channels enhances estimation accuracy. Overall, our results provide further theoretical rationale for higher-order correlations in motion estimation and demonstrate how brains might extract such signals.

## II-65. Selectivity for shape and blur in area V4

Timothy Oleskiw  
Amy Nowack  
Anitha Pasupathy  
University of Washington

OLESKIW@UW.EDU  
ALN@U.WASHINGTON.EDU  
PASUPAT@UW.EDU

A complete understanding of visual perception requires concise models for how complex information like object shape is represented in the brain. Computational and psychophysical studies have argued for the importance of boundary contours in visual processing, showing how the statistics of edges in natural scenes can be exploited to efficiently segment and represent objects. Further, many neurons in primate area V4 have been shown to be selective to an object's boundary curvature, suggesting shape as an integral feature of mid-level ventral processing. However, these studies typically use shape stimuli with sharp unnatural boundaries, or natural stimuli wherein object shape is difficult to quantify. In addition to contrast, the blur of an edge, i.e. the contrast gradient across a boundary contour, reveals useful object information such as 3D structure and depth. In addition it has been shown that edge information, including the magnitude of blur at each edge, is sufficient for encoding natural images. Currently unknown however is if cortical areas like V4 encode blur along with shape. Here we present data recorded from single-unit macaque V4 neurons exhibiting shape selectivity. After manually characterizing RF properties, preferred and non-preferred shape stimuli were presented under Gaussian blur, i.e. low-pass filtered over multiple frequency bands. We report that neurons are tuned to blur, with individual cells exhibiting peak responses at various blur ranges. Surprisingly, shape selectivity appears independent of blur tuning, especially over modest blur where boundary shape is clear. Thus, population responses of V4 not only convey object shape, but also information of boundary blur, supporting a role for V4 in developing object percepts with 3D structure and depth. A novel model capturing both curvature and blur selectivity is proposed to explain our data, giving insight into how V4 responses can form a complete representation of natural scenes.

## II-66. Decoding what types of information are in the macaque face patch system

Ethan Meyers<sup>1</sup>  
Mia Borzello<sup>2</sup>  
Winrich Freiwald<sup>3</sup>  
Doris Tsao<sup>4</sup>

EMEYERS@MIT.EDU  
MIABORZ@GMAIL.COM  
WFREIWALD@MAIL.ROCKEFELLER.EDU  
DORTSAO@CALTECH.EDU

<sup>1</sup>Massachusetts Institute of Technology  
<sup>2</sup>MGH  
<sup>3</sup>Rockefeller University  
<sup>4</sup>California Institute of Technology

Faces are a biologically important class of stimuli for primates. Recent work has identified six discrete face areas in the temporal lobe of the macaque that form a network which appears to be responsible for processing information related to faces (Moeller et al., 2008; Tsao et al., 2008; Freiwald and Tsao, 2010). The vast majority of neurons in these face areas have much higher firing rates to images of faces compared to other object categories (Tsao et al., 2006; Freiwald and Tsao, 2010; Issa and DiCarlo, 2012), however it is still unclear what types of information, and consequently which visual behaviors, each face area could support. In this work we use neural population decoding analyses to better characterize what information is being represented in three of these face areas (ML/MF, AL, and AM). Our results show that ML/MF and AL contain significant visual information that can be used to discriminate between non-face objects, while decoding of non-face objects from AM responses is comparatively worse. Additionally, we find that information about face identity that is invariant to the pose of the head is largely absent from ML/MF, is stronger in AL and is strongest in AM. These findings show that the face patch system becomes selective to visual features in the more anterior patches that are useful for identifying individuals while losing information that is useful for discriminating between different non-face objects.

## II-67. Efficient coding can account for the existence of sustained and transient bipolar cell types

Tao Hu<sup>1</sup>  
Dmitri B Chklovskii<sup>2,3</sup>

TAOHU@TEES.TAMUS.EDU  
MITYA@JANELIA.HHMI.ORG

<sup>1</sup>Texas A&M University

<sup>2</sup>Janelia Farm Research Campus

<sup>3</sup>Howard Hughes Medical Institute

A salient feature of the vertebrate retina is the existence of parallel signaling channels. For example, photoreceptors innervate two types of bipolar cells - sustained and transient – so named after the time course of their response to a step-stimulus (Werblin & Dowling, 1969, Awatramani & Slaughter, 2000). What could be the functional reason for the existence of such cell types? One possibility is that the segregation into transient and sustained signaling channels is a consequence of the efficient coding hypothesis (Attneave, 1954, Barlow, 1961). Efficient coding in parallel channels can be achieved by decorrelating signals in these channels by choosing their temporal kernels in the form of features of the principal component analysis (PCA). These features can be learned by minimizing the cumulative error of representing natural stimuli. For correlated stimuli with temporally invariant covariance the features are close to Fourier functions. Because Fourier functions do not match temporal kernels of bipolar cell, in this case, the efficient coding hypothesis may appear false. Here, we demonstrate that, under a natural modification of the cost function, the efficient coding hypothesis correctly accounts for sustained and transient channels. Specifically, we argue that errors of representing older signals must be less costly than those of the more recent ones. Thus, we propose to discount, similarly to weighted PCA, representation errors in the cost function with a factor decaying exponentially into the past. By minimizing such discounted cumulative error of representing natural signals we find that the features approximate Laguerre functions and closely match temporal kernels of sustained and transient bipolar cells. Our result can be easily generalized to other early sensory processing pathways with parallel sustained and transient channels such as, for example, fruit fly auditory system.

## II-68. Characterizing retinal network adaptation in *Drosophila* using models of mean and contrast adaptation

Daniel Coca  
Carlos Luna Ortiz  
Stephen Billings  
Mikko Juusola

D.COCA@SHEFFIELD.AC.UK  
CARLOS.LUNA@SHEFFIELD.AC.UK  
S.BILLINGS@SHEFFIELD.AC.UK  
M.JUUSOLA@SHEFFIELD.AC.UK

University of Sheffield

Despite having a limited dynamic range, fly photoreceptors operate over the full environmental range of light intensities. To achieve this feat, photoreceptors implement various adaptation mechanisms to continuously adjust their sensitivity. In addition, photoreceptor responses in flies are modulated by feedback from two classes of interneuron, large monopolar cells (LMC) and amacrine cells (AC), and axonal gap-junctions, which pool the responses from six photoreceptors. Previous studies have shown that adaptation in the photoreceptor-interneuron synapses helps extend further the operating range of photoreceptors. In histamine deficient mutants *hdcJK910* the lamina interneurons fail to receive and transmit visual information and their feedback synapses can no longer modulate the photoreceptor output. By comparing the 'open-loop' system of mutant flies with the 'closed-loop' system of wild-type flies it is possible to characterize how the network helps enhance the photoreceptor response. We have previously developed a functional nonlinear model of *Drosophila*'s R1-R6 photoreceptors which incorporates a dynamic model of photoreceptor gain. This gain control model continuously adapts so that the model accurately predicts photoreceptor responses to arbitrary stimuli, measured in wild-type flies, over the full range of light intensities. The model however, cannot be used to describe responses measured in *hdcJK910* mutants. Here we propose a new control structure to model photoreceptor adaptation, which incorporates separate gains for mean intensity and contrast. By exploiting this control architecture we were able to a) model the responses of wild type

as well as blind mutants using a single model structure and b) characterize the contribution from interneurons to boosting the operating ranges of fly photoreceptors. Specifically we found that lack of feedback from visual interneurons significantly reduces the contrast gain and slows down mean adaptation in very bright conditions. We have previously developed a functional nonlinear model of *Drosophila*'s R1-R6 photoreceptors which incorporates a dynamic model of photoreceptor gain, which continuously adapts to luminance and contrast so that the model accurately predicts photoreceptor responses to arbitrary stimuli over the full range of light intensities. Here we propose a new control structure to model photoreceptor adaptation, which incorporates separate gains for mean intensity and contrast. By exploiting this control architecture we were able to a) model the responses of wild type as well as blind mutants using a single model structure and b) characterize the contribution from interneurons to boosting the operating ranges of fly photoreceptors. Specifically we found that lack of feedback from visual interneurons significantly reduces the contrast gain and slows down mean adaptation in very bright conditions.

## II-69. It cortex contains a general-purpose visual object representation

Ha Hong<sup>1</sup>  
Daniel Yamins<sup>1</sup>  
Najib Majaj<sup>2</sup>  
James DiCarlo

HAHONG@MIT.EDU  
YAMINS@MIT.EDU  
NAJIB.MAJAJ@NYU.EDU  
DICARLO@MIT.EDU

<sup>1</sup>Massachusetts Institute of Technology

<sup>2</sup>New York University

Extensive research has shown that Inferior Temporal (IT) cortex is a key brain area underlying invariant object recognition. More recently, it has been uncovered that position is also coded in IT and human LOC. Here we show that IT neurons support robust readout of a variety of object parameters that characterize scene description in the central visual field. We recorded neural responses in IT and V4 to set of 5760 images of photorealistic objects in a variety of categories, placed in complex realistic scenes, with significant variation in object position, size, and in-plane and out-of-plane rotation. Consistent with known results, IT achieves high invariant categorization and identification performance for these images. We also find that the IT representation of object position is highly robust, with units that encode location accurately across the full range of tested positions ' even across widely varying object geometries, pose and size variation, and cluttered backgrounds that make this task very challenging for lower-level visual representations. We find similarly robust IT encodings for a variety of additional object parameters, including size, pose, perimeter, and aspect ratio, for which lower-level representations appear to have effectively no decoding power. While IT exhibits the ability to discount identity-preserving variation to solve categorization tasks, it simultaneously encodes a suite of 'identity-orthogonal' dimensions, that, combined with category and object identity encodings, form a basis for a full scene description. Moreover, while the representation of object identity and category is highly distributed across IT sites, the representations for these other properties (e.g. position) is typically more sparsely encoded, with a small proportion of highly responsive sites responsible for much of the decoding capacity. Taken together, these results suggest that IT contains a general representation of the visual environment in which key object parameters have been extracted and factored.



## II-70. Interaction of color, motion and contrast in macaque V4 population activity and behavior

Ivo Vanzetta<sup>1,2</sup>  
 Maria Bermudez<sup>1</sup>  
 Thomas Deneux<sup>1</sup>  
 Frederic Barthelemy<sup>1</sup>  
 Guillaume S Masson<sup>3,4</sup>

IVO.VANZETTA@UNIV-AMU.FR  
 MARIA.BERMUDEZ@UNIV-AMU.FR  
 THOMAS.DENEUX@UNIV-AMU.FR  
 FREDERIC.BARTHELEMY@UNIV-AMU.FR  
 GUILLAUME.MASSON@UNIV-AMU.FR

<sup>1</sup>INT - CNRS UMR7289

<sup>2</sup>Aix-Marseille Universite

<sup>3</sup>Institut de Neurosciences de la Timone

<sup>4</sup>UMR7289 CNRS & Aix-Marseille Universite

It has been suggested that motion and color information are processed together in different combinations along two distinct pathways. One is specialized in the treatment of fast motion and the other in the treatment of slow motion being particularly sensitive to that of isoluminant stimuli. It has been speculated that V4 might be the neuronal substrate of the 'slow', color-sensitive one. To test this hypothesis, we used optical imaging to measure neuronal population responses in behaving macaque V4 to fast and slow, isoluminant and luminance-based drifting gratings of several contrasts. The monkey reported on the perceived motion direction. In agreement with human psychophysics, the V4 responses to slow isoluminant stimuli were stronger than to fast ones, suggesting that V4 could indeed be a neuronal substrate for the slow channel. However, the monkey's psychophysics was opposite to human: its performance was better for fast than slow moving gratings, both isoluminant and luminance-based ones. We figured that this difference in psychophysical performance between monkey and human might be due to different weights with which the slow and the fast pathway contribute to the perception of motion in macaque and human, respectively. To test this, we measured the equivalent luminance contrast of color, for slow and fast drifting gratings, in monkeys and humans. As expected, in humans the equivalent luminance contrast of color at 1 Hz was considerably larger than at 8Hz, whereas in monkey it was very similar. Our data suggest that, due to its high sensitivity to slow, isoluminant motion, V4 could indeed be part of the neuronal substrates of the slow motion processing channel identified in human psychophysics. However, our monkey psychophysical data suggest that the weights with which the slow and the fast pathways are combined into a motion percept are probably different in macaque as compared to humans.

## II-71. Evolutionary convergence in computation of local motion signals in monkey and dragonfly

Eyal Nitzany  
 Gil Menda  
 Paul Shamble  
 James R Golden  
 Ron Hoy  
 Jonathan Victor  
 Cornell University

EIN3@CORNELL.EDU  
 GM234@CORNELL.EDU  
 PSS92@CORNELL.EDU  
 JRG265@CORNELL.EDU  
 RRH3@CORNELL.EDU  
 JDVICTO@MED.CORNELL.EDU

Many species extract motion from visual signals to perform complicated and critical tasks such as navigation, obstacle avoidance, and prey capture. The visual systems of macaques and dragonflies provide an interesting contrast because both species are expert at these tasks, but clearly lack anatomical/ancestral homology. Thus, comparing them may yield insights into successful computational strategies for motion analysis. Analysis of visual motion is generally thought to occur in two stages: extraction of local motion signals, followed by their integration. Local motion signals include those that drive the classical Reichardt detector (pairwise spatiotemporal correlation, also known as Fourier (F) motion), as well as signals that the Reichardt detector cannot extract. The latter includes spatiotemporal correlations of higher order: non-Fourier motion (NF), which can arise from motion transparency

or occlusion, and Glider (G) motion, typically associated with objects that are looming (G expansion) and receding (G contraction). We recorded single unit neurons from early visual brain areas in macaques and dragonflies and found cells that were sensitive to all these kinds of motion cues (and their combinations) in both species, and in similar distributions. Moreover, higher-order motion types (NF and G) play a relatively greater role at later processing stages (V2 vs V1 in macaque, lobula vs. medulla in dragonfly). Many cells respond to two or more kinds of motion, and, across species, the pattern of these combined sensitivities is also similar. Interestingly, in the dragonfly, responses to G expansion are more closely linked to responses to F motion than in the macaque. These findings suggest that, despite their very different anatomies and evolutionary lineages, the visual systems of dragonflies and macaques use similar computations to detect local motion signals, but that later stages of processing emphasize the kinds of signals that are most important to their survival.

### II-72. Pop-out visual search of moving targets in the archer fish

Mor Ben-Tov<sup>1</sup>  
Opher Donchin<sup>1,2</sup>  
Ohad Ben-Shahar<sup>1</sup>  
Ronen Segev<sup>1</sup>

MORBENT@POST.BGU.AC.IL  
DONCHIN@BGU.AC.IL  
BEN-SHAHAR@CS.BGU.AC.IL  
RONENSGV@BGU.AC.IL

<sup>1</sup>Ben Gurion University  
<sup>2</sup>Erasmus MC

Visual search, the ability to find an object of interest against a background, needs to be accurate and fast to ensure survival. In mammals, this led to the development of a parallel search mode, pop-out, which under certain conditions enables fast detection time independent of the number of distracting objects. This pop-out search is believed to be based on a population of pop-out cells in the visual cortex that generates saliency maps based on center-surround contrast. A key question in the study of visual search is to what extent the visual cortex is a crucial component in supporting pop-out search. To address this issue, we tested the ability of the archer fish, which lacks a fully developed cortex, to perform pop-out visual search, and explored the neural mechanism underlying this behavior in its optic tectum, the main visual area in the brain. We found that the archer fish can perform pop-out visual search of moving targets. Specifically, a target moving bar that differs in motion features such as phase and speed from a set of distractors is salient for the fish and elicits reaction times independent of the number of distractors. To understand the mechanism underlying this behavior, we recorded from the optic tectum and found pop-out cells that evince stronger response when the stimulus in their receptive field differs in motion features from the stimulus outside of their receptive field. Furthermore, we found strong similarities between responses in behavioral and cellular experiments. In both, pop-out was stronger when the target and the background differed in two motion features simultaneously. Thus, pop-out cells in the optic tectum generate saliency maps that might be the basis of pop-out visual search in the archer fish. Our findings indicate possible universality of the mechanism of visual search across vertebrates.

### II-73. Bayesian inference for low rank receptive fields

Mijung Park  
Jonathan W Pillow  
University of Texas at Austin

BRIENTS@GMAIL.COM  
PILLOW@MAIL.UTEXAS.EDU

An important problem in systems neuroscience is to characterize how a neuron integrates sensory stimuli across time and space. Mathematically, this can be formalized as the problem of estimating a neuron's linear receptive field (RF) under a probabilistic encoding model. In typical reverse correlation experiments, RFs have high dimension due to the large number of coefficients needed to specify an integration profile across time and space. Estimating RFs therefore poses a variety of challenging computational and statistical problems. Here we ad-

dress these challenges by developing Bayesian methods for inferring low-rank RFs, that is, estimators for RFs parametrized as the sum of a small number of space-time separable filters. Low-rank parametrizations can substantially reduce the number of RF parameters (from thousands down to mere hundreds), and confers substantial benefits in both statistical power and computational efficiency. Our method relies on a novel prior defined by the restriction of a matrix normal prior to the manifold of low-rank matrices. This prior is governed by a pair of covariance matrices, one for columns (temporal vectors) and one for rows (spatial vectors), which allow us to build flexible, hierarchical models of temporal and spatial RF structure. Although the prior cannot be normalized analytically, it has conditionally Gaussian marginals over spatial and temporal components, and places a marginally Gaussian prior over RF coefficients. We develop two methods for inference in the resulting hierarchical model: (1) a fully Bayesian method using blocked-Gibbs MCMC sampling; and (2) a fast, approximate method that employs alternating ascent of conditional marginal likelihoods. We develop inference methods for use with Gaussian and linear-nonlinear-Poisson noise models, and show that low-rank estimates substantially outperform full rank estimates using neural data from retina and V1.

## II-74. Lateral feature-level interactions in anterior visual cortex of monkeys and humans

Chou Hung<sup>1</sup>  
 Chang Mao Chao<sup>2</sup>  
 Li-feng Yeh<sup>2</sup>  
 Yueh-peng Chen<sup>2</sup>  
 Chia-pei Lin<sup>2</sup>  
 Yu-chun Hsu<sup>2</sup>

CH486@GEORGETOWN.EDU  
 TOM33233323966@GMAIL.COM  
 LIFENG.YEH@GMAIL.COM  
 YUEPENG@GMAIL.COM  
 GOTO.LIN@GMAIL.COM  
 JASON0313@GMAIL.COM

<sup>1</sup>Georgetown University Medical Center

<sup>2</sup>National Yang-Ming University

Computational models of recognition typically represent objects via a 'bag of features', in which the spatial arrangement of model units is unimportant. However, neuroimaging and neurophysiological evidence suggest that humans and monkeys have coarsely similar representations in anterior visual cortex, including semantic and face/body maps, consistent with ecologically driven cortical clustering. Does representational similarity extend to finer scales, i.e. to specific short-range lateral interactions? We densely sampled neighboring feature columns via a 64-channel array (1.4 mm total width and depth, ~32 channels/column), inserted in macaque lateral anterior inferior temporal cortex under light anesthesia. Spike ensemble patterns were precisely tied to spatial patterns of spontaneous spike correlation. We extracted neurally defined complex features (principal component vectors), based on the ensemble's object preferences. The neurally defined features accurately predicted, across objects, functional neuroimaging patterns in human lateral occipital complex and not primary visual cortex. Informative voxels were co-localized for feature vectors from the same array, indicating a fine-scale match between macaque and human representations. The features also predicted the latency of texture filling-in perception in humans. Filling-in was fastest between textures composed of different objects that activated the same feature column, slower between textures composed of objects from unrelated columns, and slowest between textures composed of objects that activated neighboring columns, consistent with lateral inhibition. Controls (texture scrambling, low-level image statistics, analysis of category labels) showed that the effect was not due to low-level features or category overlap. Our results show the existence of specific feature-level interactions in anterior visual cortex of monkeys and humans, possibly mediated by lateral inhibition. They suggest that neighboring representations compete to represent opposing hypotheses (e.g. whether a feature is external or internal), consistent with relativistic coding and with the conjecture that local computations 'flatten the manifold' of complex features to enable robust, invariant recognition.

## II-75. Quantifying and modeling the emergence of object recognition in the ventral stream

Darren Seibert<sup>1</sup>  
Daniel Yamins<sup>1</sup>  
Ha Hong<sup>1</sup>  
James DiCarlo  
Justin Gardner<sup>2</sup>

DARREN@MIT.EDU  
YAMINS@MIT.EDU  
HAHONG@MIT.EDU  
DICARLO@MIT.EDU  
JUSTIN@BRAIN.RIKEN.JP

<sup>1</sup>Massachusetts Institute of Technology

<sup>2</sup>RIKEN Brain Science Institute

Humans can rapidly recognize objects from images cast on the retina despite massive variation due to changes in size, lighting, orientation, and position. Invariant object recognition is supported by a largely feed-forward network of neural circuitry located in the ventral stream. Substantial evidence shows that the ventral stream is a series of processing stages in which the representations in each successive area encode high-level image content (e.g. category, identity) more and more explicitly. However, the details of how these representations are formed remain unknown. Here, we use functional imaging to capture human neural responses at all ventral stages, on a set of 2500 images that has proved useful in exposing key object recognition challenges [1]. We characterize the neural representations for these stimuli – which contain a variety of 3D objects on complex backgrounds, at a broad range of position, scale, and pose views – in V1, V2, V4 and LOC. Because of the diversity and size of this dataset, we are able to quantify the stage-by-stage detailed emergence of object and category invariance in the neural representation using both performance measures and using Representational Dissimilarity Matrices (RDMs) [2]. These RDMs are a rich target for modeling: ideally, a model would exhibit internal structure that matches each RDM in sequence. To illustrate the utility of this idea, we evaluated a recent Convolutional Neural Network (CNN) model that achieves high performance on challenging object categorization tasks [3]. We find that the model layers show remarkably high similarity to V4 and LOC, indicating that performance constraints for which the model was optimized maybe have influenced the structure of intermediate and high-level ventral cortex representations. However, initial model stages are less effective at matching the lower ventral areas, suggesting that additional constraints need to be brought to bear to effectively capture cortical representations.

## II-76. Using deep networks to learn non-linear receptive field models of V2 neurons

Michael Oliver  
Jack Gallant  
University of California, Berkeley

MICHAELOLIVER@BERKELEY.EDU  
GALLANT@BERKELEY.EDU

The receptive fields of neurons in brain areas anterior to V1, such as V2 and V4, are high dimensional and non-linear. It has therefore proved difficult to model these neurons. To date there have been no successful attempts to model V2 neurons in terms of their joint spatio-temporal-chromatic response properties. Deep networks of rectified linear units have recently emerged as a powerful modeling framework for fitting large nonlinear models. We propose to leverage DropConnect, a new development in the training of deep neural networks, to learn interpretable and highly predictive models of foveal V2 neurons from their responses to natural color movies. The trained networks have the desirable property that the first layer learns an interpretable set of features. We use a simple marginalization procedure interpret the rest of the network in terms of these features and their interactions. Furthermore we present a closed form solution for inference in the case of rectified linear hidden units that improves upon DropConnect's sampling-based inference procedure.

## II-77. Four-layer sparse coding model of natural images that reproduces tuning properties in V1, V2 and V4

Haruo Hosoya<sup>1,2</sup>  
Aapo Hyvarinen<sup>3,1</sup>

HARUO.HOSOYA@GMAIL.COM  
AAPO.HYVARINEN@HELSINKI.FI

<sup>1</sup>ATR

<sup>2</sup>JST Presto

<sup>3</sup>Helsinki University

Although sparse coding models of natural images have been successful in reproducing properties of V1, extrastriate areas have proven more challenging, and direct quantitative comparisons to neurophysiological experiments are rare. Here, we extend existing sparse coding models and investigate the connection to experimental findings on V2 and V4. We trained a four-layer sparse coding model with natural movies, based on a multinomial model in conjunction with temporal coherence. After learning, the model exhibited response properties qualitatively and quantitatively compatible with several major neurophysiological results: 1) Model learned standard tuning and invariance properties related to orientation, frequency, and phase similar to V1 simple and complex cells in the first two layers. 2) Third layer represented subfield orientation integration with a distribution biased to smaller orientation differences consistent with a V2 study by Anzai et al. 3) Third layer further exhibited tuning to angles with response specificity to one componential orientation as in a V2 study by Ito and Komatsu. 4) Fourth layer exhibited tuning properties to position-specific curvatures with a distribution biased to acute convexities as in a V4 study by Pasupathy and Connor. However, despite these selectivities, the internal representations of third and fourth layers looked much more complex than simple geometric shapes. Indeed, the units in these layers responded much more strongly to parts of natural images that had particular textures or complex contours. In fact, geometric features were often parts of such naturalistic features, producing accidental tuning properties to geometric quantities.

## II-78. Low-dimensional structure of neural correlations in cortical microcircuits

Dimitri Yatsenko<sup>1</sup>  
Kresimir Josic<sup>2</sup>  
Alexander S Ecker<sup>3</sup>  
Emmanouil Froudarakis<sup>1</sup>  
R. James Cotton<sup>1</sup>  
Andreas S. Tolias<sup>1</sup>

YATSENKO@CNS.BCM.EDU  
JOSIC@MATH.UH.EDU  
ALEXANDER.ECKER@UNI-TUEBINGEN.DE  
EFROUD@CNS.BCM.EDU  
RCOTTON@CNS.BCM.EDU  
ATOLIAS@CNS.BCM.EDU

<sup>1</sup>Baylor College of Medicine

<sup>2</sup>University of Houston

<sup>3</sup>University of Tuebingen

Ambitious projects currently under way aim to record the activity of ever larger and denser subsets of neurons in behaving animals. It is commonly anticipated that correlations measured in such recordings will uncover aspects of functional organization of neural circuits. However, estimation and interpretation of large correlation matrices from finite recordings can be challenging. Estimation can be improved by regularization: the biasing of the estimate toward a low-dimensional, less variable approximation. The amount of improvement depends on how closely the reduced approximation captures the dominant interactions with fewest terms. Therefore, the selection of the most efficient estimator is an empirical question dependent on the system under investigation. We compared regularized correlation matrix estimators biased toward four respective families of low-dimensional correlation structures: independent, latent factors, sparse partial correlations, and sparse partial correlations + latent factors. The estimators' performance was evaluated on the noise correlation matrices from spatially compact groups of 91–314 neurons (average 206) in mouse visual cortex (31 sites in 24 animals) acquired by in vivo fast 3D random-access two-photon imaging of calcium signals. Recordings lasting between 15 and 20 minutes were deconvolved

and binned at 150 ms. Each estimator was optimized and evaluated by nested cross-validation. We found that both sparse partial correlations and latent factors were required for efficient estimation. In addition to reducing estimation error, optimized estimates provided compact representations of the correlation structure decomposed into a sparse network of pairwise interactions and global diffuse fluctuations. The density of positive interactions in these networks decreased rapidly with distance and with difference in orientation preference, whereas negative interactions were less selective. In ongoing studies, we seek to relate these elements of the correlation structure to circuits' anatomical organization, stimulus conditions, and brain states.

## II-79. Intrinsic networks in visual cortex are mediated by local and inter-areal rhythmic synchronization

Chris Lewis<sup>1</sup>  
 Conrado Bosman<sup>2</sup>  
 Thilo Womelsdorf<sup>2</sup>  
 Pascal Fries<sup>1</sup>

CHRIS.LEWIS@ESI-FRANKFURT.DE  
 CONRADO.BOSMAN@GMAIL.COM  
 THIWOM@YORKU.CA  
 PASCAL.FRIES@ESI-FRANKFURT.DE

<sup>1</sup>Ernst Strungmann Institute  
<sup>2</sup>Donders Institute

A classical approach to brain research investigates brain function as stimulus-response mapping. In this framework, variance in the brain's response to identical stimulation has been considered noise, because it does not correspond to extrinsic variables. However, the observed variance is in part attributable to intrinsic network dynamics, which are highly structured and impact behavior (Fox, 2007). The variability observed during both evoked and spontaneous activity occurs in a coordinated fashion across functional networks spanning essentially all spatial scales from individual neurons to the whole brain (Luczak, 2009, Fukushima, 2012, Fox, 2007). However, how intrinsic dynamics are coordinated from the activity of neuronal populations to whole brain networks remains a mystery. Rhythmic synchronization of neuronal activity is a compelling mechanism for inter-areal interactions. Local synchronization likely enhances neuronal impact through coincident input (Salinas, 2001) and inter-areal synchronization aligns temporal windows of excitability, rendering communication effective (Womelsdorf, 2007). We investigated whether rhythmic synchronization underlies intrinsic interaction between early and intermediate visual areas. Using high-resolution multi-area electrocorticography in awake monkeys, we found that intrinsic variation during stimulation (i.e. noise correlations), as well as in spontaneous activity during passive fixation matches the pattern of stimulus induced inter-areal activity. Local and inter-areal synchronization in specific frequency bands reproduced the retinotopic organization of early and intermediate visual areas. These band-limited dynamics exist despite the fact that the spontaneous activity observed during fixation shows no deviation from a characteristic 1/f spectrum. Further, using a Bayesian decoder, we show that the frequencies at which we observe retinotopically specific structure in intrinsic activity correspond to the frequency bands that have the greatest stimulus selectivity. These results indicate that rhythmic synchronization underlies similar spatial and spectral patterns of inter-areal coupling across behavioral states and suggests a mechanism by which local spontaneous activity may be coordinated across distributed whole brain networks.

## II-80. Spectral noise shaping: a multipurpose coding strategy in dorsal and ventral visual pathways

Grant Mulliken  
 Robert Desimone  
 Massachusetts Institute of Technology

GRANTM@MIT.EDU  
 DESIMONE@MIT.EDU

A well-established neurophysiological phenomenon observed during processing of a visual stimulus is the simultaneous decrease in low-frequency and increase in high-frequency local field potential (LFP) power. Suppression

of spontaneous low-frequency LFP activity might reflect a decrease in the noise floor in order to enhance the signal-to-noise ratio (SNR) of a rate code used to encode low frequency stimuli (e.g., Shin 1993, Mar et al., 1999, Chacron et al., 2005). We explored whether this postulated spectral shaping coding strategy is carried out in neurons in areas V4 and LIP of monkeys trained to perform an attentional tracking task. We first modeled spectral noise shaping in the frequency domain using topologies inspired from known recurrent inhibitory circuits in the visual cortex. During stimulus processing, we found that the spike and LFP spectra were shaped in a manner consistent with predictions made by 1st and 2nd order noise shaping models, respectively. Moreover, the observed shaping of the gamma LFP power (40-90 Hz) was modal and explained by a 2nd order system that exhibited resonance, which could be used for temporal coordination of neural activity. Pairwise signal and noise correlations for spike and LFP pairs were computed at multiple timescales, and their SNR spectra followed a power law relationship, consistent with noise shaping. Finally, the suppression of low frequency quantization noise power and increase in signal-dependent gamma power led to phase-dependent coding. Both mutual information (between spike count and stimulus position) and signal correlations were maximal at the peak phase of the theta rhythm and, by contrast, the trough of the gamma rhythm. Finally, the noise shaping model also gives rise to divisive normalization in the frequency domain.

## II-81. Spatiotemporal integration of optic flow and running speed in V1

Maneesh Sahani<sup>1</sup>

Marius Pachitariu<sup>1</sup>

Adil Khan<sup>2</sup>

Jasper Poort<sup>2</sup>

Ivana Orsolic<sup>3</sup>

Georg Keller<sup>4</sup>

Sonja Hofer<sup>3</sup>

Thomas Mrsic-Flogel<sup>3</sup>

MANEESH@GATSBY.UCL.AC.UK

MARIUS@GATSBY.UCL.AC.UK

ADILGKHAN@GMAIL.COM

J.POORT@UCL.AC.UK

IVANA.ORSOLIC@UNIBAS.CH

GEORG.KELLER@FMI.CH

SONJA.HOFER@UNIBAS.CH

THOMAS.MRSIC-FLOGEL@UNIBAS.CH

<sup>1</sup>Gatsby Computational Neuroscience Unit, UCL

<sup>2</sup>University College London

<sup>3</sup>University of Basel

<sup>4</sup>Friedrich Miescher Institute

The theoretical framework of predictive coding (PC) suggests that sensory cortical neurons report the difference between bottom-up inputs and top-down predictions of these inputs. On the other hand, the sensory cue-integration (SCI) framework implies that top-down predictions should integrate with noisy bottom-up inputs to improve the estimation of external stimulus parameters; under Gaussian uncertainty, this implies a positively-weighted combination of mismatched inputs. Neurons in mouse primary visual cortex (V1) respond to optic flow (a bottom-up, external input) and to the running speed of the animals (a top-down, internal signal). Recent results have suggested that these signals may be mostly combined in an additive fashion as the SCI framework predicts (Saleem et al 2013), although a small proportion of neurons (13%) respond to large mismatches between optic flow and running speed in agreement with the PC framework (Keller et al 2012). However, neither study modelled the full dependence of firing on the time-course of running speed or optic flow. We used the generalized-linear statistical modelling framework to fit temporal neural filters to optic flow and running speed in an open-loop setting where the optic flows presented to mice were not locked to running speed. Most neurons were found to be sensitive to the positive or negative derivative of both optic flow and running speed. Acceleration and deceleration events, both visual and running-related, explained five times more variance in the neural response than did the speed of the optic flow or running. Indeed tuning curves to speed were non-monotonic and complicated, while the tuning curves to acceleration increased or decreased more consistently. We also found that the tuning to the acceleration of optic flow and running had reliably opposite signs in 56% of neurons, in accordance with the PC prediction, although 17% of neurons did combine the two signals additively.

## II-82. Learning related changes of stimulus representations in mouse visual cortex

Jasper Poort<sup>1</sup>  
 Adil Khan<sup>1</sup>  
 Marius Pachitariu<sup>2</sup>  
 Abdellatif Nemri<sup>1</sup>  
 Ivana Orsolic<sup>3</sup>  
 Julija Krupic<sup>1</sup>  
 Marius Bauza<sup>1</sup>  
 Georg Keller<sup>4</sup>  
 Maneesh Sahani<sup>2</sup>  
 Thomas Mrsic-Flogel<sup>3</sup>  
 Sonja Hofer<sup>3</sup>

J.POORT@UCL.AC.UK  
 ADILGKHAN@GMAIL.COM  
 MARIUS@GATSBY.UCL.AC.UK  
 A.NEMRI@GMAIL.COM  
 IVANA.ORSOLIC@UNIBAS.CH  
 J.KRUPIC@UCL.AC.UK  
 M.BAUZA@UCL.AC.UK  
 GEORG.KELLER@FMI.CH  
 MANEESH@GATSBY.UCL.AC.UK  
 THOMAS.MRSIC-FLOGEL@UNIBAS.CH  
 SONJA.HOFER@UNIBAS.CH

<sup>1</sup>University College London

<sup>2</sup>Gatsby Computational Neuroscience Unit, UCL

<sup>3</sup>University of Basel

<sup>4</sup>Friedrich Miescher Institute

How does the cortical representation of a sensory stimulus change as it becomes behaviourally relevant for the animal? We explored this question in mouse primary visual cortex by using chronic two-photon imaging of genetic calcium indicators to follow neural populations in head-fixed behaving animals. Mice were trained to discriminate between two gratings with different orientations presented on the walls of a virtual corridor through which they ran. Within a week, mice learnt this go-no go task with >90% accuracy, by selectively licking in response to only one grating orientation. With learning, we found an increase in the proportion of cells representing the rewarded stimulus, as well as an increase in the discriminability of neural responses to the rewarded and non-rewarded stimuli. The time course of this increase was closely correlated with the day-to-day improvement in behavioural discrimination. We used a generalized linear statistical modelling framework to rule out that these results could be explained by the changes in running behaviour that occurred across learning. Furthermore, by examining responses of the same cells across days, we found that their activation patterns in response to the gratings became more reproducible between neighbouring training sessions with learning, indicating that the representation of these visual stimuli was stabilized as they became behaviourally relevant. This was confirmed by a decoding analysis which showed that a decoder trained on a session late in the learning process could also perform with high accuracy on subsequent sessions, whereas on early sessions this was not the case. We thus observe a dynamic period of rearrangement, expansion and separation of stimulus representations during the learning process, which is followed by stabilization. This has important implications for our understanding of how learning and experience shape the cortical processing of sensory input.

## II-83. Relationship between natural image statistics and lateral connectivity in the primary visual cortex

Philipp Rudiger<sup>1</sup>  
 Jean-Luc Stevens<sup>1</sup>  
 Bharath Chandra Talluri<sup>2</sup>  
 Laurent Perrinet<sup>3</sup>  
 James Bednar<sup>1</sup>

P.RUDIGER@ED.AC.UK  
 JLSTEVENS@ED.AC.UK  
 BHARATHCHANDRA2007@GMAIL.COM  
 LAURENT.PERRINET@UNIV-AMU.FR  
 JBEDNAR@INF.ED.AC.UK

<sup>1</sup>University of Edinburgh

<sup>2</sup>Institute for Adaptive and Neural Computation

<sup>3</sup>Institut de Neurosciences de la Timone

One of the fundamental questions of neuroscience is how the cortex encodes past experiences in its patterns



of connectivity. In primary visual cortex (V1) it is clear that the visual environment affects the orientation map structure [1]. However, little is known about the effect of the environment on the patterns of long-range lateral connectivity in V1 that have been implicated in surround modulation effects, such as contour integration [2]. Studies combining anatomy with electrophysiology have yielded confusing results, finding that horizontal connections are elongated along the axis of preferred orientation in some species, but not all. Given the rarity of this difficult type of experiment, it remains unclear whether these results reveal genuine species differences, or merely reflect differences in rearing environments, which are difficult to control and vary between species and laboratories. Here we first analyze databases of natural images and those from different types of laboratory animal cages, showing that these databases have important differences in the statistics of edge co-occurrences. Using methods from Perrinet et al. [3], we characterize these differences in terms of an expected 'association field' indicating the likelihood of occurrence of each pair of edges. We then generate predicted patterns of long-range lateral connectivity for animals reared in each of these environments, by training a Hebbian developmental model of V1 on each database. This model is based on a mechanistic formulation of a self-organizing map algorithm [4], extending it to have a more realistic pattern of lateral connectivity suitable for making direct comparison with experimental data [5]. The resulting predictions (a) help us understand the different results found from studies of different species, (b) suggest that animals reared in these different environments will have different performance on visual tasks, and (c) can be tested in future experiments raising animals in controlled environments.

## II-84. Can retinal ganglion cell dipoles seed iso-orientation domains in the visual cortex?

Manuel Schottdorf<sup>1</sup>  
 Stephen J. Eglén<sup>2</sup>  
 Fred Wolf<sup>1</sup>  
 Wolfgang Keil<sup>3</sup>

MANUEL@NLD.DS.MPG.DE  
 S.J.EGLEN@DAMTP.CAM.AC.UK  
 FRED@NLD.DS.MPG.DE  
 WKEIL@MAIL.ROCKEFELLER.EDU

<sup>1</sup>Max Planck Institute for Dynamics and Self-Organization

<sup>2</sup>University of Cambridge

<sup>3</sup>Rockefeller University

It has been argued that the emergence of repetitive orientation preference maps (OPMs) in the striate cortex (V1) of carnivores and primates can be explained by a so-called statistical connectivity model. This model assumes that input to V1 neurons is dominated by feed-forward projections originating from a small set of retinal ganglion cells (RGCs). The typical spacing between adjacent cortical orientation columns preferring the same orientation then arises via Moire-Interference between hexagonal ON/OFF RGC mosaics. While this Moire-Interference critically depends on long-range hexagonal order within the RGC mosaics, a recent statistical analysis of RGC receptive field positions found no evidence for such long-range positional order. Here, we introduce two new methods to generate RGCs mosaics that within the statistical connectivity framework yield aperiodic OPMs. First, we employ spatially correlated noise on hexagonal RGC mosaics. Second, we present a novel pairwise interacting point process (PIPP) to infer model RGC mosaics that within the statistical connectivity framework yield experimentally measured OPMs from cat V1. We find that spatially correlated noise does not obstruct a Moire-Effect between ON/OFF lattices on short scales, leading to OPMs with locally repetitive arrangements of orientation columns. Interestingly, the resulting OPMs exhibit column spacing heterogeneities similar to experimental observations. However, unrealistically strong hexagonality of RGC mosaics is still essential to obtain aperiodic OPMs. In contrast, PIPP mosaics exhibit realistic local positional statistics when compared to experimental mosaics. Furthermore, these mosaics show specific angular correlations between ON/OFF RGC pairs (dipoles) that match the spatial correlation of orientation preferences found in cortical OPMs. In data from both cat beta cell mosaics (>290 RGCs) as well as primate parasol receptive field mosaics (>200 RGCs) we find that RGC dipole angles are spatially uncorrelated. These findings cast doubt on the hypothesis that subcortical structures alone define the layout of OPMs in V1.

## II-85. Cortical maps of ON and OFF visual responses are the substrate of the cortical orientation map

Jens Kremkow  
 Jianzhong Jin  
 Stanley J Komban  
 Yushi Wang  
 Reza Lashgari  
 Jose-Manuel Alonso  
 State University of New York

JENS@KREMKOW.DE  
 JJIN@SUNYOPT.EDU  
 JKOMBAN@SUNYOPT.EDU  
 YUSHIWANG@SUNYOPT.EDU  
 REZALASHGARI@GMAIL.COM  
 JALONSO@SUNYOPT.EDU

ON and OFF visual channels converge for the first time on simple cells in the primary visual cortex. While the orientation preference of an individual simple cell (Lampl et al., 2001) and a cortical column (Jin et al., 2011) can be predicted by the retinotopic arrangement of ON and OFF receptive sub-fields, it is still unclear how ON and OFF retinotopy relates to the topographic organization of cortical orientation maps. A relation between ON/OFF retinotopy and the orientation map has been predicted by several models of cortical development (Miller, 1994; Nakagama et al., 2000; Paik and Ringach, 2011) but never demonstrated. To address this fundamental question, we measured how orientation and ON/OFF retinotopy changes through the horizontal dimension of middle cortical layers. To that end, we tangentially introduced a 32-channel multi-electrode array through the cat primary visual cortex and measured orientation tuning with moving bars and ON/OFF receptive fields with sparse noise stimuli. Consistent with the predictions from computational models, we demonstrate that cortical maps of ON and OFF retinotopy and orientation are closely related. Both the difference between ON and OFF receptive fields (ON-OFFd) and the dipole formed by the retinotopic ON and OFF centers-of-mass (ONc-OFFc dipole) were good predictors of the cortical orientation preference (ON-OFFd:  $r=0.75$ ,  $p<0.001$ , median difference = 16 deg; ONc-OFFc dipole,  $r=0.5$ ,  $p<0.001$ , median difference = 25 deg). Furthermore, the retinotopy of neighboring channels was more scattered for ON than OFF receptive fields in such a way that ON receptive fields seem to 'move around' OFF receptive fields to establish the orientation preference (ONc-to-ONc distance / OFFc-to-OFFc distance = 1.4,  $p<0.001$ ). These results suggest that cortical maps of ON and OFF visual responses are the substrate of the cortical orientation map and that OFF visual responses act as the anchors of cortical retinotopy.

## II-86. V1 laminar-specific activity patterns evoked by receptive field, near and far surround stimulation

Maryam Bijanzadeh<sup>1</sup>  
 Lauri Nurminen<sup>2</sup>  
 Jennifer Ichida<sup>2</sup>  
 Sam Merlin<sup>2</sup>  
 Kenneth Miller<sup>3</sup>  
 Alessandra Angelucci<sup>2</sup>

MA.BIJANZADEH@GMAIL.COM  
 LARSNURMINEN@GMAIL.COM  
 JENNIFER.ICHIDA@HSC.UTAH.EDU  
 SAMMERLIN7@GMAIL.COM  
 KEN@NEUROTHEORY.COLUMBIA.EDU  
 ALESSANDRA.ANGELUCCI@HSC.UTAH.EDU

<sup>1</sup>Moran Eye center

<sup>2</sup>University of Utah

<sup>3</sup>Columbia University

In primary visual cortex (V1), stimuli in the receptive field (RF) surround modulate neuronal responses to stimuli in the RF, a property that may serve figure-ground segregation and/or visual saliency. The neural circuits for surround modulation (SM) are unknown. We previously suggested that SM arises from two distinct and overlapping mechanisms generated by different circuits: near-SM generated by geniculocortical feedforward and intra-V1 horizontal connections, and far-SM by inter-areal feedback. As feedforward, horizontal and feedback circuits terminate in distinct V1 layers, stimuli confined to the RF, near or far surround should evoke distinct laminar patterns of neural activity. We tested this prediction by recording, using linear electrode arrays, local field potentials (LFPs) and spikes across macaque V1 layers to gratings in the RF, near or far surround. Current source density (CSD)

analysis was applied to LFP signals, and the onset latency of CSD responses was estimated at each cortical depth. The earliest current sinks and spikes to RF stimulation were seen in layer (L) 4C rapidly followed by L6. This laminar activity pattern is consistent with thalamic afferents terminating in these layers. Near surround stimulation activated first the infragranular (IG) and then the supragranular (SG) layers and last L4C. This laminar pattern is consistent with near surround signals propagating horizontally to the RF faster within IG than SG layers. These longer latencies in SG layers may reflect the extra synapse that feedforward signals must traverse prior to reaching horizontally-projecting cells in SG layers. Far surround stimulation instead first activated the upper SG and IG layers almost simultaneously and last L4C. This laminar pattern is consistent with fast-conducting FB circuits terminating in SG and IG layers. These results support the hypothesis that RF, near and far SM are generated by different circuits, with horizontal connections critical for near-SM and feedback for far-SM.

## II-87. Stereoscopic surface disambiguation and interpolation in the macaque primary visual cortex

Jason M Samonds<sup>1</sup>  
 Brian Potetz<sup>2</sup>  
 Christopher Tyler<sup>3</sup>  
 Tai Sing Lee<sup>1</sup>

SAMONDJM@CNBC.CMU.EDU  
 POTETZ@GOOGLE.COM  
 CWT@SKI.ORG  
 TAI@CNBC.CMU.EDU

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>Google

<sup>3</sup>Smith-Kettlewell Eye Research Institute

Binocular disparity tuning measured in the primary visual cortex (V1) is described well by the disparity energy model, but not all aspects of disparity tuning are fully explained by the model. We developed a model that includes recurrent connections inferred from neurophysiological observations on spike timing correlations (Samonds et al., J Neurosci 2009) that can explain existing data on disparity tuning dynamics. This model also predicted sharper disparity tuning for larger stimuli and reduced sharpening and strength of inverted disparity tuning to anti-correlated stereograms, where dots of opposite luminance polarity are matched between the left- and right-eye images. We tested these predictions for neurophysiological recordings in macaque V1 and the observed dynamics of disparity tuning matched our model simulations. Such recurrent interactions among disparity-tuned neurons in V1 could play a significant role in cooperative stereoscopic computations by altering responses to favor the most likely depth interpretation for a given image pair. Psychophysical research has shown that binocular disparity stimuli displayed in one region of the visual field can be interpolated into neighboring regions that contain ambiguous depth information. Guided by these studies, we generated stimuli to test whether neurons in V1 interact in a similar manner. We found that unambiguous binocular disparity stimuli displayed in the surround indeed modified the responses of neurons when either ambiguous binocular or uniform featureless stimuli were presented within their receptive fields. The delayed timing of the behavior suggests that recurrent (lateral or feedback) interactions among V1 neurons play an important role in carrying out cooperative stereoscopic computations.

## II-88. Visuotopic anisotropy in functional connectivity among local cortical sites in human visual areas

Jungwon Ryu  
 Sang-Hun Lee  
 Seoul National University

RJUNGWON@SNU.AC.KR  
 VISIONSL@SNU.AC.KR

Recent fMRI studies repetitively reported that low-frequency spontaneous cortical activity fluctuates in synchrony among anatomically connected cortical areas. Inspired by this robust relationship, we searched for visuotopic

anisotropy in intrinsic connectivity that possibly subserves corresponding anisotropies in human vision by detailing the structure of functional connectivity between individual voxels in human visual cortex. To define visuotopic coordinates of individual voxels, their population receptive fields (pRFs) were mapped onto the visual field. Then, we acquired fMRI measurements under two resting-state ('eye-closed' and 'fixation with gray background') and two stimulation-state (periodic stimulation of 'orientation-filtered' and 'spatial frequency-filtered' images of white-noise patterns) conditions. To quantify intrinsic connectivity from these raw measurements, we extracted the moment-to-moment fluctuations that were not accounted for either by external stimulation, by trivial MR artifacts, or by non-neuronal physiological responses ('noise time series (nTS)'). Having defined pRFs and nTSs for individual cortical voxels, we examined how correlations in nTS change as a function of visuotopic distance and geometrical relationship of given voxel pairs. As expected, the correlation decreased monotonically with an increasing visuotopic distance. Intriguingly, the correlation, when visuotopic distance is controlled, was preferentially higher for voxel pairs whose pRFs belong to the geometry of concentric rings or radial lines in visuotopic space. Furthermore, what distinguished the two anisotropies was the eccentricity of voxel pairs' pRFs: The concentric anisotropy was dominant in the fovea ( $< 2$  degree) region whereas the radial anisotropy was dominant in the periphery ( $> 4$  degree). These eccentricity-dependent anisotropies in nTS correlation were invariant to the cognitive states and stimulus types that varied across the different viewing conditions. This robust anisotropy indicates the bias in intrinsic connectivity at the early visual cortex.

## II-89. Visualizing the similarity and pedigree of neuronal ion channel models for NEURON.

William Podlaski<sup>1</sup>  
 Rajnish Ranjan<sup>1</sup>  
 Alexander Seeholzer<sup>2</sup>  
 Henry Markram<sup>1</sup>  
 Wulfram Gerstner<sup>1</sup>  
 Tim Vogels<sup>3,1</sup>

WILLIAM.PODLASKI@EPFL.CH  
 RANJAN.RAJNISH@EPFL.CH  
 ALEX.SEEHOLZER@EPFL.CH  
 HENRY.MARKRAM@EPFL.CH  
 WULFRAM.GERSTNER@EPFL.CH  
 TIM.VOGELS@CNGB.OX.AC.UK

<sup>1</sup>Ecole Polytechnique Federale de Lausanne

<sup>2</sup>LCN, EPF Lausanne

<sup>3</sup>University of Oxford

Building biologically accurate neuron models often requires the use of ion channel dynamics. While there is an abundance of ion channel models in the literature, many being available through internet services like modelDB (<http://senselab.med.yale.edu/ModelDB>), it can be quite difficult to determine which channel models are most appropriate for a specific neuronal simulation at hand. Furthermore, their family relations, i.e., how a model was derived or inspired from another ion channel model, are often opaque. To address these issues we have performed an exhaustive literature search and created a buffet-style database that lists all available reference information for each of the over 2000 available channel models. We established the presumed channel type as well as its pedigree. Additionally, we tested the physiological performance of each channel in standardized voltage clamp protocols from which we derived a benchmark similarity measure between model channels as well as similarity between model output and experimentally obtained recordings. We visualize family relations between various channels by creating genealogical trees and by using clustering methods. We will provide easy access to this information through inclusion in modelDB and will provide protocols to test and integrate future models into our database. Our database will become an invaluable resource for experimentalists and modelers alike, because it brings order to the existing jungle of available ion channel models and allows for maintenance, documentation, and easy access of currently existing models, and, importantly, the integration and evaluation of new ion channel models in the future.

## II-90. Scalable semi-supervised framework for activity mapping of the leech nervous system

Edward Frady  
William Kristan  
University of California, San Diego

EFRADY@UCSD.EDU  
WKKRISTAN@UCSD.EDU

Technical progress in imaging and multi-unit recordings have enabled the recording of hundreds to thousands of neurons simultaneously during relevant behavioral states. Advances in voltage-sensitive dyes allow us to record a significant fraction of the neurons in the leech nervous system with unprecedented sensitivity and temporal resolution (Miller et al., 2012). We recorded from hundreds of neurons (~80% of neurons on ventral surface of the ganglion) during two behaviors – swimming and shortening, and utilize anatomical and functional properties of the neurons to piece together a canonical activity map of the leech ganglion. With these new techniques, the ability to obtain a holistic view of the nervous system is tantalizingly close, but making sense of all of this data presents an entirely new challenge. Here we present a semi-supervised framework for developing an activity map of the leech nervous system. Our approach incorporates computational and machine-learning techniques that can scale to handle extremely large data sets, but retains the sophistication of manual analysis. The framework consists of 3 stages: 1.) Automated extraction of cellular signals using PCA-ICA (Mukamel et al., 2009; Hill et al., 2010) and manual removal of artifacts. 2.) Abstraction of component signals into a useful high-dimensional feature space with learned and manually designed features. 3.) Clustering and bi-partite graph matching with low-confidence manual intervention to identify homologous cells across animals. With these techniques, we have identified dozens of previously uncharacterized neurons, and mapped the activity of a large fraction of all neurons during multiple behaviors.

## II-91. The cellular origin of neuronal avalanches

Tim Bellay<sup>1,2</sup>  
Andreas Klaus<sup>3</sup>  
Saurav Seshadri<sup>4</sup>  
Dietmar Plenz<sup>3</sup>

BELLAYT@MAIL.NIH.GOV  
AKLAUS.MAIL@GMAIL.COM  
SAURAV.SESHADRI@NIH.GOV  
PLENZD@MAIL.NIH.GOV

<sup>1</sup>National Institutes of Health

<sup>2</sup>Brown University

<sup>3</sup>National Institute of Mental Health/NIH

<sup>4</sup>National Institute of Mental Health

Resting activity in the cortex of mammals is composed of neuronal avalanches, a scale-free organization of intermittent neuronal cascades. Theory and experiments have shown that networks with neuronal avalanches maximize numerous aspects of information processing such as dynamic range and memory capacity. Avalanches have been measured in the spontaneous local field potential of cortex in vitro, in rodents and non-human primates in vivo, and non-invasively in humans using magnetoencephalography and functional magnetoresonance imaging. Despite such large-scale and robust measurements, the cellular origin of neuronal avalanches remains unknown. Here we selectively expressed the genetically encoded calcium indicator (GECI) YC2.6 in pyramidal cells of rats using in utero electroporation (E15.5). We imaged spiking activity in L2/3 pyramidal neurons from superficial cortical layers under awake head-restrained conditions. At the single neuron level, instantaneous firing rates and interburst intervals followed Poissonian statistics, in line with numerous reports on irregular spiking of cortical neurons in vivo. In contrast, spiking activity of neuronal groups emerged as neuronal avalanches demonstrating a precise spatio-temporal organization of neuronal group activity and balanced propagation within superficial layers (Beggs & Plenz 2003). Deviations from awake, ongoing dynamics induced by isoflurane anesthesia destroyed spike avalanches. We also demonstrate that spike avalanches from pyramidal neurons can arise intrinsically from within cortex using a reduced cortical slice-culture preparation. Our results for the first time locate the origin of neuronal avalanches to groups of pyramidal neurons in superficial cortex. Our findings of Poisson-like statistics

at the single neuron level and the emergence of avalanche dynamics at the group level should constrain cortex models on resting state dynamics and inform mechanisms on balanced propagation that are disrupted in disease states.

## **II-92. Microcircuit activity is patterned topologically and reveals features of underlying connectivity**

Brendan Chambers  
Alexander Sadoovsky  
Jason MacLean  
University of Chicago

BRENDANCHAMBERS@UCHICAGO.EDU  
SADOVSKY@UCHICAGO.EDU  
JMACLEAN@UCHICAGO.EDU

Patterned circuit activity has been shown to encode sensory inputs (Luczak et al 2007), motor outputs (Churchland et al 2012) and behavioral choice (Harvey et al 2012). However it remains unclear how synaptic connectivity gives rise to structured microcircuit dynamics. Here we investigate the relationship between activity flow through cortical microcircuitry and its relationship to synaptic connectivity, using a multidisciplinary approach. We generated functional wiring diagrams from spiking activity in local populations of up to 1000 neurons spanning multiple layers and columns, recorded using 2-photon Ca<sup>2+</sup>-imaging in mouse primary sensory cortex. We inferred pairwise functional relationships using an empirical Bayes method (Fig 1). Null expectations were established using a population of Poisson neurons matched to the firing rate statistics of each dataset. Compared to our experimental results, nulls did not feature strong connectivity, long continuous paths, or high-degree nodes (Fig 2). Thus, local circuit activity contains significant patterns of reliable sequential firing beyond what is expected by chance alone, reflecting underlying connectivity. We hypothesized that pairwise lagged firing reflects strong underlying synaptic connectivity although we did not predict that every functional edge corresponds to a synaptic connection (Gerstein et al., 1978). A spiking conductance-based network model was used to explore the relationship between synaptic versus functional connectivity. To mimic experimental sampling constraints on observability and temporal precision, 60% of cells in the model were occluded and its output spiking was down-sampled. Mapping the output of these model networks, we found that reliable sequential firing corresponds to underlying synaptic connectivity more frequently than expected by chance, especially for the strongest synaptic connections. A second category of functional links corresponds to indirect connectivity, via multiple synapses. We note that reliable identification of synaptic connectivity is strongly related to input structure. These results establish heuristics for mapping microcircuit dynamics onto underlying synaptic architectures.

## **II-93. Scalable ROI detection for calcium imaging**

David Pfau<sup>1</sup>  
Jeremy Freeman<sup>2,3</sup>  
Misha Ahrens<sup>2,3</sup>  
Liam Paninski<sup>1</sup>

DAVID.PFAU@GMAIL.COM  
FREEMANJ11@JANELIA.HHMI.ORG  
AHRENSM@JANELIA.HHMI.ORG  
LIAM@STAT.COLUMBIA.EDU

<sup>1</sup>Columbia University

<sup>2</sup>Janelia Farm Research Campus

<sup>3</sup>Howard Hughes Medical Institute

The latest advances in calcium imaging make it possible to monitor responses in tens of thousands of neurons simultaneously. At this scale, it is impractical to hand-label every neuron, particularly for high throughput experiments. Existing methods for automatic region of interest (ROI) detection either ignore most spatial structure, making it difficult to scale to large data sets, or ignore temporal information. We develop a method for automated cell body detection that applies sparse PCA to local patches and is easily adapted to distributed computing architectures. We present results on light sheet microscopy imaging of the larval zebrafish expressing GCaMP5

pan-neuronally. We are able to find the full shape of active neurons, including neuronal processes in highly active cells. Some faint, sparsely firing neurons are missed by the analysis, but almost all are at the edge of visibility. We also present results on GCaMP6-labeled neurons in mouse somatosensory cortex. The analysis suggests that patch-wise sparse PCA is an efficient and effective solution for ROI detection in a variety of fluorescent imaging data despite the differences in imaging methods and animal physiology.

## II-94. Investigating mechanisms of drug action at the hippocampal ensemble level in behaving mice.

Tamara Berdyjeva<sup>1</sup>  
Stephani Otte<sup>2</sup>  
Leah Aluisio<sup>1</sup>  
Laurie Burns<sup>2</sup>  
Yaniv Ziv<sup>3</sup>  
Christine Dugovic<sup>1</sup>  
Kunal Gosh<sup>2</sup>  
Mark J Schnitzer<sup>3</sup>  
Tim Lovenberg<sup>1</sup>  
Pascal Bonaventure<sup>1</sup>

TBERDYY6@ITS.JNJ.COM  
STEPHANI@INSCOPIX.COM  
LALUISIO@ITS.JNJ.COM  
LAURIE@INSCOPIX.COM  
YZIV@STANFORD.EDU  
CDUGOVIC@ITS.JNJ.COM  
KUNAL@INSCOPIX.COM  
MSCHNITZER@GMAIL.COM  
TLOVENBE@ITS.JNJ.COM  
PBONAVE1@ITS.JNJ.COM

<sup>1</sup>Janssen LLC

<sup>2</sup>Inscopix

<sup>3</sup>Stanford University

The miniature, integrated fluorescence microscope permits imaging studies of multiple neurons' individual dynamics in deep brain structures in freely behaving animals. Here, we demonstrate the use of this technology in drug discovery by tracking somatic calcium dynamics in hundreds (200-800 per animal) of hippocampal neurons of pharmacologically manipulated behaving mice. At the same time, we monitored EEG, EMG, temperature, and locomotor activity to control for common confounding factors of drug action and reduce ambiguity in the interpretation of observed drug effects. To demonstrate the approach, we studied the effects of the sleep-inducing drug Zolpidem. We found that Zolpidem strongly suppresses neuronal activity in hippocampus: the frequency of calcium transients (Fr) was significantly lower relative to the vehicle control (average Zolpidem Fr/vehicle Fr = 0.27, sdv = 0.04, n = 5 mice). While Zolpidem (GABAA agonist) is a commonly prescribed sleep medication, it has adverse memory-related side effects. In contrast, an Orexin receptor antagonist promotes sleep through different mechanisms (decreasing the output of active neurons) and has fewer memory side effects in preclinical species. We hypothesized that, unlike Zolpidem, Orexin antagonist would not lower hippocampal neuronal activity below physiological level. To rule out the differences in drugs' effects that are due to a potentially differing physiological states (i.e., REM/NREM ratios), we compared neuronal activity in the epochs matched by physiological states identified by multimodal recordings. Therapeutic drugs for cognitive and psychiatric disorders are often characterized by their mechanism of action at the molecular levels. Here we provide a proof-of-principle for a new approach to understand mechanism of drug action in terms of information processing in large neuronal ensembles in relevant brain circuits. Understanding mechanisms of action through the neural coding perspective may provide explanations to differences in drug efficacy, kinetics, and side effects that traditional techniques have come short to reveal.

## II-95. Phase coding in posterior parietal cortex predicts neural suppression in a gaze-anchoring task

Maureen Hagan  
Bijan Pesaran  
New York University

HAGAN@CNS.NYU.EDU  
BIJAN@NYU.EDU

Reach accuracy greatly benefits from saccadic eye movements that foveate the reach target. Double-step tasks in humans in which subjects are asked to make a reach and saccade followed closely in time with a second saccade to another target find that gaze is anchored to the first target until the reach is completed and reaction times for the second saccade were significantly longer. This suggests that there is saccadic suppression around the time of reach completion that ensures the eye will continue to foveate the reach target, maximizing reach accuracy. We trained subjects to perform a similar double-step task as well as a task involving two saccades and no reach. As in human studies, we find latencies for the second saccade are significantly longer if the second saccade target appears within 300 ms of reach completion. Recordings from the posterior parietal cortex (PPC) find that neural activity in area LIP, an area associated with guiding saccadic eye movements, is suppressed around the time of the reach and shows a suppressed response to saccade targets in the neurons' response field when the target appears close in time to the completion of the reach. Furthermore, the latency of the second saccade immediately after a reach can be predicted by the firing rate of area LIP neurons on a trial-by-trial basis. Neurons in the Parietal Reach Region (PRR) that are tuned for the reaching movement preferentially fire in phase with the beta-band of the local field potential (LFP) in area LIP prior to the reach. Around the timing of the reach, the spiking activity becomes out of phase with area LIP and may act as a gating mechanism for the suppression of activity in area LIP. These results suggest a role for the PPC in the neural mechanisms of 'gaze anchoring.'

## II-96. PV+ inhibition regulates feedforward activity in auditory cortex.

Alexandra Moore  
Michael Wehr  
University of Oregon

HARTMAN@UOREGON.EDU  
WEHR@UOREGON.EDU

In auditory cortex, inhibition is known to shape stimulus selectivity and regulate network excitability. Of the various inhibitory cell types present in the cortex, those belonging to the parvalbumin-positive (PV+) subclass are particularly well-positioned to regulate the spiking output of pyramidal cells. Based on our recent finding that PV+ neurons in auditory cortex are well-tuned for frequency, we hypothesized that PV+-mediated inhibition serves chiefly to regulate the gain and timing of pyramidal responses, and has only a minor impact on tuning. To test this hypothesis, we suppressed PV+ neurons using Arch in transgenic mice, and compared the spiking responses of single pyramidal neurons with and without PV+ inhibition. Suppressing PV+ activity increased the responsiveness of 80% of recorded neurons, primarily imparting a multiplicative gain increase (output gain control), but in most cases also producing an additive shift (input gain control). We initially hypothesized that the synaptic change underlying this effect was a direct decrease in inhibitory synaptic input to pyramidal neurons, which would allow existing excitatory input, then 'unchecked', to be converted into a larger spiking response. To test this, we used whole-cell voltage clamp methods to measure the effect of PV+ suppression on excitatory and inhibitory synaptic inputs to pyramidal neurons. Surprisingly, PV+ suppression increased both inhibition and excitation in recorded neurons, indicating that the effects of PV+ gain control propagate through a feedforward network: increases in spiking output are in most cases caused by an increase in recurrent excitatory input — a powerful, indirect consequence of silencing PV+ interneurons. We propose that by regulating the output gain of a subset of principal cells, PV+ neurons serve as indirect regulators of input gain for most of the cortical network, dampening the spread of activity through intracortical connections.



## II-97. Observing the unobservable: a note on stop-signal reaction-time distributions

Tobias Teichert  
Vincent Ferrera  
University of Pittsburgh

TEICHERT@PITT.EDU  
VPF3@COLUMBIA.EDU

Inhibitory control is an important component of executive function that allows organisms to abort behavioral plans on the fly as new sensory information becomes available. Current models treat inhibitory control as a race between Go- and a Stop processes that are mediated by partially distinct neural substrates, i.e., the direct and the hyper-direct pathway of the basal ganglia. The fact that finishing times of the Stop process (Stop-Signal Reaction Time, SSRT) cannot be observed directly has precluded a precise comparison of the functional properties that govern the initiation (GoRT) and inhibition (SSRT) of a motor response. To solve this problem, we developed a non-parametric framework to measure the trial-by-trial distribution of SSRT. A series of simulations verified that the non-parametric approach yields accurate estimates of the entire SSRT distribution from as little as 750 trials. Our results show that in identical settings, the distribution of SSRT is very similar to the distribution of GoRT albeit somewhat shorter, wider and significantly less right-skewed. The ability to measure the precise shape of SSRT distributions opens new avenues for research into the functional properties of the hyper-direct pathway that is believed to mediate inhibitory control.

## II-98. A simplified Bayesian model accounts for procedure-related variability in perceptual thresholds

Ofri Raviv  
Yonatan Loewenstein  
Merav Ahissar  
Hebrew University

OFRI.RAVIV@MAIL.HUJI.AC.IL  
YONATAN.LOEWENSTEIN@MAIL.HUJI.AC.IL  
MSMERAVA@PLUTO.MSCC.HUJI.AC.IL

Performance in many psychophysical tasks varies with seemingly insignificant variations of the protocol. Specifically, in two alternative forced choice discrimination tasks, the two stimuli to be discriminated in each trial may be drawn from a wide range (a protocol referred to as No-Reference), or alternatively, they can be chosen such that one of the stimuli is always a constant 'reference' stimulus. When a reference stimulus is used, it can appear always in the first interval of each trial (Reference-First protocol), it may be randomly placed (Reference protocol), or as most commonly practiced, it can appear in a random interval, and always be lower than the comparison stimulus (Reference-Lower protocol). These protocols consistently yield different results, in a way that is not consistent with the classical signal detection theory account of perception. We present a simple model in which observers incorporate prior knowledge about stimuli statistics into their decision. This model can be viewed as an approximation of a Bayesian optimal strategy to overcome perceptual and memory noises, but instead of representing the full prior distribution, observers only track its mean. This simple model accounts for the results of a wide range of variations in the statistics of the stimuli: the performance benefit from the existence of a reference, compared to the No-Reference protocol; for performance in Reference-First being better than in the other protocols; and for interval biases favoring opposite trial types in the Reference and Reference-Lower protocols. These findings suggest that the commonly used psychophysical paradigms do not measure perceptual noise alone, but are greatly influenced by the sequence of preceding stimuli. Moreover, the larger effect of the first stimulus of previous trials, compared to the second, more recent one, constrains the underlying neural mechanism, and suggests a specific marker for detecting it.

## II-99. A behavioral tracking paradigm for estimating visual sensitivity with dynamic internal models

Kathryn Bonnen  
 Johannes Burge  
 Jacob L Yates  
 Jonathan W Pillow  
 Lawrence K Cormack

University of Texas at Austin

KATHRYN.BONNEN@UTEXAS.EDU  
 JBURGE@MAIL.CPS.UTEXAS.EDU  
 JLYATES@UTEXAS.EDU  
 PILLOW@MAIL.UTEXAS.EDU  
 CORMACK@UTEXAS.EDU

Visual sensitivity is a fundamental quantity that reveals much about the neural mechanisms underlying vision. In traditional psychophysics, sensitivity is measured using time-intensive, trial-based experiments: targets are presented at various physical strengths, and an observer indicates whether (or where) a target is present. After many such trials, sensitivity can be computed and compared to that of an ideal (or biologically inspired) observer model. However, the time required to collect enough data to measure sensitivity is often severely limiting. Here we introduce a novel framework for estimating visual sensitivity using a continuous target-tracking task in concert with a dynamic internal model of human visual performance. Real-time tracking is a fun, easy, ecologically relevant task that rapidly generates large quantities of data. In our experiments, observers tracked the center of a 2D Gaussian luminance blob with a mouse. The blob moved in a random walk in a field of dynamic additive Gaussian luminance noise. Each frame can be considered a 'micro-trial' in which the observer indicates that they can see the blob in its new position by moving the mouse in the corresponding direction; a mere five seconds of tracking (say) thus yields 300 data points. To estimate sensitivity, we analyze tracking using a dynamic linear model (a Kalman filter). This model assumes that observers behave as Bayesian ideal observers, combining prior beliefs with noisy measurements estimate target position over time. We derive maximum likelihood estimates of the Kalman filter parameters from tracking data. This model provides a good description of tracking performance, and its parameters can be directly related to visual sensitivity ( $d'$  for example). Together, this framework can drastically reduce the amount of time (and boredom) required to assess certain aspects of visual function.

## II-100. Modeling and exploiting co-adaptation in brain-computer interfaces

Josh Merel<sup>1</sup>  
 Antony Qian<sup>1</sup>  
 Roy Fox<sup>2</sup>  
 Liam Paninski<sup>1</sup>

<sup>1</sup>Columbia University

<sup>2</sup>Hebrew University

JSM2183@COLUMBIA.EDU  
 AJQ2103@COLUMBIA.EDU  
 ROYF@CS.HUJI.AC.IL  
 LIAM@STAT.COLUMBIA.EDU

In a closed-loop brain-computer interface (BCI), an adaptive decoder is used to learn parameters for decoding the user's neural response in order to control a prosthesis. Sensory feedback to the user provides information which permits the neural tuning to adapt in a manner which can improve control. The process of joint optimization that occurs dynamically as the adaptive decoder adjusts parameters and the user adapts their control scheme is called co-adaptation. We make use of a simulated BCI task called online prosthesis simulation (OPS) to evaluate a model of the co-adaptation process which we developed, and attempt to exploit this model to improve the learnability of the BCI system. In previous work, we have related the co-adaptation problem to a multi-agent formulation of the linear quadratic Gaussian control problem. In simulation, we demonstrated an algorithm which could improve the learnability of the BCI. Here we make use of an OPS to test our model in a human psychophysics-style task which captures important features of the BCI setting. We seek to characterize how well our framework captures the co-adaptation process as well as constraints on the changes the user can make to their control scheme. Cursor control experiments using the OPS allow us to examine rates at which the user learns as well as the task circumstances in which learning by the user meaningfully contributes to decreases in task error. We also seek to assess the extent to which the co-adaptation which occurs naturally in the presence of an adaptive decoder can

be improved by using a decoder update rule which anticipates the user adaptation. The goals of this project are both to understand scientifically the co-adaptation phenomenon and to exploit it algorithmically, enabling more well-performing BCI applications.

### III-1. Behavioral and neural evidence that mice infer probabilistic models for timing

Yi Li<sup>1,2</sup>  
Joshua T Dudman<sup>1,2</sup>

LIY11@JANELIA.HHMI.ORG  
DUDMANJ@JANELIA.HHMI.ORG

<sup>1</sup>Janelia Farm Research Campus  
<sup>2</sup>Howard Hughes Medical Institute

Animals learn both whether and when a reward will occur. Neural models of timing posit that animals learn the mean time until reward perturbed by a fixed relative uncertainty. By contrast, optimal decisions require a subject to infer probability distributions that incorporate arbitrary uncertainty. Although animals can successfully forage for rewards in highly variable environments, it remains unclear whether mice can infer quantitative, probabilistic models of reward timing. We describe a novel behavioral paradigm in which mice must adapt their waiting time on unrewarded trials to the standard deviation of the reward delay distribution to maximize reward rate. We find that mice are able to behave optimally even in the presence of substantial variation ( $CV > 0.6$ ). We show that a model in which tens of trials of reward timing information is used to infer the standard deviation best described the behavioral data. Moreover, mice adapt their behavior to multiple features of the reward delay distribution consistent with the inference of a probabilistic model of reward timing. The dorsomedial striatum (DMS) is thought to be critical for the initiation and appropriate timing of voluntary behaviors. We found that inactivation of the DMS abrogated task performance. Thus, to further investigate whether and how DMS is involved in processing highly variable reward timing beyond the initiation of reward seeking actions, we implanted electrode arrays into both hemispheres of the DMS. We find that activity in the DMS peaked when animal initiated and terminated trials. Surprisingly, there was little sustained or ramping activity during the delay period as mice waited for reward delivery despite robust timing behavior. We suggest that these data are consistent with a model in which the DMS contributes to the updating and reading out of expected reward timing information by corticostriatal loops.

### III-2. High-frequency activity in the medial orbitofrontal cortex encodes positive outcomes

Ashwin Ramayya<sup>1</sup>  
Ashwini Sharan<sup>2</sup>  
James Evans<sup>2</sup>  
Michael Sperling<sup>2</sup>  
Timothy Lucas<sup>1</sup>  
Michael Kahana<sup>1</sup>

ASHWINRAMAYYA@GMAIL.COM  
ASHWINI.SHARAN@JEFFERSON.EDU  
JAMES.EVANS@JEFFERSON.EDU  
MICHAEL.SPRLING@JEFFERSON.EDU  
TIMOTHY.LUCAS@UPHS.UPENN.EDU  
KAHANA@PSYCH.UPENN.EDU

<sup>1</sup>University of Pennsylvania  
<sup>2</sup>Thomas Jefferson University

Both human functional magnetic resonance (fMRI) and non-human primate electrophysiology studies implicate the orbitofrontal cortex (OFC) in representing the value of stimuli in the environment [3]. However, there exists a discrepancy between these studies—whereas human fMRI studies have typically shown value-related signals in medial OFC [1], non-human primate electrophysiology studies typically identify value-related activity from lateral OFC neurons, rather than medial OFC neurons [3]. Part of this discrepancy may be due to methodological limitations of fMRI studies, particularly that they have reduced sensitivity in lateral OFC regions which overlay the nasal sinuses. Here, we obtained intracranial electroencephalography (iEEG) from patients undergoing monitoring for

intractable epilepsy as they performed a two–alternative probability learning task. We focused our analyses on changes in high-frequency activity (HFA, 60–200 Hz power) a neural signal which has shown to be positively correlated with BOLD activity and neuronal spiking [2], but does not suffer from reduced sensitivity in lateral OFC regions. We found that HFA was elevated following positive compared to negative outcomes in the medial OFC, but not the lateral OFC, or neighboring prefrontal regions, suggesting that the human medial OFC plays a specialized role in outcome valuation. Thus, the discrepancy between human fMRI and animal electrophysiology studies cannot solely be explained by methodological differences, but rather may reflect a true cross-species difference.

### III-3. Altered neuronal avalanches at the population and cell level in a rat model of schizophrenia

Saurav Seshadri<sup>1</sup>  
Tim Bellay<sup>2,3</sup>  
Samantha S Chou<sup>1</sup>  
Dietmar Plenz<sup>4</sup>

SAURAV.SESHADRI@NIH.GOV  
BELLAYT@MAIL.NIH.GOV  
SAMANTHA.CHOU@NIH.GOV  
PLENZD@MAIL.NIH.GOV

<sup>1</sup>National Institute of Mental Health

<sup>2</sup>National Institutes of Health

<sup>3</sup>Brown University

<sup>4</sup>National Institute of Mental Health/NIH

Schizophrenia (SZ) is a psychiatric disorder whose pathophysiology is thought to involve impaired cortical balance of excitation and inhibition (E/I), as well as altered dopaminergic signaling, which may underlie cognitive symptoms such as working memory deficits. The resting state activity of cortex in humans, non-human primates, and rodents has been shown to organize as neuronal avalanches. Avalanche dynamics optimize numerous aspects of information processing, such as information capacity and dynamic range; are highly sensitive to the E/I balance; and exhibit an inverted U-shaped profile of dopaminergic dependency similar to that identified for working memory function. We therefore hypothesized that network-level disturbances in SZ may correlate with deviations from neuronal avalanche dynamics. To address this, we used the neonatal phencyclidine (PCP) model of SZ. Rat pups were injected with PCP (s.c., 10 mg/kg) at postnatal day 7, 9, and 11, and subjected to behavioral testing at 4-6 weeks. In confirmatory behavioral tests, PCP-treated rats showed hyperlocomotion in response to a PCP challenge and impairment in novel object recognition. Avalanches were first studied at the population level by recording the local field potential in frontal cortex of adult rats under light anesthesia (1% isoflurane) using a microelectrode array (A8x4, NeuroNexus). Ongoing LFP activity deviated from criticality, exhibiting dynamics consistent with a supercritical regime of propagation. To confirm these observations at the cellular level, rats were in utero electroporated with the genetically encoded Ca<sup>2+</sup> indicator YC2.60 prior to birth and drug treatment, and subsequently fitted with a head bar and cranial window to enable 2-photon imaging of the cortical pyramidal neurons. This allows us to continuously observe the activity of individual neurons in the network, with the aim of confirming our LFP results. These results suggest that studying disruption of neuronal avalanches in disease states may be a translatable tool for psychiatric disorders.

### III-4. Modeling attention-induced drop of noise correlations by inhibitory feedback

Tatjana Wiese<sup>1</sup>  
Marlene Cohen<sup>2</sup>  
Brent Doiron<sup>2</sup>

TATJANA42@GMAIL.COM  
COHENM@PITT.EDU  
BRENT.DOIRON@GMAIL.COM

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>University of Pittsburgh

An important challenge in systems neuroscience is to understand the mechanisms underlying cognitive processes like attention. While many studies report attentional effects on the activity of individual neurons and neural populations, how neural circuits give rise to these effects is not yet understood. It is well established that attention increases trial-averaged firing rates and stimulus response gains in primate visual cortex (V4) [McAdams and Maunsell 1999]. Furthermore, recent population recordings demonstrate that attention decreases trial-to-trial covariability (i.e. 'noise correlations') [Cohen and Maunsell 2009]. We model attention as a top-down modulation of excitability of recurrently coupled excitatory and inhibitory neurons. Using a stochastic mean-field rate model, we show that under certain conditions, such a modulation can cause an increase in excitatory and inhibitory neuron firing rates and response gains, and a concurrent decrease in excitatory neuron covariability. Specifically, attention can decorrelate activity of neuron pairs over trials by recruiting inhibition to bring the network into a more stable state, consistent with recent theories of surround suppression [Ozeki et al. 2009, Tsodyks et al. 1997]. The increased stability attenuates the transfer of correlated input fluctuations to correlated output firing activity. We further demonstrate that this mechanism works not only for mean-field models, but also for spiking networks of leaky integrate-and-fire neurons. Our theories give clear predictions on how top-down and bottom-up inputs are distributed over excitatory and inhibitory populations in visual cortex (V4).

### III-5. A single net excitation strength measure broadens the scope of the 2-layer model

Lei Jin  
Bartlett W. Mel

LEIJ@USC.EDU  
MEL@USC.EDU

University of Southern California

A crucial step in understanding cortical function is to develop simplified models of individual neurons that capture their core computational capabilities while suppressing superfluous biophysical details. We previously showed that the input-output behavior of a dendritic subtree whose branches emanate from a main trunk or soma can be described by an abstract 2-layer model (2LM), where the first layer is composed of a set of independent dendritic 'subunits' with multi-dimensional sigmoidal input-output functions, and the second layer, corresponding to the soma, sums the first-layer outputs and feeds the resulting 'current' through the cell's axo-somatic  $f$ - $I$  curve. Despite the extreme simplification of a real neuron that it represents, the 2LM can predict the average firing rate of a biophysically-detailed compartmental model with remarkable accuracy. Our current 2LM uses spatial basis functions representing excitation density along a dendrite to predict the time-averaged response of the dendrite to an arbitrary spatial pattern of excitation. The scheme is limited, however, not yet addressing the fact that excitation strength at a given dendritic location has at least 3 components: (1) the number of spines activated  $N$ ; (2) the frequency of pre-synaptic firing  $F$ , and (3) the peak synaptic conductance  $G$ . How do these 3 types of strength interact over their physiological ranges? In this work, we used compartmental simulations to test — and validate — the hypothesis that these 3 strength components can be collapsed into a single net excitation strength variable which can be used as a generic excitatory input to the 2LM. This modularization of excitation strength — essentially a form of 'data hiding' — is valuable in that, when prepended to our existing 2LM, significantly broadens the parameter space over which accurate predictions of average spike rates can be generated essentially by hand calculations and table lookups.

## III-6. Low-dimensional models of neural population recordings with complex stimulus selectivity

Evan Archer<sup>1,2</sup>  
Jonathan W Pillow<sup>2</sup>  
Jakob Macke<sup>1,3</sup>

EARCHER@UTEXAS.EDU  
PILLOW@MAIL.UTEXAS.EDU  
JAKOB.MACKE@GMAIL.COM

<sup>1</sup>Max Planck Institute

<sup>2</sup>University of Texas at Austin

<sup>3</sup>Bernstein Center for Computational Neuroscience, Tuebingen

Modern experimental technologies such as multi-electrode arrays and 2-photon population calcium imaging make it possible to record the responses of large neural populations (up to 100s of neurons) simultaneously. These high-dimensional data pose a significant challenge for analysis. Recent work has focused on extracting low-dimensional dynamical trajectories that may underlie such responses. These methods enable visualization of high-dimensional neural activity, and may also provide insight into the function of underlying circuitry. Previous work, however, has primarily focused on models of a population's intrinsic dynamics, without taking into account any external stimulus drive. We propose a new technique that integrates linear dimensionality reduction of stimulus-response functions (analogous to spike-triggered average and covariance analysis) with a latent dynamical system (LDS) model of neural activity. Under our model, the population response is governed by a low-dimensional dynamical system with nonlinear (quadratic) stimulus-dependent input. Parameters of the model can be learned by combining standard expectation maximization for linear dynamical system models with a recently proposed algorithms for learning quadratic feature selectivity. Unlike models with all-to-all connectivity, this framework scales well to large populations since, given fixed latent dimensionality, the number of parameters grows linearly with population size. Simultaneous modeling of dynamics and stimulus dependence allows our method to model correlations in response variability while also uncovering low-dimensional stimulus selectivity that is shared across a population. Because stimulus selectivity and noise correlations both arise from coupling to the underlying dynamical system, it is particularly well-suited for studying the neural population activity of sensory cortices, where stimulus inputs received by different neurons are likely to be mediated by local circuitry, giving rise to both shared dynamics and substantial receptive field overlap.

## III-7. Scalable nonparametric models for binary spike patterns

Il Memming Park<sup>1</sup>  
Evan Archer<sup>2,1</sup>  
Kenneth Latimer<sup>1</sup>  
Jonathan W Pillow<sup>1</sup>

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

<sup>1</sup>University of Texas at Austin

<sup>2</sup>Max Planck Institute

Probabilistic models for binary spike patterns provide a powerful tool for understanding statistical dependencies in large-scale neural recordings. Maximum entropy (or "maxent") models, which seek to explain dependencies in terms of low-order interactions between neurons, have enjoyed remarkable success in modeling such patterns, particularly for small groups of neurons. However, these models are not computationally tractable for large populations, and low-order maxent models have been shown to be inadequate for some datasets. To overcome these limitations, we propose a family of "universal" models for binary spike patterns, where universality refers to the model's ability to describe arbitrary distributions over all  $2^m$  binary patterns. Universal binary models consist of a well-behaved parametric base distribution; they combine the flexibility of a histogram with the parsimony of a parametric model. Universal models automatically trade off between parametric and nonparametric behavior, depending on the amount of data. We derive computationally efficient inference methods using the cascaded logistic base distribution (Park et al. 2013, Pachitariu et al. 2013), which permits tractable scaling to large populations. The cascaded logistic model assumes the dependencies among neurons take a simple cascaded form,

equivalent to a series of logistic regressions, with each neuron depending only upon those that came before. We prove a condition for the equivalence between the cascaded logistic and the 2nd-order maxent (or "Ising") model, making cascade-logistic a reasonable choice for base distribution in a universal model. Specifically, any Ising model with pentadiagonal interactions is also a cascaded logistic model. We evaluate the performance of these models with neural data, and will compare to the "Reliable Interaction Model", a recently proposed framework for building flexible models of binary spike patterns within the maximum entropy formalism (Ganmor et al. 2011)

### III-8. High frequency deep brain stimulation elicits neural restoration via loop-based reinforcement

Sabato Santaniello<sup>1</sup>  
Erwin B. Montgomery<sup>2</sup>  
John T. Gale<sup>3</sup>  
Sridevi Sarma<sup>1</sup>

SSANTAN5@JHU.EDU  
EMONTGOM@UAB.EDU  
GALEJ@CCF.ORG  
SREE@JHU.EDU

<sup>1</sup>Johns Hopkins University

<sup>2</sup>University of Alabama Birmingham

<sup>3</sup>Cleveland Clinic

High frequency deep brain stimulation (HF-DBS) of the basal ganglia (BG) is a recognized treatment for movement disorders in Parkinson's disease. Despite extensive assessment of its clinical effects, the mechanisms of HF-DBS remain elusive. The therapeutic merit of HF-DBS is hypothesized to stem from increasing the regularity of firing patterns in the BG, but, even though this is consistent with experiments in humans and animal models of Parkinsonism, it is unclear how the pattern regularization would originate from HF-DBS of the BG. To address this question, we constructed a network model of the cortico-BG-thalamo-cortical loop in normal and Parkinsonian conditions (880 neurons from cortex, striatum, GPI, and thalamus) and we used it to investigate the effects that subthalamic (STN) DBS has proximally to the stimulation target and distally through antidromic and orthodromic mechanisms for several DBS frequencies (20-180Hz). For each structure, our model matched data from non-human primates in normal and Parkinsonian conditions, with and without DBS, collected in several laboratories in the past twenty years. We found that (i) neither the proximal nor the distal effects individually account for pattern changes observed experimentally in the loop; (ii) the impact of their combined effects on the firing patterns increases with the stimulation frequency and becomes significant for HF-DBS; (iii) perturbations evoked proximally and distally propagate along the loop, rendezvous in the striatum, and, for HF-DBS, positively overlap (reinforcement), thus causing larger post-stimulus activation and more regular patterns in striatum. Reinforcement is maximal for clinically relevant 130Hz DBS and restores a more normal activity in the nuclei downstream (i.e., lower beta-band spectra in GPI, higher thalamic relay fidelity). Results suggest that reinforcement in striatum may be pivotal to initiate pattern regularization and to restore the function of the BG nuclei downstream, and might stem from frequency-selective resonant properties of the loop.

### III-9. Theory for the firing statistics of neurons driven by non-Poissonian synaptic input

Tilo Schwalger<sup>1</sup>  
Benjamin Lindner<sup>2</sup>

SCHWALGER@GMX.DE  
BENJAMIN.LINDNER@PHYSIK.HU-BERLIN.DE

<sup>1</sup>Ecole Polytechnique Federale de Lausanne

<sup>2</sup>Humboldt-Universität zu Berlin

A fundamental problem in theoretical neuroscience is to characterize the spiking statistics of a neuron for a given synaptic input. In a stationary state of activity, the relation between the statistics of synaptic input and spike output is essential for a theoretical understanding of the spontaneous activity of single neurons, neural populations as

well as their capability to process information. For mathematical tractability, presynaptic input is often modeled by Poisson processes, which possess a white noise spectrum. However, realistic input is generally not white but presynaptic spikes are temporally correlated, e.g. due to refractoriness, bursting, adaptation or due to a signal encoded in the spikes. Additionally, short-term synaptic plasticity shapes the temporal structure of synaptic currents. Yet, results on the spiking statistics in response to general, correlated input are still lacking. Here, we put forward analytical expressions for the interspike-interval (ISI) statistics of a perfect integrate-and-fire neuron model driven by a Gaussian process with an arbitrary correlation function. We use a weak-noise expansion of a multi-dimensional Fokker-Planck equation and solve the associated first-passage-time problem, which allows us to calculate the coefficient of variation (CV), the ISI correlation coefficient as well as the ISI distribution and its skewness in terms of the input correlation function. Our novel theory accurately predicts the ISI statistics induced by massive spike input as well as short-term plasticity. We find that presynaptic input that is more regular than Poisson ( $CV < 1$ ) causes negative ISI correlations and less skewed ISI distributions, resembling spike patterns commonly attributed to adapting neurons. In contrast, presynaptic bursting ( $CV > 1$ ) exhibits positive correlations and more skewed ISI distributions. Finally, we extend our formulas to nonlinear neural oscillators using the phase-response curve. Our results bear important implications for a self-consistent analysis of network activity and the interpretation of experimental ISI statistics.

### III-10. Error correction in the retinal population code

Jason Prentice<sup>1</sup>  
Olivier Marre<sup>2</sup>  
Mark Ioffe<sup>1</sup>  
Gasper Tkacik<sup>3</sup>  
Michael Berry<sup>1</sup>

JASONSP@PRINCETON.EDU  
OLIVIER.MARRE@GMAIL.COM  
MIOFFE@PRINCETON.EDU  
GASPER.TKACIK@IST.AC.AT  
BERRY@PRINCETON.EDU

<sup>1</sup>Princeton University

<sup>2</sup>Institut de la Vision - INSERM

<sup>3</sup>Institute of Science and Technology, Austria

Retinal ganglion cell spiking is unreliable over stimulus repeats. However, redundancy induced by correlations between cells may confer error-correcting properties on the population code. Analyzing multi-electrode array recordings of 100+ salamander retinal ganglion cells, we identify collective activity modes that encode the stimulus less noisily than any individual ganglion cell, supporting the error-correction hypothesis. We employ a novel variant of a hidden Markov model. The model classifies the spike timing response into a sequence of discrete states (the collective modes; 20 ms time bin), with cells having mode-dependent firing rates and correlations. Model inference is tractable while providing an excellent statistical description of the data. The model is trained using ganglion cell spike times, and no other knowledge of the stimulus such as repeat structure. It therefore captures information accessible to downstream brain regions. Modes were highly reproducible across stimulus repeats, in spite of fluctuations in the detailed population spiking pattern. The information efficiency (mutual information / output entropy) was higher for collective modes than for individual cells across several stimulus classes. The improvement was particularly pronounced under natural movies: the majority of modes were 80% - 97% efficient, while cells achieved no higher than 73%. We also assessed the representation of stimuli by modes. Under a white noise checkerboard, collective modes typically had compact receptive fields consistent with the intersection of several individual cells' receptive fields. Under a randomly moving bar stimulus, mode identity allowed excellent reconstruction of the bar's complete trajectory from a linear decoder (93% correlation between real and estimated bar position). Our model is computationally efficient and scalable to datasets of hundreds of neurons. Our work suggests that the retinal population output is formatted to support downstream extraction of relevant stimulus features against a noisy background.



### III-11. The nature of cortical variability

I-Chun Lin  
Michael Okun  
Matteo Carandini  
Kenneth Harris  
University College London

I.LIN@UCL.AC.UK  
M.OKUN@UCL.AC.UK  
MATTEO@CORTEXLAB.NET  
KENNETH.HARRIS@UCL.AC.UK

The responses of sensory cortex to repeated presentations of a stimulus are highly variable. This variability is correlated between neurons, and places constraints on information transmission. Understanding its nature is therefore critical to understanding the operation of sensory cortex. Neuronal variability is usually studied at the level of neuronal pairs by measuring noise correlations. A key property of noise correlations is that they depend on signal correlations: they are larger between neurons of similar sensory tuning. Here we show that cortical variability can be understood by studying an entire population of neurons simultaneously. We introduce a simple model of population activity that explains the structure of trial-to-trial variability in neuronal populations of visual cortex. The model comprises two sources of variability: a multiplicative component that invests all neurons in proportion to their instantaneous sensory drive (i.e., a variable response gain), and an additive component that is independent of the sensory drive but affects individual neurons to different degrees. This model captures the variability of individual neurons in a single trial and accurately predicts pairwise correlations. For example, it captures the precise dependence of noise correlations on tuning similarity. In addition, the model gives accurate predictions of the spike counts of individual neurons on single trials, using parameters fit through cross-validation. Our results show that cortical noise arises from a variable multiplicative gain change, whose effects dominate those of the commonly described additive noise. We speculate that these effects may arise from an interaction of sensory input with top-down control from higher order cortices.

### III-12. From sounds to meaning: neural representation of calls in the avian auditory cortex

Julie Elie  
Frederic Theunissen  
University of California, Berkeley

JULIE.ELIE@BERKELEY.EDU  
THEUNISSEN@BERKELEY.EDU

Understanding how the brain extracts meaning from vocalizations is a central question in auditory research. Communication sounds distinguish themselves not only by their acoustical properties but also by their information content. Here, we are developing the birdsong model to investigate how the auditory system extracts invariant features carrying symbolic information of the acoustic signals and categorize communication sounds according to their social meanings. Songbirds have been used extensively to study vocal learning but the communicative function of vocalizations and their neural representation has yet to be examined. In our research, we first generated a library containing the entire zebra finch vocal repertoire and organized communication calls along 9 different categories. We then investigated the neural representations of these semantic categories in the primary and secondary auditory areas of 6 zebra finches. To decrypt the neural computations underlying the classification of these calls into semantic categories, we used a combination of optimal decoding methods and encoding models of the neural response that took into account both the acoustical properties of the sounds and their semantic grouping. Both decoding and encoding analyses showed that neural responses in higher auditory areas can be more effectively explained by models that describe sounds in terms of their semantic content rather than just their acoustical features. The optimal decoding method revealed that for 2/3 of the units, neural discriminability of semantic groups is higher than what could be expected from the spectro-temporal features of the stimuli only. The encoding model revealed that many neural responses are best explained by non-linear transformations of spectro-temporal sound patterns and that these non-linearities emphasize the semantic grouping of calls. Combining these results with the anatomical properties of cells (positions and spike shapes) gives new insight into the neural representation of meaningful stimuli in the avian auditory neural network.

### III-13. Stimulus-specific adaptation in the auditory system — models and data

Israel Nelken  
 Leila Khouri  
 Bshara Awwad  
 Itai Hershenhoren  
 Amit Yaron  
 Tohar Yarden  
 Hebrew University

ISRAEL@CC.HUJI.AC.IL  
 LEILA.KHOURI@MAIL.HUJI.AC.IL  
 BSHARAAWWAD@GMAIL.COM  
 ITAI.HERSH@MAIL.HUJI.AC.IL  
 AMITYAR@GMAIL.COM  
 TOHARYAR@GMAIL.COM

Stimulus-specific adaptation (SSA) is the reduction in the response to a common stimulus that does not generalize, or only partially generalizes, to other, rare stimuli. SSA has been proposed to be a correlate of 'deviance detection', an important computational task of sensory systems. SSA is ubiquitous in the auditory system: it is found both in cortex and in subcortical stations, and it has been demonstrated in many mammalian species as well as in birds. A number of models have been suggested in the literature to account for SSA in the auditory domain. The simplest is the adaptation of narrowly-tuned modules (ANTM) model, in which a number of narrowly-tuned inputs, each of which showing simple fatigue, converge on a post-synaptic neuron. Here, we develop analytically some predictions of the ANTM model for the responses to various tone sequences as well as for the responses to more complex, wideband but spectrally balanced stimuli called 'tone clouds'. We then show that these predictions generally hold in the inferior colliculus (IC), a major midbrain auditory center and currently the first station of the auditory pathway in which SSA has been found; but that all of them fail in auditory cortex. We conclude that SSA in IC can be accounted to a large degree by simple convergence of narrowly-tuned inputs, accounting for the association between SSA and the non-lemniscal subdivisions of the IC. On the other hand, an account of cortical SSA requires more complex mechanisms, resulting in true sensitivity to sensory deviance.

### III-14. Modeling sound pulse counting with phasic inhibition and nonlinear integration

Richard Naud<sup>1</sup>  
 Dave Houtman<sup>1</sup>  
 Gary Rose<sup>2</sup>  
 Andre Longtin<sup>1</sup>

RNAUD@UOTTAWA.CA  
 DAVEHOUTMAN@HOTMAIL.COM  
 GARY.ROSE@M.CC.UTAH.EDU  
 ALONGTIN@UOTTAWA.CA

<sup>1</sup>University of Ottawa  
<sup>2</sup>University of Utah

The nervous system must perform many types of computations. A simple computation that frogs execute routinely is counting the number of consecutive sound pulses in a conspecific call that occur with precise and regular timing. Cells signaling that a threshold number of pulses have occurred have been found in the midbrain of anurans. These counting cells will not respond if a single inter-pulse interval is a few milliseconds longer than the baseline interval. What intrinsic or network mechanisms can give rise to such pulse/interval counting? Comparing simplified neuron models with previously published in vivo membrane potential recordings, we identify biophysical processes that can explain the observations. First we consider short-term facilitation of excitation and non-linear subthreshold dynamics such as provided by persistent sodium currents. We find that these processes reproduce some features of uninterrupted call data, whether they are considered one at a time or together. Second, we consider a model of phasic inhibition made of onset and offset inhibition. Phasic inhibition enhances reset because a longer interval will engender onset and possibly offset inhibition. Combining phasic inhibition with either synaptic facilitation, nonlinear subthreshold dynamics or both can qualitatively reproduce the array of recordings for different pulse patterns – including those with pauses that reset the counting – as well as the effect of pharmacologically attenuating inhibition. These results support the hypothesis that prior segmentation of sound via phasic on and off responses underlies the emergence of features such as pulse counting and duration selectivity in the auditory midbrain.

## III-15. How do insects see motion? Using connectomics, physiology and behavior to constrain possible models

Arjun Bharioke<sup>1,2</sup>  
Dmitri B Chklovskii<sup>1,2</sup>

BHARIOKEA@JANELIA.HHMI.ORG  
MITYA@JANELIA.HHMI.ORG

<sup>1</sup>Janelia Farm Research Campus  
<sup>2</sup>Howard Hughes Medical Institute

Motion detection is an essential visual computation and has been studied extensively. Many models have been proposed, yet it is still not clear which one is implemented in the insect brain. Analysis of the fly visual system's connectome (Takemura et al. (2013)) revealed that motion is detected by combining inputs from multiple photoreceptors onto a single cell, via two neuron types: Mi1 and Tm3. These two channels provide spatially displaced inputs to the motion detecting cell (Takemura et al. (2013)), in agreement with existing correlation-type models (Hassenstein, Reichardt (1956); Barlow, Levick (1961)) (HR; BL). However, the sign of the synaptic input from each channel remains unknown, complicating the choice between the two models. From the connectome, we observe that the receptive field components mediated by the two channels overlap spatially, with the Mi1 component almost completely contained within the Tm3 component (Takemura et al. (2013)). In light of this and other recent electrophysiological and behavioral evidence, we analyze the viability of HR and BL models in describing fly vision. We find that, given overlapping receptive field components, HR detectors do not show stronger responses to stimuli moving in the preferred vs. the opposed, null direction (PD; ND). In contrast, BL detectors do. This suggests that the HR model may not be viable in the fly, a result additionally supported by recent measurements of the relative delay between Mi1 and Tm3 (Behnia, Clark, Clandinin, Desplan, unpublished). We then characterize the possible presence of rectification within the remaining BL model. Joesch et al (2010) and Clark et al (2011) showed that the motion sensitive neurons downstream of Mi1 and Tm3, respond preferentially to ON vs. OFF stimuli. We find that only a subset of BL models with rectification show this preference. Finally, within this subset, we provide testable predictions of behavioral and physiological responses in specific cells.

## III-16. Predicting IT and V4 neural responses with performance-optimized neural networks

Daniel Yamins  
Ha Hong  
Darren Seibert  
James DiCarlo

YAMINS@MIT.EDU  
HAHONG@MIT.EDU  
DARREN@MIT.EDU  
DICARLO@MIT.EDU

Massachusetts Institute of Technology

Though substantial evidence shows that higher ventral cortex supports the capacity for rapid view-invariant object recognition, processing mechanisms in these areas remain poorly understood. Here, we hypothesized that an effective way to produce models of individual IT and V4 neuronal responses would be to match the higher ventral stream's population-level "behavior". If IT developed to support invariant object recognition behavior, networks that perform invariant object recognition better tasks should also more effectively capture responses of high-level visual neurons. To test this we used high-throughput computational techniques to evaluate thousands of candidate architectures within a class of biologically-plausible feedforward neural network models. We measured categorization performance and IT neural explained-variance for each model, finding a strong correlation between a model's performance and IT predictivity. This result has a clear implication: to find highly IT-like models, performance optimization may be an effective strategy. We next identified a highly optimized feedforward network model that matches IT performance on a range of challenging recognition tasks. Critically, even though we did not explicitly constrain this model to match neural data, its output layer turned out to be highly predictive of IT neural responses — a 90% improvement over comparison models and comparable to the state-of-the-art models of lower-level areas such as V1. Moreover, the penultimate model layer is highly predictive of V4 neural responses, showing that performance optimization imposes biologically consistent constraints on intermediate feature representations

as well. A common assumption in visual neuroscience is that understanding the qualitative structure of tuning curves in lower cortical areas is a necessary precursor to explaining higher visual cortex. Our results indicate that it is useful to complement this “bottom-up” approach with a “top-down” perspective in which behavioral metrics are a sharp and computationally tractable constraint shaping individual neural tuning curves in both higher and intermediate cortical areas.

### **III-17. Adaptation in V1 populations reflects optimal inference of natural movie statistics**

Michoel Snow<sup>1</sup>  
Ruben Coen-Cagli<sup>2</sup>  
Odelia Schwartz<sup>1</sup>

MICHOEL.SNOW@MED.EINSTEIN.YU.EDU  
RUBEN.COENCAGLI@UNIGE.CH  
ODELIA.SCHWARTZ@EINSTEIN.YU.EDU

<sup>1</sup>Albert Einstein College of Medicine

<sup>2</sup>University of Geneva

Perception and neural processing of a stimulus are dramatically influenced by what has been observed in the past, its temporal context. The properties and goals of adaptation have been studied extensively, but it remains unclear whether they can be explained as an optimization to the temporal structure of the natural environment. We address this question by hypothesizing that primary visual cortex learns statistical regularities of natural movies and uses adaptation to reduce their redundancy. We consider a Mixture of Gaussian Scale Mixtures (MGSM; Coen-Cagli et al., 2012) generative model, which we previously applied to spatial context. We develop a model that captures statistical dependencies, for a neuronal population, between the responses of oriented filters to multiple frames of natural movies. Given an ensemble of movies, we optimize model parameters, including the prior probability that past and present frames are dependent, as well as the covariance matrix quantifying such dependency. As new stimuli are presented, the prior is continually updated to reflect the changing environment. We assume cortical responses to novel stimuli represent Bayesian inferences in the model. In the MGSM, this amounts to a flexible divisive normalization; the filter response to the present (test) input is divisively normalized by responses to past (adapter) and present inputs, but only to the extent that they are inferred dependent. The model reproduces classical effects of tuning curve repulsion (shift in preferred orientation away from the adapter), because the largest probability of dependency in the population, and therefore maximal suppression, arises when adapter and test are matched, irrespective of the neuron’s preferred orientation. We further consider how adaptation is updated in the model over time, using random grating sequences with a bias for a particular orientation (Benucci et al. 2013). The updated prior tracks the bias level, thus equalizing the time averaged population responses.

### **III-18. Heisenberg: a tool for mapping complex cell receptive fields**

Urs Koester<sup>1</sup>  
Bruno Olshausen<sup>2</sup>

URS@BERKELEY.EDU  
BAOLSHAUSEN@BERKELEY.EDU

<sup>1</sup>Redwood Center

<sup>2</sup>University of California, Berkeley

The majority of V1 cells display some degree of phase-invariance in their response [Mechler & Ringach 2002]. Thus linear filter models fail to predict most V1 responses well. On the other hand, Fourier methods such as FFT power / phase-separated Fourier transform [David & Gallant, 2004, 2005] discard location information. This necessitates a priori knowledge of the receptive field envelope. Subunit models [Vintch et al., 2012; McFarland et al., 2013] alleviate this problem and capture local, phase-invariant responses, but require solving a non-convex, high-dimensional optimization problem requiring clever optimization and regularization schemes. We propose a spatial spectrogram method dubbed Heisenberg, that combines the robustness and simplicity of linear methods with much of the flexibility of subunit methods. Following the method of [Nishimoto et al., 2006]

we create an overcomplete tiling of phase-invariant filters by computing the local Fourier power in overlapping patches of the stimulus window. Looking at spike-triggered averages in this representation, we note that receptive fields are approximately separable in the pixel and frequency domain. Thus they are well captured by a rank one approximation that takes the form of an outer product between a spatial envelope and local Fourier power spectrum. The model is linear (and hence the estimation convex) in both the spatial and the frequency filter, making it easy to estimate by alternating updates in the two filters. The rank one approximation acts as a form of regularization that reduces the dimensionality of the problem sufficiently to make it tractable even for small datasets. Applied to multi-electrode recordings from a population of V1 cells, the method provides state of the art performance predicting responses to a natural movie stimulus.

### III-19. Are spikes unitary signals? How axonal noise produces synaptic variability

Ali Neishabouri

Aldo A. Faisal

Imperial College London

M.NEISHABOURI10@IMPERIAL.AC.UK

A.FAISAL@IMPERIAL.AC.UK

The waveform of the action potential is generally assumed to be conserved as it travels from the soma to the synapse. However, this view is based on data from classically studied axons, which are very large (0.1-1mm) compared to average diameter of axons in mammalian nervous systems (0.2-0.3 $\mu$ m). This has prompted us to reconsider the role of axons as faithful transmission channels. Post-synaptic potential (PSP) variability is typically attributed to mechanisms inside the synapse, yet in the many thin axons of our brain, the action potential (AP) waveform and thus the Ca<sup>++</sup> signal controlling vesicle release at the synapse will be significantly affected by the inherent stochasticity of ion channels. We investigate to what extent fluctuations in the AP waveform cause the observed PSP variability. We show, using both biophysical theory and stochastic simulations of CNS and PNS axons from vertebrates and invertebrates, that channel noise in axons below 1 $\mu$ m causes random changes in the AP waveform. AP height and width, both experimentally characterised parameters of post-synaptic response amplitude, vary e.g. by 6 mV and 1.5 ms as a single AP propagates in 0.3  $\mu$ m diameter axons (cortical axon collaterals). We show how AP height and width variabilities increase with a — power-law as diameter decreases and translate these fluctuations into post-synaptic response variability using biophysical data and models of synaptic transmission. We find that e.g. in synapses innervated by 0.2 $\mu$ m diameter unmyelinated axons (cerebellar parallel fibres) axonal noise alone can account for half of PSP variability. Axonal variability may have considerable impact on synaptic response variability. However, it can be easily confounded with synaptic variability in many theoretical STDP models or experimental frameworks investigating synaptic transmission through paired-cell recordings or extracellular stimulation of the presynaptic neuron and intracellular postsynaptic recordings.

### III-20. Optimal population codes with limited input information have finite tuning-curve widths

Guillaume Dehaene<sup>1,2</sup>

Jeffrey Beck<sup>3</sup>

Alexandre Pouget<sup>1</sup>

GUILLAUME.DEHAENE@GMAIL.COM

JEFF.BECK@DUKE.EDU

ALEX.POUGET@GMAIL.COM

<sup>1</sup>University of Geneva

<sup>2</sup>Universite Paris-Descartes

<sup>3</sup>University of Rochester

Sensory variables are encoded in the brain via the coordinated activity of large populations of variable neurons. Several factors influence the quality of such codes: noise correlations, neuronal density, peak firing rate and tuning curve width. Previous studies have used efficient coding arguments to derive the optimal relationship between

these variables in the presence of resource constraints. Most of these studies, however, assume that the population of neurons receives an infinite amount of information and that the only source of behavioral variability comes from Poisson noise. Here we revisit this issue but with two new constraints: (1) we seek efficient probabilistic population codes, that is, codes that represent probability distributions as opposed to scalar estimates of the variables, and (2) we consider the regime in which the information entering the network is limited in the sense that it saturates for large number of neurons. The second constraint is relevant for the brain as sensory inputs only provide limited information about latent variables (e.g., speech is not infinitely informative about the identity of the words being uttered). We show that under these constraints, neurons have an optimal tuning curve width. By contrast, previous studies had concluded that optimal tuning curves should be infinitely narrow, which is obviously not the case in vivo. Importantly, the optimal width is predicted solely by amount of input information and does not depend on either the prior over possible stimulus values or resource constraints such as limits on the total number of neurons or total number of spikes. Finally, our analysis predicts that when input information follows a Weber scaling, the tuning curve width should be proportional to the preferred value of the neurons, which is indeed what has been found experimentally for variables like numerosity or speed.

### III-21. Optimal compensation for neuron death

David G.T. Barrett<sup>1,2</sup>  
Sophie Deneve<sup>1</sup>  
Christian Machens<sup>3</sup>

DAITHIOBAIREAD@GMAIL.COM

SOPHIE.DENEVE@ENS.FR

CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

<sup>1</sup>Ecole Normale Supérieure

<sup>2</sup>Champalimaud Centre for the Unknown

<sup>3</sup>Champalimaud Neuroscience Programme

The brain has an impressive ability to withstand neural damage. Diseases that kill neurons can go unnoticed for years, and acute conditions such as silent stroke have little impact on neural function. How does the brain compensate for such damage, and what are the limits of this compensation? We propose that neural circuits optimally compensate for neuron death, thereby preserving their function as much as possible. This proposal is a simple generalisation of the efficient coding principle, from the intact brain state to the damaged brain state. We show that this compensation can explain tuning curve shape changes induced by neuron silencing across a variety of systems, including the primary visual cortex. Specifically, we calculate the firing rate tuning curves in a neural system that is optimised to provide an efficient representation of a signal (Barrett et al. 2013). Then, we artificially silence some neurons by constraining their firing rates to be zero and again, we calculate efficient coding tuning curves. The difference between the firing rates in the damaged and the intact models is optimal compensation, because it preserves efficient coding as much as possible. Furthermore, we find that optimal compensation can be implemented through the dynamics of networks with a tight balance of excitation and inhibition, without requiring synaptic plasticity. In previous work, we found that the balance of excitation and inhibition is synonymous with efficient coding (Boerlin et al.). Here, we find that balanced network dynamics implements optimal compensation, by decreasing excitation (or inhibition) following the death of inhibitory (or excitatory) neurons. The limits of this compensatory mechanism are reached when excitation and inhibition can no longer balance each other. This limit demarcates a recovery boundary, where signal representation fails, and where diseases may become symptomatic.

### III-22. Optimal sensory coding by populations of ON and OFF neurons

Julijana Gjorgjieva<sup>1</sup>  
Markus Meister<sup>2</sup>  
Haim Sompolinsky<sup>3</sup>

GJORGJIEVA@FAS.HARVARD.EDU

MEISTER@CALTECH.EDU

HAIM@FIZ.HUJI.AC.IL

<sup>1</sup>Harvard University

<sup>2</sup>California Institute of Technology

<sup>3</sup>Hebrew University

Sensory responses of a population of neurons generally split into two types: ON cells that respond to stimulus onset and OFF cells that respond to stimulus offset. The ON-OFF dichotomy exists in different sensory modalities and species including the vertebrate retina [Kuffler, 1953], invertebrate motion vision [Joesch et al, 2010], thermosensation [Gallio et al, 2011], and chemosensation [Chalasanani et al, 2007], suggesting a strong evolutionary pressure driving the emergence of separate channels for signaling stimulus increments and decrements. We test the hypothesis that signal divergence into ON and OFF channels serves to optimize the encoding of the stimulus. We study a population of  $N$  cells that respond to a common stimulus with a binary firing rate: ON (OFF) neurons fire Poisson spikes with a fixed mean count whenever the stimulus intensity is above (below) their threshold, and zero otherwise. We compare the coding efficiency of different mixtures of ON and OFF cells using two measures: the mutual information between population spike counts and stimulus, and the mean square error of the optimal linear stimulus estimator. Counter to our expectations, we find that the mutual information is identical for any combination of ON and OFF cells once the thresholds of all cells are optimized. However, the total mean spike count is lowest for the population with equal numbers of ON and OFF cells, making this arrangement optimal in terms of bits per spike. For the mean square error criterion using a linear decoder, the optimal mixture depends on both the noise level and any asymmetries in the stimulus distribution. Our theory makes predictions about the optimal distribution of thresholds in sensory systems that use a population of neurons to encode the same stimulus variable. We are now testing these predictions on experimentally recorded retinal ganglion cells in the vertebrate retina.

### III-23. Distinguishing different efficient coding principles from tuning curve statistics

Zhuo Wang  
Alan Stocker  
Daniel Lee

University of Pennsylvania

WANGZHUO@SAS.UPENN.EDU  
ASTOCKER@SAS.UPENN.EDU  
DDLEE@SEAS.UPENN.EDU

Efficient coding theories argue that the tuning curves of neural populations are adapted to optimally encode the stimulus distributions. Previous theories have relied on knowing the input stimulus distributions in order to predict observed tuning curve characteristics. These theories compute the optimal tuning characteristics based upon optimality principles such as maximizing mutual information (Infomax) or minimizing squared error (L2-min), typically assuming a one-dimensional stimulus variable. However, is it possible to determine whether the neural population response represents an efficient code even when the input distribution is unknown and if so, which optimality principle the code obeys? On one hand we show that for one-dimensional stimulus variables the tuning widths are related to the neural density in a manner that does not permit a distinction between codes that maximize mutual information and minimize L2 loss. On the other hand, we demonstrate that efficient codes can be differentiated for two- and higher-dimensional stimulus variables according to a statistical signature that relates neural tuning widths to neural density. In particular, Infomax solutions can be distinguished from L2 solutions by projecting their tuning widths along a special set of axes computed from the neural density statistics. We believe that our analysis provides a useful tool to further test the efficient coding hypothesis especially in situations where tuning characteristics of neurons can be measured but the stimulus distributions are unknown.

## III-24. Second order phase transition describes maximally informative encoding in the retina

David Kastner<sup>1,2</sup>  
 Stephen A Baccus<sup>2</sup>  
 Tatyana Sharpee<sup>3</sup>

DAVID.KASTNER@EPFL.CH  
 BACCUS@STANFORD.EDU  
 SHARPEE@SALK.EDU

<sup>1</sup>Ecole Polytechnique Federale de Lausanne

<sup>2</sup>Stanford University

<sup>3</sup>Salk Institute for Biological Studies

Neural circuits use diverse populations of neurons to perform complex computations. However, we lack a detailed understanding of how this diversity benefits the nervous system. Theoretical arguments, based upon the maximization of information transmitted about incoming stimuli, have proven successful in accounting for properties of single neurons, and they predict the usefulness of encoding with uniform populations of neurons. Here we apply information maximization to understand the encoding of distinct populations of neurons in the retina. We show that the maximally informative solution for populations of neurons exhibits a second-order phase transition with respect to an effective temperature, defined as the inverse of the average slope of the neural response functions. The broken exchange symmetry between neurons associated with this transition describes two, recently discovered, subtypes of fast Off retinal ganglion cells. It also explains the homogeneity of On cells, which have shallower slopes, mapping them onto an effective temperature beyond the transition where it is no longer optimal to split neurons into different populations. We show that the maximally informative solution exhibits the same set of singularities and scaling laws near the critical point as expected for the Ising model of magnetism, a classic model of phase transitions in physical systems. Our measurements in the retina matched experimental values observed in this class of physical systems. The difference in spiking thresholds played the role analogous to magnetization, refining previous suggestions that the balance between excitation and inhibition might be analogous to magnetization. While these results were motivated by properties of retinal cells, they should describe the emergence of different neuronal classes elsewhere in the nervous system. Connecting specific neuronal properties to well-established theories of critical systems represents a strong step towards understanding the universality of various, and sometimes controversial, observations of criticality in the nervous system.

## III-25. Sensor placement for sparse sensory decision-making

Bingni Brunton  
 Annika Eberle  
 Bradley Dickerson  
 Steven Brunton  
 J. Nathan Kutz  
 Tom Daniel

BBRUNTON@UW.EDU  
 EBERLE10@UW.EDU  
 BDICKER@U.WASHINGTON.EDU  
 SBRUNTON@UW.EDU  
 KUTZ@UW.EDU  
 DANIELT@UW.EDU

University of Washington

Organisms are remarkably adept at interacting with high-dimensional physical systems in nature, yet they sometimes rely on information gathered through a handful of sensory organs. For example, in addition to visual inputs, a flying insect adjusts to disturbances in the air using tens to hundreds of strain sensors on its wings, despite the millions or billions of states required to characterize the surrounding fluid. This strikingly efficient sensory-motor performance is possible in part because natural signals are inherently compressible, having relatively low-dimensional features underlying high-dimensional dynamics. In addition, it is rarely necessary to reconstruct details; instead, the animal has only to decide on the state of the environment corresponding to the required motor response. Consequently, a few well-placed sensors may capture variations in these features. Here we ask how a fixed budget of sensors should be placed to optimally inform decision-making. Specifically, we explored how an enhanced sparsity for classification may be exploited by campaniform sensilla (strain sensors) on the surface of a flapping insect wing. We took an array of strain measurements from models of a hawkmoth



wing in two distinct behavioral states: stationary flapping, versus flapping under body rotation. This large set of measurements was refined to select a much smaller subset of key locations that best serve detection of forces present only under body rotation. This analysis was repeated by imposing neuronal properties corresponding to experimentally observed nonlinearities of campaniform sensilla encoding. We found that, with neuronal encoding, efficient detection of body rotation could be achieved with only 2–5 strain sensors, located primarily at the wing's center. Importantly, this analysis provides insight into biological sensor placement and design, constrained by the knowledge that these sensors follow neuronal encoding, but requires no a priori knowledge of physical features, biomechanics, or access to the high-dimensional dynamical system.

### III-26. The role of environmental feedback in transitions from passive to active brain processing

Christopher Buckley<sup>1,2</sup>  
Taro Toyozumi<sup>1</sup>

CHRISTOPHER.LAURIE.BUCKLEY@GMAIL.COM  
TARO.TOYOIZUMI@BRAIN.RIKEN.JP

<sup>1</sup>RIKEN Brain Science Institute  
<sup>2</sup>CCNR, Sussex University

Coherent behavior emerges from the mutual interaction between the brain, body and environment and not from within the brain alone. Sensory perception and motor action are inseparably bound by feedback mediated by the body and environment. However, it is unknown how such feedback impacts on neural processing. Whisking behavior has been the model system for understanding brain/body/environment interactions. The membrane potentials of neurons in the barrel cortex of quiet attentive mice exhibit low frequency fluctuations and nearby neurons are highly correlated both of which are suppressed with the onset of whisking. Whisking also accompanies a desensitization of neurons to passive whisker deflections and sensitization to active whisker contacts. Here we show that all these experimental results can be straightforwardly explained by appealing to the role of re-afferent neural feedback mediated by the body. Specifically we show, in a simple model, that assuming brain-whisker interactions constitute a negative feedback pathway that stabilises cortical dynamics is sufficient to account for the changes in both brain activity and sensory processing. To experimentally demonstrate that body/environment feedback is sufficient to modulate brain dynamics we analyzed calcium-imaging data from zebrafish larvae in a virtual reality environment. Neural activity was sampled from whole-brain preparations under a visual feedback and passive visual-replay condition, i.e., in the presence or absence of environment feedback, respectively. We found that: 1. neuronal fluctuations and correlations are suppressed with the onset of visually guided behavior but not during the visual replay condition; 2. The brain-environmental interaction provides a negative feedback signal to the brain; 3. The degree of the suppression of each neuron is predicted by the degree to which a neuron is involved in this feedback. In sum, these results strongly suggest that the mapping between sensory stimuli and responses is highly flexible and dependent on the context of brain/body/environment feedback.

### III-27. Calcium-dependence of vesicle replenishment at photoreceptor synapses

Caitlyn Parmelee  
Carina Curto  
University of Nebraska, Lincoln

S-CPARMEL1@MATH.UNL.EDU  
CCURTO2@UNL.EDU

To fully understand visual computation by the retina requires understanding the way signals are transformed at the very first synapse: the ribbon synapse of photoreceptor neurons (rods and cones). These synapses possess a ribbon-like structure on which approximately 100 synaptic vesicles can be stored, allowing graded responses through the release of different numbers of vesicles in response to visual input. The dynamics of these responses, however, depend critically on the ability of the ribbon to replenish itself when ribbon sites become empty after vesicle release. The rate of vesicle replenishment is thus an important factor in shaping neural coding in the

retina. Experimentally, the replenishment rate has been shown to be regulated by calcium (Ca/CaM), although the mechanisms for this process remain unclear. To determine the possible effects of Ca/CaM on the rate of vesicle replenishment, we modeled vesicle diffusion in the synaptic terminal of ribbon synapses by spherical vesicles undergoing random walks on a 3-dimensional rectangular lattice. We found that the time constant(s) governing the replenishment rate depends only on the diffusion coefficient, the density of (mobile) vesicles in the synaptic terminal, the vesicle diameter, and a 'stickiness' factor,  $s$ , corresponding to the probability of successful attachment in the event of a vesicle collision with the ribbon. We then analyzed two variations of our model to explore possible mechanisms by which Ca/CaM can regulate vesicle replenishment. The model analysis revealed a qualitative difference between the predicted kinetics of vesicle replenishment when Ca/CaM acted on vesicles as opposed to ribbon tethering sites. By comparing to data collected in the Thoreson lab at UNMC (Omaha, NE), we were able to rule out the action directly on vesicles and support the hypothesis that calcium promotes attachment of colliding vesicles onto the ribbon.

### III-28. Uncertainty modulates timing of neural interaction in PMd during reach planning

Pavan Ramkumar<sup>1,2</sup>

Brian Dekleva<sup>2</sup>

Paul Wanda<sup>2</sup>

Hugo Fernandes<sup>1</sup>

Lee E Miller<sup>2</sup>

Konrad Koerding<sup>1</sup>

PAVAN.RAMKUMAR@NORTHWESTERN.EDU

BRIANDELKEVA2017@NORTHWESTERN.EDU

PWANDA@NORTHWESTERN.EDU

HUGOGUH@GMAIL.COM

LM@NORTHWESTERN.EDU

KK@NORTHWESTERN.EDU

<sup>1</sup>Rehabilitation Institute of Chicago

<sup>2</sup>Northwestern University

Hand movements in the real world are often planned with uncertain information about where to move. Understanding the role of uncertainty in movement plans is important to improve rehabilitation therapies and brain-based prostheses. Bayesian estimation theory, which combines sources of information in proportion to their uncertainty, predicts movement behavior when uncertainty about subjective beliefs (priors) and sensory observations (likelihoods) is varied. This suggests that the brain performs some kind of probabilistic inference. Here, to study the neural mechanisms underlying probabilistic computation, we test the intriguing possibility that uncertainty modulates neural interaction during a sensorimotor decision-making task. We characterized functional connectivity in a circuit comprising 15 neurons from the dorsal premotor cortex when one monkey planned uncertain center-out reaches. Both prior (blocked) and likelihood (interleaved) uncertainties of target position were varied in the experiment systematically. Comparing Bayesian predictions with behavior suggests that monkeys combine prior and likelihood uncertainties meaningfully. Using generalized linear models (GLMs) of Poisson-spiking neurons we analyzed activity from the preparatory period and inferred post-spike-time filters (PSTF). The PSTF for a pair of neurons ( $i, j$ ) describes the temporal evolution of the gain in firing rate of neuron  $i$ , given that neuron  $j$  spiked at time zero. We found significant effects of prior and likelihood uncertainty on the peak times of the PSTF ( $P < 0.005$ ; Bonferroni corrected for 900 multiple comparisons). However, no significant effects of uncertainty were found for the peak amplitude or the area under the curve of the PSTF ( $P > 0.05$ ; corrected). Our results suggest that reach planning under uncertainty is mediated by the timing of signaling between pairs of critical neurons implicated in planning complex movements. This work bridges normative models of decision making with behavior and brain activity, thereby constraining neural models of Bayesian inference.

### III-29. Bayesian inference of neural connectivity via approximate message passing

Alyson Fletcher<sup>1</sup>  
Sundeep Rangan<sup>2</sup>

AFLETCHER@SOE.UCSC.EDU  
SRANGAN@POLY.EDU

<sup>1</sup>University of California, Santa Cruz

<sup>2</sup>New York University

Key to understanding information processing in neural circuits is a systematic characterization of synaptic connectivity. New technologies such as calcium imaging and high-density electrode arrays provide parallel recordings of the spiking activity of large numbers of neurons. Computational methods to infer connectivity from these measurements using a Bayesian expectation-maximization (EM) approach are effective but prohibitively expensive for large problems. This work introduces a computationally efficient alternative based on the approximate message passing (AMP) methodology, which employs Gaussian approximations of loopy belief propagation. In the proposed approach, the spike activities are modeled as outputs of a nonlinear dynamical system involving the synaptic weights. The filtered spike history and synaptic weights are estimated jointly using an EM algorithm, where both the E and M steps are simplified using AMP. We show that our method is computationally scalable to large numbers of neurons with much lower cost than Markov chain Monte Carlo (MCMC) and extended Kalman filter methods. Additional unknowns, such as parameters of nonlinearities in neuron models, can be estimated alongside spike histories and synaptic weights, giving provably consistent estimates despite little increase in computational complexity. Furthermore, our technique can readily incorporate prior information on the unknown parameters, such as sparsity of neural connectivity and hence of synaptic weights.

### III-30. Reconstructing functional connectivity in complete neural population by randomized sparse sampling

Yuriy Mishchenko  
Toros University

YURIY.MISHCHENKO@GMAIL.COM

Reconstructing neural connectivity from functional imaging is important current challenge in neuroscience. One of the key problems existing in this approach today is practical impossibility to fully image the complete neural population comprising even one neural circuit. This leads to a well known problem of unobserved correlated neural inputs producing 'phantom' connections between unrelated neurons. It had been recently suggested that by functionally imaging small parts of a neural population in such a way that a full coverage of the population can be achieved cumulatively by all observations, it may be still possible to reconstruct the entire neural population's connectivity matrix (Keshri et al., arxiv 1309.3724). The estimation problem in these settings is solved by the standard Expectation Maximization (EM) algorithm, thus resulting in an attractive method for functional connectivity estimation in complete neural populations. An important aspect of this proposal is that, even though the entire neural population is covered by the observations cumulatively, at no point of time the complete input-output information is available. It is, therefore, interesting if the correct and unbiased connectivity matrix should be indeed recoverable in this manner. We study the proposed method in the simplest settings of a linear neural activity model to produce answers to this question. We show that the proposed estimation approach is indeed asymptotically consistent and unbiased, despite the lack of complete input-output information. We also establish the connections with the earlier autoregressive hidden input models and analyze the method's convergence speed and the effects of hidden population size and coupling strength. We furthermore identify important drop in the performance of the EM procedure in the full model for large hidden population sizes and also show that the EM estimator in the full model has a large number of local maxima, creating possibility of poor convergence.

**III-31. Classical conditioning via inference over observable situation contexts**

Nisheeth Srivastava  
Paul Schrater  
University of Minnesota

NISHEETHS@GMAIL.COM  
SCHRATER@UMN.EDU

Courville and colleagues have demonstrated an alternative explanation for classical conditioning as resulting for organisms inferring latent causes that predict the occurrence of reinforcing stimuli, which suggests radically alternative functionality for the neural correlates of choice-related behavior. In this work, we demonstrate that it is possible to explain classical conditioning-related data using an even simpler framework than latent cause models, making simpler assumptions about the computational capabilities of subjects. Whereas earlier latent cause models assume that animals try to predict the causal structure of their environment, we assume that they have a more modest goal - predicting co-occurrence patterns in environmental stimuli. The primary finding of our research is that predicting stimulus co-occurrence patterns in a Bayes-optimal manner explains classical conditioning without having to invoke latent causes. Simulated experiments with a standard Bayesian implementation of this model show that it is capable of explaining a broader range of effects than any previous theory of classical conditioning. By simplifying the mathematical structure of statistical modeling of conditioning and demonstrating its ability to explain a large set of experimentally observed effects, our work advances Bayes-optimal inference about stimulus co-occurrence as a rational principle explaining classical conditioning. The model also gives a novel interpretation for the neuroanatomical substrates in classical conditioning.

**III-32. Tractable, optimal high dimensional Bayesian inference**

Madhu Advani  
Surya Ganguli  
Stanford University

MSADVANI@STANFORD.EDU  
SGANGULI@STANFORD.EDU

Modern neuroscience has entered the era of high dimensional data. It is often the case that we can simultaneously record  $N = O(100)$  to  $O(1000)$  neurons but only for a limited number  $P$  trials per condition, which may be the same order of magnitude as the dimensionality of data ( $N$ ). This high dimensional data scenario carries with it the curse of dimensionality; in essence it is exceedingly difficult to analyze such limited amounts of high dimensional data by estimating large probabilistic models. We develop analytical methods and algorithms to combat this curse for a wide variety of (Bayesian) regression problems. For example, consider the problem of estimating functional connectivity between neurons. This is effectively a high dimensional regression problem in which an unknown pattern of  $N$  synaptic weights onto a noisy neuron are to be inferred, based on  $P$  measured inputs and outputs to the neuron, and where we may have an appropriate prior over synaptic weights. Maximum likelihood (ML) or maximum a posteriori (MAP) estimation are almost ubiquitous methods for solving this regression problem. We consider instead the optimal tractable estimator; this calculation involves optimizing an arbitrary loss function over the unknown synaptic weights. We used methods from statistical mechanics to find this optimal loss. Intriguingly, we find the optimal function is neither ML, nor MAP, but involves a smoothed version of the log-likelihood function and the prior where the degree of smoothing depends on the ratio  $P/N$ . The optimal loss function enjoys substantial improvements in squared error relative to ML and MAP. These results indicate that widely cherished Bayesian procedures for analyzing data must be modified in the high dimensional setting. In essence, they suggest we should strive not just to be Bayesians, but smooth Bayesians.

### III-33. Short-term facilitation as a normative consequence of presynaptic spike-rate adaptation

Simone Carlo Surace  
Jean-Pascal Pfister  
University of Bern

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

Synapses are highly dynamical elements. On the hundreds of milliseconds time scale, their strength can increase (facilitate) or decrease (depress). Despite the ubiquity of this form of plasticity, called short-term plasticity (STP), the functional advantages of such a history-dependent synaptic strength remain elusive. Recently, it has been shown that short-term plasticity acts as the optimal estimator of the presynaptic membrane potential based on the observed past spike timings. A limitation of this model is the absence of an explicit mechanism for refractoriness of the presynaptic neuron. Here we expand the model by including refractoriness and adaptation effects and we show that this modification critically impacts the predictions on short-term plasticity. More specifically we show that in the presence of strong refractoriness the model predicts facilitation whereas short-term depression is predicted for neurons with weak refractoriness or even bursting properties. Those predictions can be proven analytically and are illustrated with numerical simulations. Finally, this new prediction on the link between refractoriness of the presynaptic neuron and facilitation of the downstream synapse is directly testable.

### III-34. Learning, inference, and replay of hidden state sequences in recurrent spiking neural networks

Dane Corneil<sup>1,2</sup>  
Emre Neftci<sup>3</sup>  
Giacomo Indiveri<sup>4</sup>  
Michael Pfeiffer<sup>2</sup>

DANE.CORNEIL@EPFL.CH  
NEMRE@UCSD.EDU  
GIACOMO@INI.PHYS.ETHZ.CH  
PFEIFFER@INI.PHYS.ETHZ.CH

<sup>1</sup>Ecole Polytechnique Federale de Lausanne

<sup>2</sup>ETH Zurich

<sup>3</sup>University of California, San Diego

<sup>4</sup>ETH/UZH Zurich

Learning to recognize, predict, and generate spatio-temporal patterns and sequences of spikes is a key feature of nervous systems, and essential for solving basic tasks like localization and navigation. Inferring hidden states of the environment, and tracking them over time are implicit subtasks in this process. In the realistic case of noisy or incomplete input stimuli, and only partial knowledge of the previous state, tracking of state estimates can be improved through a probabilistic transition model that describes the evolution of states over time. How such models can be encoded and learned in spiking neural networks is an important open question. Recently, different types of spiking Winner-Take-All (WTA) networks, which are inspired by canonical cortical microcircuits, have been proposed to implement Finite State Machines (FSMs) and Hidden Markov Models (HMMs) for tracking hidden states. Here we present a Spike-Timing-Dependent Plasticity (STDP)-based framework that can simultaneously learn to abstract hidden states from sensory inputs, and learn transition probabilities between these states in recurrent connection weights. We show that this model can track hidden states over time, and recover from erroneous state transitions introduced by transient noise and stochastic spiking dynamics. In particular, we show how the model can encode probabilistic automata in which state transitions are actively triggered by signals that indicate actions, but are otherwise uninformative of the previous or subsequent state. Such networks represent generative models from which probabilistic samples of likely states under a given policy can be obtained, even in the absence of sensory input.

### III-35. Unsupervised identification of excitatory and inhibitory populations from multi-cell recordings

Lars Buesing

LARS@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

The interaction between pools of excitatory (E) and inhibitory (I) neurons critically shapes the activity dynamics of cortical neural populations. It is therefore of great interest to experimentally characterize the dynamical properties of the E and I sub-populations (e.g. interaction efficacies, time constants, noise correlations) of a local cortical population in addition to characterizing their responses to external stimuli. Such properties can in principle be estimated from multi-electrode array or Ca-recordings. However, there are two important obstacles. First, the neurons labels (i.e. if a neurons is E or I) are often unknown. Second, even with modern, large-scale recording techniques, only a small fraction of a local neural population can be recorded from, resulting in the fact that the observed activity is dominated by unobserved processes instead of direct, pair-wise interactions (sometimes called "common input problem"). Here we present a probabilistic generative model for cortical population recordings that addresses these two problems and that yields an unsupervised clustering of the observed neurons into E and I populations. Spike counts are modelled as Poisson variables whose log-intensities are generated by a latent, low-dimensional Linear Dynamical System (LDS). This LDS is a discrete-time, stochastic, linearized version of a (possibly multi-dimensional) Wilson-Cowan system describing the activity of interacting E and I populations. We show in proof-of-principle simulations that based on our model we able to jointly infer the unobserved LDS dynamics and the cell labels from artificial data with no model mismatch. As an additional, technical contribution, we present a method for doing variational inference in this model, which includes as a sub-routine variational inference for Linear Dynamical Systems (LDSs) with exponential family observations. The latter is of practical interest for a broad class of probabilistic models in computational neuroscience.

### III-36. The size-weight illusion is not anti-Bayesian after all

Megan A. K. Peters

PETERS.MEGAN@GMAIL.COM

Ladan Shams

LSHAMS@PSYCH.UCLA.EDU

University of California, Los Angeles

When we lift two objects of the same weight but different sizes, we expect the larger to be heavier, but the smaller feels heavier. This Size-Weight Illusion (SWI) cannot be explained by differential motor forces, and has been labeled 'anti-Bayesian' because it opposes simple Bayesian predictions utilizing a 'larger is heavier' prior. However, we report that this previously-considered Bayesian framework for the SWI neglects crucial information available to observers. We present a competitive prior model, in which objects' expected weight relationship is based not only on their size difference, but also on their inferred density relationship. Until now, studies of the SWI have assumed observers believe two similar-looking (same visual material) objects have equal density. However, our previous work demonstrates this assumption is flawed: smaller objects in the environment are denser than larger objects, and humans' perceptual systems appear cognizant of this relationship. Therefore, we postulate that when faced with two similar-looking objects, the nervous system examines competing hypotheses regarding their relative densities. We measured observers' (1) expectations under 'equal density' and 'smaller is denser' scenarios and (2) SWI magnitude. The model reproduces the basic illusion, and explains two findings reported in the literature: (1) illusion magnitude grows with larger size differences, and (2) after training with 'small-heavy' and 'large-light' objects, illusion magnitude is attenuated and ultimately reversed. Moreover, the model predicts (a) the degree to which an individual believes smaller objects are denser correlates with his SWI magnitude; and (b) decreased haptic precision will decrease SWI magnitude. Our psychophysical experiments validated both predictions. These findings demonstrate that the SWI is not anti-Bayesian as previously claimed, unifying the SWI with the sensorimotor literature and demonstrating that it, too, can be explained by competitive-prior Bayesian inference that can account for other illusions such as the ventriloquist and sound-induced flash illusions.

### III-37. How well can saliency maps tell where people look

Matthias Kummerer<sup>1,2</sup>  
 Matthias Bethge<sup>1,2</sup>

MATTHIAS.KUEMMERER@BETHGELAB.ORG  
 MATTHIAS@BETHGELAB.ORG

<sup>1</sup>University of Tuebingen

<sup>2</sup>Bernstein Center for Computational Neuroscience

Recently it has been suggested to use spatial point processes for modeling saliency. Here, we show how the framework of point processes can be used for the evaluation of saliency maps offering several advantages over other measures previously used. More specifically, we show how arbitrary saliency maps can be translated into point process models for which we can compute and compare the model likelihoods on measured sequences of fixations. Using this approach we reevaluate several influential models based on their likelihoods and compare the results to previous evaluations. Using the same dataset as Judd2009 we first show that the total amount of mutual information that can be extracted from an image about the spatial structure of fixations yields 2.2 bits/fixation. The best performing model can explain 65% of this total information. For comparison, we also fit a most simple model which ignores the image content altogether but only captures the observers' center bias. Remarkably, this model already explains 54% of the total information showing that all saliency models are very limited in capturing any structure beyond the center bias. In a second step, we extend the point process modeling into the temporal domain. By doing so, we can show that the often assumed mechanism of inhibition of return would rather decrease performance. Instead, by introducing self-reinforcing point processes we can increase the likelihood of the models by 15% which is just as much information as the best saliency models can extract from the content of an image. In conclusion, likelihood comparisons offer a new way of systematically dissecting how much information about the location of fixations is provided by different types of information. In particular, our approach reveals that most of the information inherent to the content of an image is not captured by any of the saliency models yet.

### III-38. Sampling orientations at different contrasts

Agnieszka Grabska-Barwinska  
 Peter Latham

AGNIESZKA@GATSBY.UCL.AC.UK  
 PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

How neurons represent probability distributions is a matter of debate. Two major hypotheses are: (i) neurons code for parameters of a probability distribution; and (ii) they represent virtually any kind of distribution by sampling. The parametric coding hypothesis is well established, with the probabilistic population code as its most popular instantiation (Ma et al. 2006). In contrast, sampling is a relatively new method. It proposes that neural activity represents samples from a probability distribution over some variable of interest. In vision, the variable of interest is often the intensity of an image component; e.g. a Gabor patch (Hoyer & Hyvarinen 2002, Orban & Lengyel 2011). However, this is inconsistent with the observation that orientation tuning is contrast-invariant (Pouget et al. 2013), since the probability distribution, and thus the tuning curves, should broaden as contrast decreases. It is also inconsistent with the observation that neurons with similar receptive fields have similar responses, since neurons with similar receptive fields are expected to silence each other because of 'explaining away'. We address both criticisms of sampling by redefining what neurons code for. We propose that an orientation selective V1 neuron (in particular, its membrane potential) provides samples for a range of orientations, rather than a single one, as suggested before. This is in agreement with the standard view that neurons report the projection of an image onto their receptive field. Under this scheme, contrast-invariant orientation response arises naturally in broadly selective neurons (in a biologically-plausible range of membrane potential tuning width) and there is no 'explaining away' (there are no strong inhibitory interactions between similar neurons). Our proposal makes the strong experimental prediction that membrane potential is not quite contrast invariant. Instead, narrow tuning curves should broaden as contrast increases and broad tuning curves should narrow.

### III-39. Sampling methods for spike inference from calcium imaging data

Eftychios Pnevmatikakis  
Josh Merel  
Liam Paninski  
Columbia University

EFTYCHIOS@STAT.COLUMBIA.EDU  
JSM2183@COLUMBIA.EDU  
LIAM@STAT.COLUMBIA.EDU

Calcium imaging is an increasingly popular technique for large scale data acquisition in neuroscience. A key problem in the analysis of calcium imaging data is the extraction of the spike times from the noisy calcium signal which in addition is sampled at a low rate and has slower dynamics compared to the fast dynamics of neural spiking. In this paper we propose Bayesian methods for sampling from the joint posterior distribution of the spike train and the problem parameters given the fluorescence observations. We present two efficient methodologies for spike sampling. The first is a discrete time binary sampler that samples whether a spike occurred at each timebin using either Gibbs sampling or Hamiltonian Monte Carlo. By exploiting the weak interaction between spikes at distant timebins, we show that a full sample can be obtained with just  $O(T)$  complexity, where  $T$  is the number of timebins, and that parallelization is also possible. Our second sampler operates in continuous time and samples the number of spikes and the spike times at arbitrary resolution using Metropolis-Hastings techniques. We use a proposal distribution for shifting the spike times around a local neighborhood that is based on the resulting signal residual. This proposal distribution enables fast mixing and tractable inference; each full sample is obtained with just  $O(K)$  complexity where  $K$  is the total number of spikes, rendering this algorithm particularly efficient for recordings that are sparse and/or are obtained at a fine resolution. Moreover, in high-SNR conditions it enables spike super-resolution within each timebin. Our methods provide a principled approach for estimating the posterior distribution of the spike trains and the model parameters, especially in low-SNR conditions, and can also be used in cases where the calcium measurements are intermittent or are available in the form of random projections.

### III-40. Predictive balance of cortical membrane potentials by STDP of inhibition

Christian Albers  
Maren Westkott  
Klaus Pawelzik  
University of Bremen

CALBERS@NEURO.UNI-BREMEN.DE  
MAREN@NEURO.UNI-BREMEN.DE  
PAWELZIK@NEURO.UNI-BREMEN.DE

The dynamical state of early sensory areas of cerebral cortex remains poorly understood. In particular, neither the functional roles of the detailed balance of excitation and inhibition in cortex are known nor have been identified the physiological mechanisms that achieve and maintain this peculiar effect. Here we investigate the hypothesis that the cortex self-organizes into a state at which predominantly deviations from expected inputs are signaled and processed while predictable inputs become suppressed. We show that this operating point is a natural consequence of spike-timing-dependent plasticity of inhibitory synapses (ISTDP), when it acts in concert with postsynaptic after-hyperpolarization and synaptic scaling. We rigorously prove that this mechanism realizes optimal associative learning on the basis of single spikes (Rosenblatt's Perceptron-learning), which generalizes into the temporal domain (Chronotron-learning). Furthermore, we demonstrate in simulations that ISTDP generally tends to balance all predictable membrane potential fluctuations of the postsynaptic cell. In other words, we find evidence that with ISTDP spiking neuronal networks naturally self-organize towards an associative mode of operation as well as sparse predictive coding, depending on the structure of the inputs and the neuron's membrane potential dynamics. Taken together our theory predicts that the cerebral cortex tunes itself to a dynamical state in which it is particularly sensitive to non-predictable events.



### III-41. Comparing roles of feedback and feedforward inhibition in sparsening of sensory codes

Pavel Sanda<sup>1</sup>  
 Tiffany Kee<sup>1</sup>  
 Nitin Gupta<sup>2</sup>  
 Mark Stopfer<sup>3</sup>  
 Maxim Bazhenov<sup>1</sup>

SANDA@UCR.EDU  
 TGLENO02@UCR.EDU  
 NITINIITK@GMAIL.COM  
 STOPFERM@MAIL.NIH.GOV  
 MAKSIMB@UCR.EDU

<sup>1</sup>University of California, Riverside

<sup>2</sup>NIH-NICHD

<sup>3</sup>National Institutes Health

Sparse sensory codes offer several advantages, particularly in the context of associative learning. Kenyon cells of the mushroom body (MB), an area of the insect brain thought to be involved in olfactory learning, show remarkably sparse responses: a given odor elicits very few action potentials, from very few cells. KCs receive excitatory input from the projection neurons (PNs) of the antennal lobe (AL). PNs also project to another area of the brain, the lateral horn (LH). In locusts, feedforward inhibition from the LH to the KCs was thought to maintain the sparseness of KC responses across a wide range of odor concentrations. New experimental evidence, however, shows that the primary source of inhibition to KCs is a giant GABAergic neuron (GGN). KCs provide excitatory input to GGN, which, in turn, inhibits the KCs, forming a negative feedback circuit that could also regulate the sparseness of KCs. To examine whether the feedback inhibition would have different functional properties compared to the feedforward inhibition, we devised a network model of the locust olfactory system, including local neurons of the AL, PNs, KCs and GGN. The properties of GGN and other neurons in the model were tightly constrained by *in vivo* experimental data. We simulated two variants of the model, with feedforward or feedback inhibition to the KCs. We found that both feedforward and feedback models could maintain the sparse responses of KCs across a range of odor concentrations and provide phase-locking of spikes with the local field potential, as observed experimentally. The feedback model, however, produced a more stable discrimination of odors than the feedforward model. The feedback model reached optimal classification within 300 ms of stimulus presentation, consistent with the experimental observations. Our results illuminate the conditions in which feedback inhibition is advantageous for maintaining sparse sensory responses.

### III-42. Persistent activity through multiple mechanisms in a spiking network that solves DMS tasks

Jason Fleischer  
 The Neurosciences Institute

FLEISCHER@NSI.EDU

We have constructed a large-scale spiking neural simulation that displays stimulus-related persistent activity during visual delayed match-to-sample (DMS) tasks. This persistent activity enables multi-item working memory and forms the basis of decision making during DMS tasks. The network incorporates three known biological mechanisms for generating persistent activity: (1) dense recurrent connectivity (Amit and Brunel, 1997), (2) short-term synaptic plasticity (Mongillo et al., 2008), and (3) NMDA receptors (Wang, 1999). In real nervous systems, all three of these mechanisms likely interact to support persistent activity. Therefore we characterized how parameters in the simulation that control these mechanisms affect persistent activity, and we found that the most robust persistent activity resulted from parameter values that reflect our current understanding of connectivity between pyramidal cells in cortex. To enable DMS task decision making, we constructed a matching mechanism based upon a segregation of visual and memory-related inputs onto AMPA and NMDA receptors of postsynaptic neurons, and we examined the ability of the network to perform visual DMS tasks. This matching mechanism is simple, effective and consistent with the experimental evidence for silent synapses and the unequal distribution of AMPA and NMDA receptors in different cortical layers. We characterized the capacity of the network to store multiple items presented serially. We also examined the relationship between the serial position of a stimulus

and the persistence of its activity pattern. We found that depending on parameters the network exhibited either a primacy or a recency effect but not both. Amit, D. J., and Brunel, N. (1997). *Cereb Cortex* 7, 237–252. Mongillo, G., Barak, O., and Tsodyks, M. (2008). *Science* 319, 1543–1546. Wang, X. J. (1999). *Journal of Neuroscience* 19, 9587–9603.

### **III-43. Oscillations emerging from noise-driven networks with electrical synapses and subthreshold resonance**

Claudia Clopath<sup>1</sup>  
Tatjana Tchumatchenko<sup>2</sup>

C.CLOPATH@IMPERIAL.AC.UK  
TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

<sup>1</sup>Imperial College London

<sup>2</sup>Max Planck Institute for Brain Research

Oscillations play a critical role in cognitive phenomena and have been observed in many brain regions, but their emergence from the interplay between the single biophysics and the connectivity has only recently attracted attention. Experimental evidence indicates that classes of neurons that engage in global network oscillations exhibit many properties that could promote oscillations and synchronization. One of them is subthreshold resonance, in which the voltage response to oscillating input peaks at preferred input frequencies. Subthreshold resonance has been reported particularly in inhibitory neurons in a variety of regions, many of which support oscillations. Furthermore, inhibitory interneurons are key players in global oscillations are connected through strong electrical gap-junctions that are implicated in network synchrony. Typically, these two properties are studied separately and it is not clear which is the dominant determinant of global network rhythms. The aim of the present work is to provide an analytical understanding of how these two effects destabilize the fluctuation-driven state and lead to an emergence of global synchronous oscillations. We propose an analytically tractable framework to address both phenomena and clarify how the oscillation frequency and amplitude of synchronous oscillations is shaped by single neuron resonance, electrical and chemical synapses when emerging from a fluctuation-driven network state. Both, the presence of gap-junctions and subthreshold resonance are necessary for the emergence of synchronous oscillations from the fluctuation-driven state. Our results are in agreement with several experimental observations and offer a much-needed conceptual link connecting a collection of disparate effects observed in networks of neurons.

### **III-44. Self-supervised learning of neural integrator networks**

Erik Nygren  
Robert Urbanczik  
Walter Senn

NYGREN@PYL.UNIBE.CH  
URBANCZIK@PYL.UNIBE.CH  
SENN@PYL.UNIBE.CH

University of Bern

Many functions of the brain require information to be kept or accumulated over long timescales, such as storing the current position of the eyes or accumulating evidence in a decision making process. The fact that a single neuron is very forgetful (~10ms memory) suggests that there must exist neural structures in the brain which can keep information for much longer timescales. Although various neuronal recordings suggest the existence of such integrator networks, we still have no profound understanding about how they integrate and how they are formed. In fact, to achieve a stable integration, a fine tuning of the recurrent weight matrix (< 0.5% accuracy) is required, or additional stabilization mechanisms have to be introduced. Here we consider a different approach to obtain stable neural integrators. Instead of imposing a specific wiring matrix to a neural network, we use a self-supervised learning paradigm for the feedforward and recurrent synapses projecting to the dendrites of 2-compartment model neurons. The activity of these neurons are shaped by somatic teacher input, and the dendritic input learns to reproduce the somatic firing. The neurons are trained by projecting the derivative of a signal to the

dendrite, and the original signal to the soma. After learning, the original signal is successfully reproduced from its derivative projected to the dendrite. Hence, the neurons learned to temporally integrate the dendritic input and represent this integral by the instantaneous somatic firing rate. This self-supervised learning paradigm forms stable neural integrators with a minimal number of neurons ( $\sim 200$ ) using only AMPA conductances. The plasticity paradigm represents a candidate for a biologically plausible learning in a recurrent network.

### III-45. Nonlinear threshold dynamics implements multiplicative adaptation in pyramidal neurons

Christian Pozzorini  
Skander Mensi  
Oliver Hagens  
Wulfram Gerstner

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

Ecole Polytechnique Federale de Lausanne

At odds with the predictions of standard spiking neuron models, neocortical pyramidal neurons have recently been shown to maintain high sensitivity to input fluctuations even at large suprathreshold inputs. Overall, this phenomenon remains poorly understood. Here, we recorded the in vitro response of fast-spiking and pyramidal neurons to fluctuating currents generated by varying the magnitude of both the input mean and variance. In contrast with fast-spiking neurons, pyramidal neurons saturated at relatively low firing rates but maintained sensitivity to input fluctuations over the entire range of mean input that were tested. A simple analysis of the experimental data revealed that, in pyramidal neurons, firing threshold and firing rate were positively correlated. However, while evoking larger responses, we surprisingly found that increased input fluctuations were always associated with reduced firing thresholds. To understand this behavior and study its implications on neural coding, we designed a new spiking model in which the firing threshold is nonlinearly coupled with the subthreshold membrane potential and also linearly dependent on the spiking history. All model parameters were extracted from the intracellular recordings using a new nonparametric maximum likelihood method. The model captured the experimental data and was able to accurately predict the occurrence of individual spikes. Moreover, the extracted threshold dynamics was consistent with a biophysical model accounting for fast and slow sodium-channel inactivation. Our results demonstrate that a nonlinear interaction between spike-dependent and spike-independent threshold movements adaptively modifies the effective timescale on which single neurons operate, providing an explanation for the enhanced sensitivity of pyramidal neurons to input fluctuations. Consistent with this explanation, the firing threshold of fast-spiking neurons did not depend on the spiking history. Overall our results demonstrate and explain how, depending on the input statistics, neocortical pyramidal neurons adapt their coding strategy operating either as integrators or differentiators.

### III-46. Asymmetry of electrical synapses: measurement and impact in small networks

Julie Haas  
Joe C Brague  
Jessica Severson  
Lehigh University

JULIE.HAAS@GMAIL.COM  
JOEBRAGUE@GMAIL.COM  
JESS7SON@GMAIL.COM

Electrical synapses are often assumed to exert synchronizing influences within neuronal networks. However, recent studies have demonstrated the potential for these synapses to instead desynchronize coupled networks (Vervaeke et al., 2010). When combined with chemical synapses, the effects of electrical synapses of the output patterns of networks can be dynamically diverse. In addition, it has been shown that full rectification of signals by electrical synapses can dramatically alter the parameterscape, or the full complement of behaviors expressed by

a network (Guitierrez and Marder, 2013). Here we characterize asymmetry, or rectification, in electrically coupled neurons of the thalamic reticular nucleus (TRN), which provide the main source of inhibition to thalamocortical relay cells. We demonstrate the functional impact of asymmetry on spike timing in coupled neurons using dual-cell recordings, and show that this impact is often substantial, leading to variations in spike time onset of tens of ms. We introduce a measure of the impact of asymmetry on spike timing. Finally, we evaluate the effects of electrical synaptic asymmetry in a model of two cells coupled by an electrical synapse, which we provide with inputs of varied amplitude and timing. Together, these results demonstrate that asymmetry of electrical synapses alters the dynamic range and computational role that these synapses play within a neuronal network. In the TRN, the directionality of electrical synapses may serve as a tuned and plastic basis for the control of intra-TRN excitation.

### **III-47. Balanced networks in space**

Robert Rosenbaum  
Jonathan Rubin  
Brent Doiron

RJROSENBAUM@GMAIL.COM  
JONRUBIN@PITT.EDU  
BRENT.DOIRON@GMAIL.COM

University of Pittsburgh

Networks of cortical neurons exhibit strong spatiotemporal variability. Balanced network models – characterized by a dynamically stable balance between strong excitatory and inhibitory synaptic currents – offer an appealing theoretical framework for studying this neural variability since they produce intrinsically noisy and asynchronous dynamics with statistical features similar to those observed in cortical recordings. Previous studies of balanced networks assume a spatially homogeneous or discretely clustered topology, but it is well-known that the probability of connection between cortical neurons decays with their separation in physical space or, for some sensory systems, feature space. The question of how and whether networks with spatially dependent connection probabilities can achieve a balanced, asynchronous state has not been previously addressed. We derive experimentally testable conditions on the strength and spatial profile of connection probabilities that must be satisfied for a recurrent network of excitatory and inhibitory neurons to maintain a balanced asynchronous state in the limit of large network size. Specifically, we find that external inputs must be broader than recurrent excitation which in turn must be broader than recurrent inhibition, suggesting that classical network models relying on broad recurrent inhibition cannot coexist with a balanced state. We show that, despite requiring localized inhibition, balanced networks can produce a functional surround-suppression in which the spatial profile of a neural response is narrower than that of external input. Furthermore, we show that chaotic spatiotemporal dynamics that emerge in spatially balanced networks can give rise to moderate trial-to-trial correlations in contrast to spatially homogeneous balanced networks which exhibit extremely small correlations. This observation could potentially resolve some aspects of the ongoing debate over the magnitude of correlations in the spontaneous state.

### **III-48. Derivation of leaky integrate-and-fire and STDP from a signal processing cost function**

Dmitri B Chklovskii<sup>1,2</sup>  
Tao Hu<sup>3</sup>

MITYA@JANELIA.HHMI.ORG  
TAOHU@TEES.TAMUS.EDU

<sup>1</sup>Janelia Farm Research Campus

<sup>2</sup>Howard Hughes Medical Institute

<sup>3</sup>Texas A&M University

While much is known about physiological properties of a neuron, its computational role is poorly understood. In the absence of such understanding, modeling neuronal function on a biophysical level requires measuring prohibitively many physiological parameters. Previous work addressed the computational role of either activity dynamics (Yu, 2007, Pfister et al., 2010, Boerlin and Deneve, 2011) or synaptic learning (Oja, 1982, Hyvarinen

and Oja, 1998, Rao and Sejnowski, 2001, Bohte and Mozer, 2004, Hennequin et al., 2010, Clopath and Gerstner, 2010) but not both at the same time. Here, we make a step towards assembling a unified algorithmic framework of a spiking-neuron computation that can account for both kinds of neurophysiological properties. Such framework should allow modeling the function of neural circuits on a higher level of abstraction, which does not require measuring all the physiological parameters, as well as enable designing neuromorphic electronics. We derive mathematically both activity dynamics and synaptic learning by treating a neuron as a robust online signal processing device. Specifically, we postulate that a neuron represents the spike trains from neurons presynaptic to it by the corresponding synaptic weights times the outgoing spike train convolved with a temporal kernel. We formalize this postulate in a cost function that combines cumulative squared representation error and regularizers of synaptic weights and neuronal activity. Such cost function can be minimized online by alternating minimization steps with respect to neuronal activity, synaptic weights, and the temporal kernel. Amazingly, these steps correspond, respectively, to the leaky integrate-and-fire dynamics, the spike-timing dependent plasticity (STDP) time course, and the temporal kernel learning. In this framework, the STDP time course mirrors cross-correlation of pre- and postsynaptic spike trains. Our theoretical framework makes several predictions, several agreeing with experimental observations, others awaiting verification.

### III-49. Summation and phase decorrelation of odor responses in an olfactory bulb microcircuit model

Aditya Gilra  
Upinder Bhalla

ADITYAG@NCBS.RES.IN  
BHALLA@NCBS.RES.IN

National Centre for Biological Sciences

The vertebrate olfactory bulb is a primary sensory region that receives olfactory receptor input and transforms it via a unique two-layer inhibitory dendro-dendritic circuitry whose role in odor coding is still unclear. Using a detailed microcircuit model of the olfactory bulb we show that the periglomerular-cell layer implements linear summation, while the granule-cell layer implements respiratory phase decorrelation of odor responses. Both these computations are mechanistically distinct from previous studies. We used large-scale parallel simulations, to model three data-intensive single-unit in-vivo tetrode recording studies comprising responses to pulsed multi-odor stimuli at different concentrations, pairwise comparisons of respiratory phase-responses, and morphed odor responses in freely-breathing animals. Thus our model directly draws upon, and provides a mechanistic explanation for large-scale physiological data. Our model is built in a hierarchical manner, using detailed anatomical data, intra-cellular single- and dual-cell recordings, and the afore-mentioned large-scale in-vivo recordings. We incorporate detailed morphology in our model to predict that mitral lateral dendrites are output structures and deliver inhibition to distant mitral cells via granule cell columns around the distant cells' somas. This long-range inhibition forms the mechanism for the observed phase-decorrelation of sister mitral cell responses (Dhawale, et al. 2010). We estimate that there are up to 20 of these strongly inhibiting connections per mitral cell. In contrast, the periglomerular cells play a key role in our model in linearizing the mitral input-output transformation to account for the observed linear summation of odor responses (Gupta and Bhalla, in revision, Khan et al. 2008). We feel that our study will enliven some long-standing debates about olfactory bulb computations, while also serving as an illustration for harvesting the kind of large-scale data that is emerging in neuroscience.

## III-50. Dynamics of olfactory bulb output generates concentration-invariant cortical odor representations

Merav Stern<sup>1,2</sup>  
L.F. Abbott<sup>1</sup>  
Kevin M. Franks<sup>3</sup>

MERAV.STERN@MAIL.HUJI.AC.IL  
LFA2103@COLUMBIA.EDU  
FRANKS@NEURO.DUKE.EDU

<sup>1</sup>Columbia University  
<sup>2</sup>Hebrew University  
<sup>3</sup>Duke University

Ensembles of piriform cortex neurons likely encode both an odor's concentration and its identity, regardless of concentration. Recent studies indicate that olfactory bulb mitral and tufted cells (MTs) fire at specific times during the inhalation phase of the respiration cycle. We asked whether piriform cortex can use the temporal structure of MT spiking to generate odor-specific ensembles that can simultaneously encode both identity and concentration. We therefore developed a model that recapitulates odor-evoked MT firing patterns. Consistent with experimental findings, low odorant concentrations sparsely activate model MTs toward the end of the inhalation phase of the simulated respiration cycle. Increasing odorant concentration advances the onset of spiking in these cells and recruits additional MTs at later phases of the cycle. This information is fed into a model cortical network of excitatory (pyramidal) and inhibitory cells whose organization resembles piriform cortex. Specifically, afferent cortical inputs from MTs are random, diffuse and overlapping, and pyramidal cells form recurrent excitatory synaptic connections with both other pyramidal cells and inhibitory interneurons across the cortex. In the absence of recurrent circuitry, increasing odorant concentration recruits more pyramidal cells as the number of spiking MTs increases during the sniff. However, recurrent excitation from the earliest cortical responders suppresses cells responding selectively to weakly activated MTs that fire later in the sniff cycle by driving strong, scaled cortical feedback inhibition. Recurrent cortical circuitry thereby helps normalize the cortical response to generate a more concentration-invariant ensemble of active cortical pyramidal cells. Interestingly, recurrent circuitry also turns the steadily ramping MT input activity into a sharp onset pulse in cortex, allowing for rapid odor identification, and increases the coherence of the cortical response, especially at higher odorant concentrations. Recurrent circuitry can therefore shape activity in piriform cortex to generate a spatiotemporal representation of both odor identity and concentration.

## III-51. Inferring synaptic conductances from spike trains with a point process encoding model

Kenneth Latimer<sup>1</sup>  
E. J. Chichilnisky<sup>2</sup>  
Fred Rieke<sup>3,4</sup>  
Jonathan W Pillow<sup>1</sup>

LATIMERK@UTEXAS.EDU  
EJ@STANFORD.EDU  
RIEKE@UW.EDU  
PILLOW@MAIL.UTEXAS.EDU

<sup>1</sup>University of Texas at Austin  
<sup>2</sup>Stanford University  
<sup>3</sup>University of Washington  
<sup>4</sup>Howard Hughes Medical Institute

Sensory neurons respond to stimuli in a manner that depends on the integration of excitatory and inhibitory synaptic inputs. However, statistical models for spike responses tend to describe stimulus integration with a linear filter and relegate all nonlinear processing steps to a single post-integration nonlinearity. Here we introduce a novel point process model, an extension of the well-known generalized linear model (GLM), that allows us to characterize the tuning of a neuron's excitatory and inhibitory synaptic inputs from its extracellularly recorded spike responses to a visual stimulus. We validate our method using direct intracellular measurements of excitatory and inhibitory conductances in parasol retinal ganglion cells in primate retina. Our work makes two novel theoretical contributions: first, we show that the standard generalized linear encoding model is mathematically equivalent to a conductance-based model in which the stimulus linearly modulates excitatory and inhibitory conductances

in an equal and opposite “push-pull” fashion; second, we relax these assumptions to obtain a more flexible, biophysically realistic model in which conductances have distinct tuning and nonlinear stimulus dependence. The resulting conductance-based model can produce shunting as well as hyperpolarizing inhibition, and can exhibit changes in gain and membrane time constant due to changes in total conductance. We apply our method to neural responses to a full-field temporal noise stimulus. We find that the stimulus-dependence of both excitatory and inhibitory conductances can be well described by a linear-nonlinear cascade, with the filter driving inhibition exhibiting opposite sign and a slight delay relative to the filter driving excitation. We show that the model fit to extracellular spike trains can predict excitatory and inhibitory conductances with nearly the same accuracy as a model fit directly to intracellularly measured conductances.

### III-52. Emergent connectivity in a spiking model of visual cortex with attentional feedback

Philip Meier  
 Dimitry Fisher  
 Jayram Moorkanikara Nageswaran  
 Botond Szatmary  
 Eugene Izhikevich  
 Brain Corporation

PHILIP.MEIER@BRAIN CORPORATION.COM  
 FISHER@BRAIN CORP.COM  
 NAGESWARAN@BRAIN CORP.COM  
 SZATMARY@BRAIN CORP.COM  
 IZHIKEVICH@BRAIN CORP.COM

We present a hierarchical model of the ventral visual system, composed of layers of spiking neurons. The structure is inspired by the primate visual system. All synaptic weights (except Working Memory) are learned from visual experience via STDP. Area IT acts as an object classifier. The entire network is implemented in spiking neurons and trained via visual experience. The network contains spiking neurons arranged in seven layers. There are two layers ('layer 4' and 'layer 2/3') per visual area in simulated areas V1, V2, and IT; as well as a single-layer spiking working memory area connected to area IT. Distinct mechanisms of homeostasis act in individual neurons, in lateral connections between neurons, and in layer-wide inhibitory feedback. The network connectivity is emergent; the weights of synapses in the cortical layers are learned by adapting to the statistics of the visual input – naturalistic movies. Different STDP rules are used for learning feed-forward, lateral, and feedback synaptic weights. In order to explore the strengths and weaknesses of the system, we require that the network performs a delayed-match-to-sample task. The system must saccade to the location of the cued object, despite the simultaneously presented distractor. Success on this task requires sufficient visual feature learning, working memory, attention, saccadic control, and integration between these mechanisms. The spiking activity in area IT during the testing phase encodes the cued object. However, we find that it is difficult for attentional feedback to sufficiently bias V1 or V2, such that the spatial spike density induces the superior colliculus to reliably read out the position of the cued object, and ignore the distractor object. In other words, to perform the task in our network, the difference in salience of the objects must be overcome by the top-down attentional modulation.

### III-53. A spiking neural network for autoencoder learning

Kendra Burbank  
 University of Chicago

BURBANK@GALTON.UCHICAGO.EDU

Visual object recognition occurs through a hierarchical process of computations in multiple areas of visual cortex. We currently lack a good theoretical understanding of what computations occur in each area, of what computations would be optimal given the final goal of object recognition, or of how neuronal networks could learn to perform the required computations. Recently, new methods for training deep hierarchical networks have seen great empirical success in the machine learning community. These techniques involve an initial unsupervised 'pre-training' followed by a supervised fine-tuning step. In the simplest cases, the pre-training uses an 'autoencoder' learning rule.

In this work, we asked whether biological neural networks could plausibly implement such autoencoder learning. We created a two-layer neural network with feedforward and feedback (but no lateral) connections. Feedforward connections were trained using a spike-timing dependent plasticity (STDP) rule. By contrast, because feedback connections in cortex typically synapse far from the soma, we modeled feedback learning using a temporally reversed STDP rule inspired by plasticity observed experimentally at distal synapses. We showed analytically that our network will exactly implement autoencoder learning in the case of probabilistic neurons, discrete time-steps, and a simplified STDP rule. Next, we implemented our network with spiking neurons, visual inputs, and timing of activations designed to resemble activity in visual areas V1 and V2. We showed that this network learned to reconstruct its inputs, and compared the precise learning to that given by the analytical autoencoder. These results suggest that cortical networks may be well-suited for using unsupervised learning to learn representations of visual input. Moreover, they imply that machine learning concepts such as Deep Belief Networks may have direct relevance to learning in mammalian brains. Finally, they point to an important but perhaps unappreciated role for feedback in the learning of feedforward connections.

### III-54. Evidence for critical slowing down in neurons near spiking threshold

Christian Meisel<sup>1</sup>  
Andreas Klaus<sup>2</sup>  
Christian Kuehn  
Dietmar Plenz<sup>2</sup>

CHRISTIAN@MEISEL.DE  
AKLAUS.MAIL@GMAIL.COM  
CK274@CORNELL.EDU  
PLENZD@MAIL.NIH.GOV

<sup>1</sup>National Institute of Mental Health

<sup>2</sup>National Institute of Mental Health/NIH

The transition of a neuron from quiescence to spiking constitutes an abrupt shift in its dynamics. The dynamical systems' approach has yielded significant insights into the underlying dynamics identifying two main mechanisms, saddle-node and Hopf bifurcation, for the transition from subthreshold activity to spiking. In both cases, the system is expected to recover more slowly from small perturbations upon approaching the bifurcation point, a phenomenon called critical slowing down. Although theory provides precise predictions on signatures of critical slowing down when the bifurcation point to spiking is approached, experimental evidence has been lacking. Here we provide the first experimental evidence for signatures of critical slowing down in cortical pyramidal neurons near their spiking threshold. Using whole-cell patch-clamp recordings in acute slices of rat cortex, we show that the recovery rate of the subthreshold membrane potential to small perturbations decreases while variance and autocorrelation increase upon approaching the spiking threshold amounting to conclusive evidence for critical slowing down. We then compare the scaling of these metrics to the scaling relations expected for saddle-node and Hopf bifurcation. Theory predicts power-law scaling for variance and recovery rate with exponents 1/-1 for Hopf and 0.5/-0.5 for saddle-node bifurcation. In contrast, we experimentally find an exponent of 1 for the increase in variance and -0.5 for the decrease in recovery rate. Finally, we show that all signatures of critical slowing down are completely abrogated when blocking voltage-gated sodium channels by bath application of tetrotoxin. Our findings are therefore twofold: 1) the transition to spiking is governed by a critical transition exhibiting slowing down and 2) the precise scaling during critical slowing down is insufficiently described by either saddle-node or Hopf bifurcation alone. Together these results provide novel insights into our understanding of the bifurcation structure that underlies the transition to spiking in single neurons.



## III-55. Self-organization in balanced state networks by STDP and homeostatic plasticity

Felix Effenberger<sup>1</sup>  
 Anna Levina<sup>2,3</sup>  
 Jurgen Jost<sup>4</sup>

FELIX.EFFENBERGER@MIS.MPG.DE  
 ANNA@NLD.DS.MPG.DE  
 JOST@MIS.MPG.DE

<sup>1</sup>Max Planck Institute for Mathematics in the Sciences

<sup>2</sup>Bernstein Center for Computational Neuroscience, Goettingen

<sup>3</sup>Max Planck Institute

<sup>4</sup>MaxPlanck Institute Mathematics in the Sciences

Structural inhomogeneities have a strong impact on population response dynamics of cortical networks and are believed to play an important role in their functioning. However, little is known about how such inhomogeneities could evolve in self-organized way by means of synaptic plasticity. Here, we present an adaptive model of a balanced neuronal network that combines two different types of plasticity, STDP acting on excitatory and inhibitory synapses and synaptic scaling. The plasticity rules yield both long-tailed distributions of synaptic weights and firing rates, that was previously not achieved in biologically realistic models. Although the structure in the weight distribution of the synaptic weights changes dramatically, the network resides in the balanced state without the necessity to fine tune parameters throughout the simulation. A symmetry breaking in the synaptic weight distributions leads to the appearance of a highly connected subnetwork of neurons with strong synapses that we call driver neurons. Coincident spiking activity of several driver cells can evoke population bursts and driver cells have similar dynamical properties as leader neurons found experimentally. We find that already two driver cells firing together lead to a causal increase in the network's activity and a firing of the whole group leads to a large network response. Contrary to previous studies, where the strong synapses arising due to a temporary symmetry breaking disappear after some time, our driver group remains stable. We show that local topological properties of the network define the probability of a neuron to become a driver cell. Our model is simple, robust to parameter changes and able to explain a multitude of different experimental findings in one basic network. It allows to observe the interplay between structural and dynamical properties of the emergent inhomogeneities and makes many predictions for future experimental work.

## III-56. Predicting in vivo statistics from in vitro cell parameters in a detailed prefrontal cortex model

Joachim Hass<sup>1,2</sup>  
 Loreen Hertaeg<sup>1,2</sup>  
 Tatiana Golovko<sup>1,2</sup>  
 Daniel Durstewitz<sup>1,2</sup>

JOACHIM.HASS@ZI-MANNHEIM.DE  
 LOREEN.HERTAEG@ZI-MANNHEIM.DE  
 TATIANA.GOLOVKO@ZI-MANNHEIM.DE  
 DANIEL.DURSTEWITZ@ZI-MANNHEIM.DE

<sup>1</sup>Bernstein Center for Computational Neuroscience, Heidelberg-Mannheim

<sup>2</sup>CIMH Mannheim

A major challenge in computational neuroscience is to derive neuronal network models that are on the one hand physiologically realistic, thus allowing reasonable inferences on their biological counterparts, yet on the other hand are still computationally tractable in terms of simulation time, fitting effort and mathematical analysis. Towards this goal, we constructed a computational network model of the prefrontal cortex (PFC) based on simple neuron elements, yet equipped with highly realistic anatomy and close fitting of cellular parameter distributions derived from an extensive data base of prefrontal in vitro recordings. The resulting model was reasonably fast to simulate and fit to data, yet exhibits a high degree of physiological validity, allowing for quantitative predictions. Single neurons were modeled by a simplification of the adaptive exponential integrate-and-fire model (simpAdEx) that allows for fast, efficient, and predictive fitting to large sets of standard in vitro recordings. ~200 pyramidal cells, fast-spiking and bitufted interneurons from rodent PFC were used to generate a distribution of model cells that reflects the diversity of neurons in the real PFC. These cells were embedded in a network with columnar and laminar orga-

nization and parameter distributions for connectivity, synaptic weights, kinetics and short-term plasticity extracted from an extensive literature research. Even without any tuning, the network model approximates a number of spike statistics extracted from in-vivo data from awake behaving rat PFC. However, the reproduction of in-vivo-like properties depended sensitively on how well cellular parameter distributions of the model were matched to the multivariate distributions obtained empirically. In particular, more homogeneous neuron parameters, as often assumed in network models, led to highly non-physiological statistics. These results highlight the importance of the diversity of neurons for the generation of realistic network dynamics. Support: BMBF, 01GQ1003B and DFG, Du 354/7-2 & 6-1.

### **III-57. Firing-rate dynamics from spiking network models with low-rank connectivity**

Brian DePasquale  
Mark Churchland  
L.F. Abbott

Columbia University

BDD2107@COLUMBIA.EDU  
MC3502@COLUMBIA.EDU  
LFA2103@COLUMBIA.EDU

Neuronal activity recorded in awake animals exhibits two broad features: irregular spiking and smooth firing-rate modulations that correlate with behavior. Irregular firing can be generated in spiking models with random connectivity (van Vreeswijk & Sompolinsky, 1996) and rate modulations related to task performance can arise in firing-rate models (Jaeger and Haass, 2004). Here we develop and study networks of spiking neurons that display both of these features. Smooth firing-rate dynamics can be generated in networks of irregularly spiking integrate-and-fire neurons by adding a low-rank (LR) connectivity structure to a randomly connected network (Eliasmith, 1996; Boerlin et al., 2013; Litwin-Kumar & Doiron, 2012). We use an iterative algorithm (Sussillo & Abbott, 2009) to train the LR connections. This generates smooth firing-rate modulations that allowed the network to perform interesting tasks without disrupting the irregular spiking produced by the random full-rank connectivity. Chaotic spiking in these networks generates trial-to-trial fluctuations that are suppressed by external stimuli, consistent with experimental findings (Churchland et al. 2009). As the rank of the LR connectivity is increased, more complex tasks can be performed. For example, 1000 spiking neurons can be trained to produce about 25 oscillatory outputs and to switch between them on the basis of network inputs. Interestingly, this network generates chaotic firing-rate fluctuations in the absence of input. This work leads to a novel interpretation of firing-rate models and smoothly varying signals in general. In these spiking models, the smoothly fluctuating signals can be read out from the low-rank connectivity that generates them. These signals are similar to what would be extracted from in vivo recordings using PCA or other dimensionality reduction techniques and would normally be modeled using firing-rate networks. This suggests that firing-rate models are not describing the activity of spiking neurons but rather signals arising from the connectivity between spiking neurons.

### **III-58. CA1 activity sequences emerge after reorganization of network correlation structure during learning**

Mehrab Modi<sup>1</sup>  
Ashesh Dhawale<sup>2</sup>  
Upinder Bhalla<sup>1</sup>

<sup>1</sup>National Centre for Biological Sciences

<sup>2</sup>Harvard University

MEHRAB@NCBS.RES.IN  
DHAWALE@FAS.HARVARD.EDU  
BHALLA@NCBS.RES.IN

Animals can learn causal relationships between pairs of stimuli separated in time, an ability that depends on the hippocampus<sup>1</sup>. We monitored hippocampal CA1 activity using two photon calcium imaging, as naive mice were trained on a hippocampus-dependent, temporal-memory task to associate a tone and an air-puff spaced in time.

During learning, area CA1 neurons converged to sequences of activity that tiled the interval between the stimuli. We also observed learning related changes in correlations in spontaneous activity between neuron-pairs and the spatial organization of these correlated pairs that might underlie the emergence of sequential activity during the tone-trace-puff period. Trace eyeblink conditioning is a hippocampus dependent task<sup>2</sup>, where subjects learn to associate a neutral tone (conditioned stimulus - CS) with an aversive air-puff to the eye (unconditioned stimulus - US) presented with a delay after the end of the tone. In our view, the observed activity sequences constitute a representation of the CS that bridges the CS-US temporal gap. Interestingly, cells sharing the same timing in the sequence showed an increase in spontaneous activity correlations that peaked prior to the peak in behavioral, conditioned responses. Furthermore, the relative locations of co-time-tuned cells in the field of view were random. On the other hand, considering spontaneous correlations across the population, neurons could be classified into spatially-organized clusters sharing high within-cluster correlations. Interestingly, these correlation clusters, while stable prior to training, underwent repeated changes during learning. We propose that the learning-dependent re-organization of spatially organized clusters reflect coarse, network-level connectivity changes that underlie the emergence of randomly distributed cell groups that have behaviorally relevant, stimulus or context encoding properties. 1. Fortin, N. J., Agster, K. L. & Eichenbaum, H. B.; *Nat. Neurosci.* 5, 458–462 (2002). 2. Tseng, W., Guan, R., Disterhoft, J. f. & Weiss, C.; *Hippocampus* 14, 58–65 (2004).

### III-59. Input-driven unsupervised learning in recurrent neural networks

Alireza Alemi-Neissi<sup>1</sup>

Carlo Baldassi<sup>1,2</sup>

Nicolas Brunel<sup>3</sup>

Riccardo Zecchina<sup>1,2</sup>

AALEMI@GMAIL.COM

CARLOBALDASSI@GMAIL.COM

NBRUNEL@GALTON.UCHICAGO.EDU

RICCARDO.ZECCHINA@POLITO.IT

<sup>1</sup>Human Genetics Foundation

<sup>2</sup>Politecnico di Torino

<sup>3</sup>University of Chicago

Understanding the theoretical foundations of how memories are encoded and retrieved in neural populations is a central challenge in neuroscience. A popular theoretical scenario for modeling memory function is an attractor neural network with Hebbian learning (e.g. the Hopfield model). The model simplicity and the locality of the synaptic update rules come at the cost of a limited storage capacity, compared with the capacity achieved with supervised learning algorithms, whose biological plausibility is questionable. Here, we present an on-line learning rule for a recurrent neural network that achieves near-optimal performance without an explicit supervisory error signal and using only locally accessible information, and which is therefore biologically plausible. The fully connected network consists of excitatory units with plastic recurrent connections and non-plastic inhibitory feedback stabilizing the network dynamics; the patterns to be memorized are presented on-line as strong afferent currents, producing a bimodal distribution for the neuron synaptic inputs ('local fields'). Synapses corresponding to active inputs are modified as a function of the position of the local field with respect to three thresholds. Above the highest threshold, and below the lowest threshold, no plasticity occurs. In between these two thresholds, potentiation/depression occurs when the local field is above/below an intermediate threshold. An additional parameter of the model allows to trade storage capacity for robustness, i.e. increased size of the basins of attraction. We simulated a network of 1001 excitatory neurons implementing this rule and measured its storage capacity for different sizes of the basins of attraction: our results show that, for any given basin size, our network more than doubles the storage capacity, compared with a standard Hopfield network. Our learning rule is consistent with available experimental data documenting how plasticity depends on firing rate. It predicts that at high enough firing rates, no potentiation should occur.

### III-60. Normalized Hebbian learning develops both simple and complex receptive fields from naturalistic vide

Thomas Miconi<sup>1</sup>  
 Jedediah Singer<sup>2</sup>  
 Gabriel Kreiman<sup>2</sup>

THOMAS.MICONI@GMAIL.COM  
 JEDEDIAH.SINGER@CHILDRENS.HARVARD.EDU  
 GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU

<sup>1</sup>The Neurosciences Institute  
<sup>2</sup>Harvard Medical School

We describe a simple 2-layer model of a V1 column, in which neurons in each layer develop simple and complex receptive fields, respectively, using the exact same algorithm in both layers, when exposed to naturalistic video stimuli. The algorithm is based on Hebbian learning, with divisive normalization and homeostatic threshold adaptation. Neurons in layer 1 take inputs from a 17x17 patch of video input, extracted at high frequency at the point of gaze of a human viewer (using an eye-tracking device), and processed through a difference-of-Gaussians filter to model the output of thalamic cells. Neurons at layer 2 take inputs from neurons in layer 1. For all neurons, Hebbian learning modifies synaptic weights proportionally to the product of current synaptic input and recent neuron output, with a penalty for weight magnitude ('instar' rule). Normalization divides all neural responses by the sum of all responses within each layer. Finally, homeostatic threshold adaptation increases the threshold of active cells and decreases the threshold of inactive cells. By interacting with plasticity, normalization decorrelates neuron responses and produces a sparse population coding of the inputs. Using eye-tracking allows us to meaningfully sample from the video stream at a frequency higher than video frame rate. Layer 1 cells develop localized receptive fields resembling Gabor filters with various orientations, positions and phases. Meanwhile, layer 2 neurons develop complex receptive fields, with strong synapses from layer-1 cells of similar orientation at various positions and phases, and weak synapses from cells at orthogonal orientations. To our knowledge, this is the first report of simple and complex receptive fields developing using the same biologically plausible learning algorithm in response to naturalistic video data.

### III-61. A normative framework of single-neuron computation yields Oja-like learning with realistic features

Cengiz Pehlevan<sup>1,2</sup>  
 Tao Hu<sup>3</sup>  
 Zaid J. Towfic<sup>4</sup>  
 Dmitri B Chklovskii<sup>1,2</sup>

PEHLEVANC@JANELIA.HHMI.ORG  
 TAOHU@TEES.TAMUS.EDU  
 ZTOWFIC@EE.UCLA.EDU  
 MITYA@JANELIA.HHMI.ORG

<sup>1</sup>Janelia Farm Research Campus  
<sup>2</sup>Howard Hughes Medical Institute  
<sup>3</sup>Texas A&M University  
<sup>4</sup>University of California, Los Angeles

What does a single neuron compute? Neurons have long been accepted as the basic units of neural computation, but an understanding of their computational role, say, to the level of our understanding of an electronic logic gate's computational role, is still lacking. We propose to view the neuron as a device that reduces the dimensionality of its high-dimensional synaptic input. The input is represented by the synaptic weight vector scaled by the neuron's firing rate. Minimizing the regularized representation error leads to an online algorithm, which takes alternating steps in updating the neuron's firing rate and changing its synaptic weights. Hence, our algorithm addresses both neural activity dynamics and dynamics of synaptic plasticity, providing a holistic view of neural computation. These steps reproduce well-known physiological properties of neurons such as leaky-integration of weighted synaptic activity and a nonlinear output function. Synaptic weights are updated with an Oja-like plasticity rule with an activity dependent learning rate that decreases over time. Physiological synapses are given by thresholding an internal variable, which may account for the existence of silent synapses. Features of temporal evolution of synaptic weights resemble experimental data from LTP critical periods in cortex. We show that the algorithm can

learn non-Gaussian features of an input signal and lead to sparse activity.

### III-62. Sequence generation by spatio-temporal cycles of inhibition

Jonathan Cannon  
Jeffrey Markowitz  
Nancy Kopell  
Timothy Gardner  
Boston University

THEOCEANISTEA@GMAIL.COM  
JMARKOW@CNS.BU.EDU  
NK@MATH.BU.EDU  
TIMOTHYG@BU.EDU

A theory is proposed for robust sequence generation in neuronal networks. In previous papers, rhythmic activity has been characterized as an antagonist of sequence generation. In our model, rhythm plays a key role in sequence generation. The model is based on two recent observations in songbird premotor cortex: (1) interneuron and principal cell activity are locked to different phases of the 25-35 Hz band of a stereotyped LFP, and (2) principal cells are strongly coupled via disynaptic inhibition. In the model, long sequences of principal cells fire within the constraints of mesoscopic cortical dynamics. When principal cell activity reaches a cortical subregion, it triggers strong, locally diffuse feedback inhibition on a 30ms time scale. The feedback quickly suppresses local activity, but other cortical subregions are disinhibited at any given time, allowing activity to continue and propagate from one subregion to the next as each becomes disinhibited. Due to this spatio-temporal coordination, these sequences are robust to various perturbations even in the absence of redundant feed-forward connectivity, making the model fundamentally distinct from the synfire chain. The proposed mechanism produces long, stereotyped sequences even when the connections between principal neurons are randomly generated. With the addition of single cell firing rate homeostasis, each projection neuron can be made to participate sparsely. As specific sequences are reinforced, sequence propagation becomes more temporally stable and more robust against connection noise. These conclusions are demonstrated in simulation; we also heuristically explain the tolerance for connection noise, and draw on the theory of discrete dynamical systems to analytically demonstrate timing stability. The model concisely accounts for the observed single-cell and mesoscopic dynamics in songbird, and can be generalized to provide a similar explanation for traveling waves observed in primate motor cortex.

### III-63. Adaptive shaping of feature selectivity in the awake rodent vibrissa system

He Zheng<sup>1</sup>  
Douglas Ollerenshaw<sup>2</sup>  
Daniel Millard<sup>1,3</sup>  
Qi Wang<sup>4</sup>  
Garrett B Stanley<sup>1,3</sup>

NEURON@GATECH.EDU  
D.OLLERENSHAW@GATECH.EDU  
DMILLARD6@GATECH.EDU  
QW2161@COLUMBIA.EDU  
GARRETT.STANLEY@BME.GATECH.EDU

<sup>1</sup>Georgia Institute of Technology  
<sup>2</sup>Allen Institute for Brain Science  
<sup>3</sup>Emory University  
<sup>4</sup>Columbia University

Sensory processing is context-dependent. One particularly salient example of this is the way in which sensory pathways may switch from conveying information necessary for detecting novel features of the sensory environment, to conveying information necessary for transmitting fine details and/or discriminating between different sensory features [1]. While psychophysical studies have shown that sensory adaptation induces heightened spatial acuity in tactile discrimination tasks, electrophysiological and imaging studies have shown a spatially constrained or sharpened cortical representation of repetitive stimuli, proposed as a potential mechanism for the observed enhanced spatial acuity [2-5]. The relationship between the average cortical representation and the information

conveyed about the sensory input, however, is unknown [6]. Furthermore, the link between the corresponding neural representation and behavior has not been established. In the rodent vibrissa system, we conducted single-trial ideal observer analysis of cortical activity measured using voltage sensitive dye (VSD) imaging in the anesthetized animal, combined with behavioral detection and discrimination tasks, thalamic recordings from the awake animal, and computational modeling to show that spatial discrimination performance was improved following adaptation, but at the expense of the ability to detect weak stimuli. Furthermore, we show that this change in cortical representation and information flow is a malleable phenomenon, specific to the kinetic features of the adapting stimulus. Together, these results provide direct behavioral evidence for the trade-off between detectability and discriminability, that this tradeoff can be modulated through bottom-up sensory adaptation, and that these effects correspond to important changes in thalamocortical coding properties.

### **III-64. Computational model of optimal changes in auditory cortical receptive fields during behavioral tasks**

Michael Carlin  
Mounya Elhilali

MACARLIN@JHU.EDU  
MOUNYA@JHU.EDU

Johns Hopkins University

Auditory scenes are dynamic and complex, requiring that neural resources be quickly repurposed to account for ever-changing behavioral and contextual demands. In particular, physiological studies have found that neural plasticity in mammalian primary auditory cortex can be manifested as a rapid adaptation in the shapes of spectro-temporal receptive fields (STRFs) to improve contrast between foreground and background sounds. Such top-down, attention-driven changes permit human and animal listeners to process salient sounds and ignore irrelevant or noisy sounds. As a complement to physiological studies, computational models of these adaptive changes are valuable since a better understanding of the principles underlying observed physiological patterns can aid analysis of neural data and help guide future experiments. In this study, we consider a novel computational framework to explain task-driven adaptive changes in a population of neurons characterized by their STRFs. Given a set of labeled acoustic tokens, our framework uses a discriminative probabilistic model to induce optimal changes in the shapes of STRFs that improve discrimination between acoustic categories. The framework optimizes an objective function that balances neural plasticity and discrimination, alternately updating the STRFs and the parameters of the probabilistic model. We validate our framework using an ensemble of STRFs obtained from ferret primary auditory cortex, demonstrating induced adaptive changes akin to those seen for a number of behavioral tasks described in physiological studies. Importantly, the induced changes broadly support the contrast filtering hypothesis, are spectro-temporally local, and reflect the meaning of the tokens (task valence). Furthermore, we note that because the framework applies to linear neural models, it is readily applicable to other sensory modalities. Finally, we propose a number of extensions to our model that may have practical implications for automated sound processing systems.

### **III-65. Graded memories in balanced attractor networks**

Dylan Festa  
Guillaume Hennequin  
Mate Lengyel

DF325@CAM.AC.UK  
GJE.HENNEQUIN@GMAIL.COM  
M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

Cortical circuits have separate excitatory and inhibitory (E-I) populations, operate in the balanced regime, and are able to store and retrieve memories that consist of graded patterns of activity. Previous theoretical studies have fallen short of explaining all these characteristics in a single unified model. Some models only allowed symmetric weight matrices precluding an E-I architecture. Others exploited the saturating part of nonlinear single

neuron transfer functions, leading to effectively binary memories and no balance. Yet others remained in the convex, non-saturated regime but stabilized low firing rates via a single inhibitory feedback loop, irrespective of the memory being recalled — but this inhibition could only be calibrated for a single stereotypical level of excitation, thus implying effectively binary memories again. Here we propose a novel framework in which we enforce a balanced regime and the stability of graded attractors by directly optimizing the connectivity of the network under physiological constraints. We formalized memory storage as implying two conditions: that the memorized patterns be fixed points of the dynamics, and that the dynamics be stable around those fixed points. The first objective is easily approached as it maps to a simple form of linear regression, despite the network being nonlinear, where the regressing coefficients are the elements of the synaptic weight matrix and the dependent variable is the required attractor. However, the second objective is a paradigmatic challenge in robust control studying the stabilization of systems with potentially strong positive feedback loops. Thus, we used a recently developed relaxation of the spectral abscissa which provides a smoothly differentiable measure of the stability of an attractor to perturbations. This allowed us to construct functioning E-I circuits that encode graded memories in the balanced regime. This provides the first step towards understanding the basic organizing principles of cortical memories.

### III-66. Interplay between learning-rate control and uncertainty minimization during one-shot causal learning

Sang Wan Lee  
John O'Doherty  
Shinsuke Shimojo

California Institute of Technology

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

While there is a rich computational, behavioral and neuroscience literature on how animals learn to associate stimuli and outcomes in an incremental fashion, much less studied in humans is how an individual can learn about causal relationships from only a single exposure to a cue-outcome pairing, aka 'one-shot-learning'. To address this, we developed a behavioral task in which novel or non-novel stimuli are presented in a sequence, followed by either a novel or non-novel monetary outcome. Following a block of stimulus and outcome pairings, participants are invited to provide behavioral ratings about the extent to which a given stimulus caused a particular outcome. We hypothesized that one-shot learning occurs when an agent is maximally uncertain about the causal relationship between a stimulus and an outcome (causal uncertainty). We further hypothesized that this is mediated through a rapid increase in the learning-rate for a particular cue-outcome association. To test this hypothesis we developed a novel Bayesian inference model and fit the behavioral data from 47 participants to this and a variety of alternative models. Our computational model provided a qualitatively good fit to the causal ratings generated by the participants on the task, while the alternative models failed to do so. This finding was confirmed by formal model comparisons on the behavioral data. Our results suggest that modulation of the learning rate serves to resolve uncertainty in causal relationships. Moreover, computer simulations indicate that the change in learning rate causes the reduction in the total amount of uncertainty. These findings provide a computational account for the phenomenon of one-shot learning above and beyond that previously established for incremental learning situations. The novel approach to studying one-shot learning developed here may also provide an experimental platform from which to study the neural mechanisms underpinning one-shot learning in the human brain.

### III-67. A self-consistent theory for STDP and spiking correlations in neuronal networks.

Gabriel K Ocker<sup>1</sup>  
Ashok Litwin-Kumar<sup>2</sup>  
Brent Doiron<sup>1</sup>

GKO1@PITT.EDU  
ALK@CMU.EDU  
BRENT.DOIRON@GMAIL.COM

<sup>1</sup>University of Pittsburgh

<sup>2</sup>Carnegie Mellon University

The structure of cortical networks is not random; certain patterns of connectivity are over-represented compared to a random network (Perin et al., 2011; Song et al., 2005). This connectivity patterning can be related to stimulus preference (Ko et al., 2011) and activity level (Yassin et al., 2010). How this structure arises is an important question for understanding the way cortical networks develop, learn and process information. The strength of individual synapses is shaped by both the firing rate and spike timing of pre- and post-synaptic activity through spike timing-dependent plasticity (STDP). Conversely, direct and indirect connections between neurons shape the joint statistics of pre- and post-synaptic activity. We have self-consistently related the evolution of spiking activity and network structure. Our theory combines a linear response theory for spiking correlations that takes the full network structure into account (Trousdale et al., 2012) and an adiabatic theory for the evolution of synaptic weights (Kempster et al., 1999). In contrast to previous theories of synaptic plasticity, we calculate the joint statistics of pre- and post-synaptic spiking from the specific integrate-and-fire model used and allow the spiking correlations to evolve with the synaptic weights, rather than fixing them. This allows the shape of the spiking cross-correlation functions to interact with the shape of the STDP window, giving rise to new dynamics of the synaptic weights. We investigate how STDP interacts with spiking activity to shape the structure and activity of neuronal networks. Using our self-consistent theory, we investigate how the form of different STDP rules and the intrinsic dynamics of neurons affect the generation of structured connectivity in neuronal networks.

### III-68. Generalization of memories by spatial patterning of protein synthesis

Cian O'Donnell<sup>1</sup>  
Terrence Sejnowski<sup>1,2</sup>

CIAN@SALK.EDU  
TERRY@SALK.EDU

<sup>1</sup>Salk Institute for Biological Studies

<sup>2</sup>Howard Hughes Medical Institute

Despite the decades-old knowledge that the consolidation of both memories and synaptic plasticity requires the synthesis of new proteins, the computational benefit of this cellular process for learning remains unknown. Interestingly, protein expression during synaptic plasticity has been found to be spatially restricted to the activated neurons, and even to specific activated dendrites within those neurons. Using computational modeling we found that this spatial patterning of protein synthesis at both the single-cell and neural circuit levels can allow for selective memory consolidation. In this way the brain can simultaneously remember some items while forgetting others, based on each item's information content. We showed how the efficacy of this mechanism depends on the specificity of synaptic wiring on dendrites, and on the overlap of neural activity patterns at the circuit level. We used the framework to predict which specific brain circuits are the best places to look for this mechanism. Finally, we further used the framework to propose a novel model for memory generalization during slow-wave and REM sleep.



### III-69. A normative theory of structural and intrinsic plasticity in dendrites

Balazs Ujfalussy  
Mate Lengyel  
University of Cambridge

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

The dendritic tree of a cortical neuron is the locus of a dazzling array of powerful and puzzling forms of plasticity. Beside more traditional forms of synaptic plasticity, dendrites also exhibit experience-dependent forms of structural as well as intrinsic plasticity, changing the set of presynaptic neurons targeting a dendritic branch and its intrinsic nonlinearities, respectively. While synaptic plasticity has been studied extensively, little is known about the functional principles underlying the other two forms of plasticity. Here we derive optimal rules for structural and intrinsic plasticity, and their biological approximations, based on the idea that they tune sub-cellular connectivity and dendritic nonlinearities to the statistics of population activity in the presynaptic ensemble to allow for robust computations in spiking circuits. Our theory predicts that presynaptic neurons from the same ensemble innervate the same dendritic branch, and hence the stabilization or elimination of a newly formed synapse depends on the evidence that it belongs to the same ensemble as the neighboring synapses. We show that computing this evidence requires two signals readily available at the synapse: one that is local to the synapse (glutamate bound to the NMDA receptors) and another that is local to the dendritic branch (its membrane potential). Both the local nature and the timing requirements of the derived plasticity rule match those observed in dendrites of cortical neurons (Gordon et al., 2006). Our theory also predicts that the nonlinearity of a dendritic branch depends on the correlations among the neurons of the presynaptic ensemble. Thus, we derive a plasticity rule to learn these correlations, and the resulting nonlinearity, from spiking inputs and show a close correspondence with experimentally observed forms of branch strength potentiation (Losonczy et al., 2008). Our approach provides a novel framework for studying plasticity within the dendritic tree from the perspective of circuit-level computations.

### III-70. Contribution of cerebellar Golgi cells to learned motor timing during the vestibulo-ocular reflex

Thomas Chartrand<sup>1,2</sup>  
Grace Q. Zhao<sup>3</sup>  
Jennifer L. Raymond<sup>3</sup>  
Mark Goldman<sup>1</sup>

TMCHARTRAND@UCDAVIS.EDU  
GQZHAO@STANFORD.EDU  
JENR@STANFORD.EDU  
MSGOLDMAN@UCDAVIS.EDU

<sup>1</sup>University of California, Davis  
<sup>2</sup>University of Cambridge  
<sup>3</sup>Stanford University

The cerebellum is known to be essential for the learning of precisely timed motor responses. Its stereotyped organization suggests that a common mechanism may underlie a wide range of cerebellar motor functions. However, a mechanistic understanding of cerebellar function has been hindered by lack of knowledge about the Golgi and granule cell network that forms the input layer of the cerebellar cortex. This network is widely hypothesized to encode signals for cerebellar motor learning, but direct experimental manipulations or recordings of the interactions occurring within this circuit during behavior have proven difficult. To bring new insights to this problem, we used a genetic knockout mouse line with targeted disruption of inhibitory Golgi-to-granule cell synapses, and combined behavioral testing with computational modeling of cerebellar motor learning. As a precisely quantifiable, well-studied paradigm for cerebellar motor learning, we tested adaptation of the vestibulo-ocular reflex (VOR), an eye movement reflex that stabilizes gaze in response to head motion. We observed deficits in the knockout mice for learned changes in the timing, but not gain, of VOR responses, directly confirming that these inhibitory connections are essential for cerebellar learned timing. We modeled this perturbation by constructing a network model of Golgi and granule cells, with relative numbers and connectivity of cell classes based on known anatomy, and with time constants based on previous studies of cerebellar electrophysiology. To model the learned motor output, we simulated a correlational learning rule for granule cell to Purkinje cell synapses, guided by a climbing fiber

error signal. We find conditions under which the Golgi cell inhibitory connections create the temporal diversity of granule cell population activity required for generating time- or phase-shifted cerebellar output. These results demonstrate how Golgi cell inhibition, through effects on granular layer network dynamics, contribute essentially to cerebellar-mediated learning of timed motor responses.

### **III-71. Stochastic synchronization in the Purkinje cells of the cerebellum**

Sergio Verduzco-Flores

SERGIO.VERDUZCOFLORES@COLORADO.EDU

University of Colorado at Boulder

An outstanding question in cerebellar research is how the plasticity modulated by climbing fiber activity affects the output of the cerebellum. The only output of the cerebellar cortex consists of the Purkinje cell axons, which make inhibitory synapses onto neurons in the deep cerebellar nuclei (DCN). A large enough hyperpolarization of DCN cells by Purkinje inhibition can activate cationic currents that generate 'rebound potentials.' Because of this, synchronization of all the Purkinje inputs that converge onto a single DCN cell can increase its firing rate, since a barrage of synchronous inputs causes a deep hyperpolarization, and between each barrage there is a time window during which the DCN neuron can fire. Consistent with this, simple spikes in nearby Purkinje cells tend to synchronize. This synchronization has been previously attributed to recurrent axon collaterals in Purkinje cells, despite the fact that these collaterals form weak and sparse connections. Another possible synchronization mechanism, known as stochastic synchronization has not been explored before. Stochastic synchronization happens when uncoupled oscillators synchronize by receiving correlated inputs, and among Purkinje cells this could be functionally more important than the coupling by recurrent collaterals, because several known plasticity mechanisms in the cerebellar cortex seem to act so as to increase or reduce stochastic synchrony when appropriate. These plasticity mechanisms are modulated by climbing fiber activity, which is considered to be the main learning signal in the cerebellum. This work explores how synaptic plasticity in the cerebellar cortex could affect three factors controlling stochastic synchrony: input correlation, homogeneity of responses, and input amplitude. We use a novel method in order to characterize the response of Purkinje cells through phase response curves despite the presence of feed-forward inhibition. We then use a computational model to explore the effect of synaptic plasticity by simulating variations in the three aforementioned factors.

### **III-72. Two forms of information found in local field potentials of the hippocampus**

Gautam Agarwal<sup>1,2</sup>

Kenji Mizuseki<sup>3</sup>

Antal Berenyi<sup>4</sup>

Gyorgy Buzsaki<sup>5</sup>

Friedrich Sommer<sup>2</sup>

GAGARWAL@BERKELEY.EDU

KENJI.MIZUSEKI@GMAIL.COM

DRBERENYI@GMAIL.COM

GYORGY.BUZSAKI@NYUMC.ORG

FSSOMMER@BERKELEY.EDU

<sup>1</sup>Redwood Center for Theoretical Neuroscience

<sup>2</sup>University of California, Berkeley

<sup>3</sup>Allen Institute for Brain Science

<sup>4</sup>University of Szeged

<sup>5</sup>New York University Medical Center

The local field potential (LFP) reflects the summed synaptic and spiking activity of neuronal populations. Different frequency bands of the LFP arise from distinct physiological sources; however, it remains unclear to what extent the LFP preserves information present in the underlying neuronal population. We analyze multi-electrode LFP recordings from the hippocampus of rats, in order to decipher the content of signals found at different timescales. We first examine how well each LFP frequency band predicts the spiking of simultaneously recorded neurons. We

filter the LFP at several frequencies (5 - 500 Hz), and at each frequency use linear regression to predict single-neuron activity. We find that predictions are best at high frequencies (~100+ Hz), and in the theta band (~8 Hz). However, the fits in the two frequency regimes differ fundamentally in the structure of their weight matrices, and in their ability to predict spiking in new behavioral contexts. Using an unsupervised learning algorithm known as ICA, we further examine the structure of activity at these two frequencies (5-11 Hz and 100-625 Hz). In the theta band, ICA learns features that span the entire electrode array, and activate at single positions of the rat's environment. At high frequencies, learned features localize to small groups of electrodes, and activate at multiple positions of the rat's environment, in a manner that corresponds to the spiking of nearby neurons. We find that information about the animal's position originates from two distinct sources. The first is the spiking activity of nearby neurons. The second is the synchronized oscillatory activity of a much larger, anatomically dispersed, population of neurons that co-activate at specific locations of the rat's environment. In each case, the multi-electrode patterns of activity predict the rat's location. Thus the LFP offers a rich substrate for dissecting multiple scales of a network's activity.

### III-73. Cluster analysis of sharp-wave ripple field potential signatures in the macaque hippocampus

Juan Ramirez-Villegas<sup>1</sup>

Nikos Logothetis<sup>1</sup>

Michel Besserve<sup>2,1</sup>

JUAN.RAMIREZ-VILLEGAS@TUEBINGEN.MPG.DE

NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE

MICHEL.BESSERVE@TUEBINGEN.MPG.DE

<sup>1</sup>Max Planck Institute for Biological Cybernetics

<sup>2</sup>Max Planck Institute for Intelligent Systems

Sharp-wave ripple complexes (SPW-Rs), transient episodes of neural activity combining a sharp wave of dendritic depolarization and a high-frequency oscillation, are a major feature of the cortico-hippocampal communication during immobility, consummatory behaviors and sleep. Experimental evidence relates these episodes to offline consolidation of memory traces. In order to allow for the wide range of network reconfigurations required by this process, different SPW-R events certainly reflect a large variety of selective interactions both within the hippocampus and with other brain regions. A better understanding of the underlying mechanisms of these interactions thus requires a finer characterization of the SPW-R events and their associated signatures over the entire brain. Using unsupervised-learning artificial neural networks and clustering techniques, we analyzed peri-event multi-channel local field potential (LFP) recordings of the hippocampus of anesthetized macaques, and extracted the electrophysiological characteristics of SPW-R events dynamics. We combined this analysis with neural event-triggered functional magnetic resonance imaging analysis in order to map the activity of the whole brain during SPW-R events. Our primary findings hint upon differentiated SPW-R complexes, whose signatures come in four classes: high frequency oscillations preceding, following or located in the peak of sharp wave dendritic depolarization, as well as high frequency oscillations without noticeable sharp-wave signature. These differentiated SPW-R LFP signatures were highly reproducible both across different animals and experimental sessions. At a larger scale, ripple-triggered fMRI activation map for most of the cerebral cortex and sub-cortical regions during ripples has been described by Logothetis et al., 2012. On top of this, our preliminary results suggest that the classes of SPW-R field potential signatures reflect differentiated cortical activation and sub-cortical deactivation maps. In light of these findings, we hypothesize that these distinct patterns of SPW-R in hippocampal LFP mark differentiated brain-wide dynamical events, possibly reflecting several underlying mechanisms of memory function.

### III-74. Priors formed over multiple time scales facilitate performance in a visual search task

Nikos Gekas<sup>1</sup>  
 Aaron Seitz<sup>2</sup>  
 Peggy Series<sup>1</sup>

N.GEKAS@SMS.ED.AC.UK  
 ASEITZ@UCR.EDU  
 PSERIES@INF.ED.AC.UK

<sup>1</sup>University of Edinburgh

<sup>2</sup>University of California, Riverside

Bayesian Inference provides a parsimonious model of perception, where sensory-stimuli are combined with prior beliefs in a near-optimal way [1]. In this framework, a question of key importance is how the visual system updates its priors through interaction with the environment. Here, we used a visual search task to explore how expectations of different time scales (from the last few trials to hours to long-term statistics of natural scenes) interact to alter perception. Previously [2, 3], we demonstrated that the statistics of past visual motion stimuli can powerfully modulate the perception of new motion directions. Here, we presented human observers with low contrast white dots in twelve possible positions equally spaced on a circle, and we asked them to identify the presence and position of the dots, and manipulated subjects' expectations by presenting stimuli at some positions more frequently than others. We compared the effect of expectations developed over a long sequence of trials, with priming, which we defined as positions being presented in the very few last trials (e.g. 3). Our findings suggest that there are strong pre-existing long-term biases regarding target positions (e.g. horizontal vs. vertical). On top of these biases, subjects quickly learned about the stimulus distribution, which improved their detection performance, and, in the absence of stimuli, perceived stimuli more often at the most frequently presented positions. Moreover, subjects significantly improved their detection performance and reported more false positives at the most frequently presented positions when these were primed compared to other positions. Our results can be modeled and understood within the Bayesian framework, in terms of a near-optimal integration of pre-existing long-term priors with rapidly learned statistical priors, which are skewed towards the very recent history of trials, and may help in understanding the time-scale of developing expectations at the neural level.

### III-75. A multi-step decision task elicits planning behavior in rats

Kevin Miller<sup>1</sup>  
 Jeffrey Erlich<sup>1</sup>  
 Charles Kopec<sup>1</sup>  
 Matthew Botvinick<sup>1</sup>  
 Carlos Brody<sup>1,2</sup>

KJMILLER@PRINCETON.EDU  
 JERLICH@PRINCETON.EDU  
 CKOPEC@PRINCETON.EDU  
 MATTHEWB@PRINCETON.EDU  
 BRODY@PRINCETON.EDU

<sup>1</sup>Princeton University

<sup>2</sup>Howard Hughes Medical Institute

A central goal of cognitive neuroscience is to understand the neural mechanisms of decision making. In recent years, much progress has been made uncovering these mechanisms, but this research has focused largely on immediate decisions. The neural mechanisms of planning (that is: use of an internal model of the structure of an environment to inform sequential decision-making) remain effectively unknown. One major reason for this is that currently available planning tasks for animals typically produce only one or a few instances of planned behavior per subject, making it impossible to apply many of the most powerful methods of modern systems neuroscience. We have adapted for rats a two-step decision-making task from the human literature (Daw et al, 2011), and show that it produces planning behavior in most animals. In the first step of this task, rats choose between two alternatives, each of which leads to one of two second-step actions becoming available with high probability (80%), and to the other with low probability (20%). In the second step of the task, the rat performs whichever action is available, and either receives or does not receive a water reward. Reward probabilities associated with the second-step actions are constantly changing, requiring rats to continually learn in order to perform well. Optimal performance requires a model-based strategy: animals need to know the transition probabilities from first- to second-step

actions in order to choose the first-step action that leads most often to the second-step action with highest reward probability. The majority of rats that learn the task adopt such a strategy, making this (to our knowledge) the first repeated-trials planning task for rats. This opens the door to many exciting experimental possibilities — immediate next steps involve inactivations and single-unit recordings in prelimbic cortex and other brain regions.

### III-76. Dissociable systems for reward and content updating revealed using cross stimulus suppression

Erie Boorman<sup>1,2</sup>  
Vani Rajendran<sup>1</sup>  
Timothy Behrens<sup>1</sup>

ERIE.BOORMAN@NDCN.OX.AC.UK  
VANI.RAJENDRAN@UNIV.OX.AC.UK  
BEHRENS@FMRIB.OX.AC.UK

<sup>1</sup>University of Oxford

<sup>2</sup>University College London

A crucial question in reinforcement learning (RL) concerns how animals learn a 'forward model' of the world that enables them to efficiently pursue rewarding goals. Forming this model necessitates learning what to expect from a choice. We present a novel approach that combines computational fMRI with repetition-suppression to investigate this question. This approach, cross-stimulus-suppression (XSS), exploits a generalization of repetition-suppression whereby sequential presentation of different stimuli induces attenuation in BOLD response if they are related by association. We designed a task in which learning a stimulus's reward value could be dissociated from learning the likely reward identity that followed if selected. When subjects received feedback, we identified two simultaneous prediction errors (PEs)—one on the value of reward (an rPE) and one on the association between a stimulus and reward identity (an iPE)—in putative ventral tegmental area (VTA) and lateral orbitofrontal cortex (IOFC), respectively. In between successive PEs, we used XSS to probe both the strength of association between specific stimuli and reward identities and the change in association from one choice to the next as a result of the intervening PEs. We found that IOFC, medial prefrontal cortex, anterior caudate, and hippocampus showed BOLD suppression proportional to the model-derived, trial-by-trial association between a stimulus and outcome. This shows that the degree of suppression in these regions reflects trial-by-trial associations between stimuli and reward identity. Furthermore, the change in suppression from one trial to the next in IOFC was proportional to the iPE, suggesting IOFC both encodes and updates associations crucial for a forward model of the reward environment. Collectively, these findings show that updating the value of stimuli in VTA is anatomically separable from updating the transition between stimuli and reward identities in IOFC, and suggest a new approach to investigate similar learning questions in humans.

### III-77. Working memory contributions to reinforcement learning impairments in schizophrenia.

Anne Collins<sup>1</sup>  
James Gold<sup>2</sup>  
James Waltz<sup>2</sup>  
Michael Frank<sup>1</sup>

ANNE\_COLLINS@BROWN.EDU  
JGOLD@MPRC.UMARYLAND.EDU  
JWALTZ@MPRC.UMARYLAND.EDU  
MICHAEL\_FRANK@BROWN.EDU

<sup>1</sup>Brown University

<sup>2</sup>University of Maryland

Previous research has shown that patients with schizophrenia show impairment on reward- and punishment-driven learning tasks. However, the observed learning behavior originates from the interaction of multiple neural processes, such as incremental basal-ganglia and dopamine-dependent reinforcement learning (RL) systems, but also prefrontal-cortical-dependent working memory (WM). Thus, the degrees to which specific systems contribute to learning impairments in schizophrenia are unclear. We recently developed a task and computational model

allowing us to separately assess the roles of (slow, cumulative learning) RL mechanisms vs. (fast, but capacity limited) WM mechanisms, in healthy adults. Here, we used this task with 48 patients and 37 healthy controls, in order to assess the specific sources of learning impairments. In 15 separate blocks, subjects learned to pick 1 of 3 actions for stimuli. The number of stimuli to learn in each block varied from 2 to 6, allowing us to separate influences of capacity limited WM from RL system. As expected, both patients and controls were able to learn in all blocks, but showed effects of set size and delay between identical stimulus repetition, confirming the presence of working memory effects. Patients performed significantly worse than controls. Fit of our working-memory reinforcement-learning model accounted better for the data than all other models. Model fits showed a significant difference between patients and controls for working-memory parameters (capacity and reliability), but not for reinforcement learning parameters (learning rate, softmax temperature). These results indicate that dysfunctional prefrontal working memory mechanisms make an inordinate contribution to learning impairment in schizophrenia.

## III-78. Temporal stimulus segmentation and decision making via spike-based reinforcement learning

Giancarlo La Camera<sup>1</sup>  
Robert Urbanczik<sup>2</sup>  
Walter Senn<sup>2</sup>

LACAMERAG@GMAIL.COM  
URBANCZIK@PYL.UNIBE.CH  
SENN@PYL.UNIBE.CH

<sup>1</sup>State University of New York at Stony Brook

<sup>2</sup>University of Bern

In theoretical and experimental investigations of decision-making, a decision maker (or 'agent' in the following) learns to make the correct decision in the presence of a stimulus cue. For every decision, the agent receives feedback about whether or not the decision was correct, and upon many repeated presentations of each stimulus, the agent learns to respond appropriately to all of them. This problem is often framed as a classification task wherein the relevant stimuli cueing the correct decisions are known to the agent: the agent knows which stimuli are action-relevant and when it is being presented with one. However, in many real-life situations it is not clear which segments in a continuous sensory stream are action relevant, and relevant segments may blend seamlessly into irrelevant ones. Then the decision problem is just as much about when to act as about choosing the right action. Here, we present a spiking neuron model of a cortical circuit (also referred to as the 'agent') which learns to classify hidden relevant segments in a continuous sensory stream of spatio-temporal patterns of spike trains. The agent has no a-priori knowledge of the stimuli, when they are being presented, and their behavioral significance — i.e., whether or not they are action-relevant. The agent is trained by the reward received for taking correct decisions in the presence of relevant stimuli. Simulation results show that by maximizing expected reward, the agent learns to distinguish behaviorally relevant segments in the input stream from irrelevant ones, performing a task akin to temporal stimulus segmentation.

## III-79. A link between corticostriatal plasticity and risk taking in humans

David Arkadir<sup>1</sup>  
Angela Radulescu<sup>2</sup>  
Deborah Raymond<sup>3</sup>  
Susan B Bressman<sup>3</sup>  
Pietro Mazzoni<sup>4</sup>  
Yael Niv<sup>2</sup>

ARKADIR@GMAIL.COM  
ANGELAR@PRINCETON.EDU  
DRAYMOND@CHPNET.ORG  
SBRESSMA@CHPNET.ORG  
PM125@COLUMBIA.EDU  
YAEL@PRINCETON.EDU

<sup>1</sup>Columbia Neurological Institute

<sup>2</sup>Princeton University

<sup>3</sup>Beth Israel Medical Center

<sup>4</sup>Columbia University

A central assumption of reinforcement learning theory is that dopaminergic prediction errors drive trial and error learning by modulating long-term potentiation (LTP) and long-term depression (LTD) in corticostriatal synapses. To test the link between corticostriatal plasticity and overt behavior, we examined learning in patients with DYT1 dystonia, a genetic movement disorder that has been linked to altered plasticity. Rodent models of DYT1 dystonia suggest a fundamental imbalance between LTP and LTD in corticostriatal synapses: in brain slices of transgenic rodents expressing the mutant TOR1A gene that causes dystonia, LTP was greater in magnitude, while LTD could not be induced. If this mechanism also manifests in humans, patients with DYT1 dystonia should show abnormal trial and error learning due to an imbalance in learning from positive and negative prediction errors. To test this hypothesis, we compared the behavior of DYT1 dystonia patients and healthy controls on an instrumental learning task previously shown to depend on prediction-error driven learning in the striatum. Participants learned to associate four different visual cues with monetary rewards. Three of the cues yielded deterministic rewards of 0', 5', and 10' respectively, while the payoff for the fourth 'risky' cue was 0' or 10' with equal probability. In accordance with the hypothesized role of corticostriatal plasticity in reward learning, DYT1 dystonia patients were more risk taking than healthy controls, exhibiting a higher probability of choosing the risky 0/10' cue over the sure 5' cue. Moreover, the degree of risk taking increased with the severity of motor symptoms. These findings help to establish the long-postulated link between corticostriatal plasticity and trial and error learning and provide new insights into DYT1 dystonia.

### III-80. Kalman filter modeling of reinforced behavior: Application to pitch and tempo learning in songbirds

Anja T. Zai<sup>1,2</sup>  
 Florian Blattler<sup>1,2</sup>  
 Anna E. Stepien<sup>1,2</sup>  
 Alessandro Canopoli<sup>1,2</sup>  
 Richard H.R. Hahnloser<sup>1,2</sup>

ZAI@INI.PHYS.ETHZ.CH  
 FLORIAN@INI.PHYS.ETHZ.CH  
 STEPIEN@INI.PHYS.ETHZ.CH  
 CANOPOLI@INI.PHYS.ETHZ.CH  
 RICH@INI.PHYS.ETHZ.CH

<sup>1</sup>ETH Zurich  
<sup>2</sup>University of Zurich

A trial-and-error process of reinforcement learning has been shown to effectively guide adaptive modifications of well-controlled and highly complex motor skills. Songbirds provide a unique model system to study the neural systems underlying such processes. Both spectral and temporal (Ali et al., 2013) aspects of birdsong can be modified independently by delivering real-time auditory feedback to a subset of produced song syllables: Birds tend to adapt their songs to escape those disruptions, demonstrating that birds can exploit trial-by-trial variation in vocalizations to escape negative reinforcement. While this widely used learning paradigm has brought many new insights, no mechanistic models of such trial-and-error reinforcement learning exist to date that would allow accurate fitting of behavioral trajectories with the goal of relating neurophysiological findings to relevant behavioral parameters such as the different sources of motor variability and the learning rate. To overcome this gap, we present a linear dynamical system to model the effect of a reinforced trial on future trials via the correlation between exploration and reward, as in the Williams model (Williams, 1992). In our system, the reward acts as a time-dependent coefficient; the learning rate and other model parameters can be estimated using the expectation maximization algorithm. The model includes one or several hidden states that represent the global outputs of different premotor brain areas and that account for reinforcement-driven behavioral changes. The mechanistic model allows us to attribute birds' behaviors to just a few intrinsic parameters and to make predictions about the underlying motor codes. The proposed model generalizes the successful Timing Variability Model (Glaze & Troyer, 2012) used to describe the temporal structure of (non-reinforced) song in two ways: first, it does not assume trial independence; and, second, it captures the influences of both reward and sensory targets on behavioral trajectories.

### III-81. A role for habenula in dopamine neurons' responses to reward omission

Ju Tian  
Naoshige Uchida  
Harvard University

JUTIAN@FAS.HARVARD.EDU  
UCHIDA@MCB.HARVARD.EDU

Learning the values of different objects is critical for survival. Midbrain dopamine neurons encode reward prediction error (RPE), the difference between actual outcome and expected outcome, a signal thought to drive associative learning. Although various computational models have been proposed to explain how dopamine neurons compute RPEs, exact mechanisms remain elusive. Recent studies have found that neurons in the lateral habenula (LH) encode negative prediction errors, the mirror image of dopamine activity. LH sends indirect inhibitory projections to dopamine neurons. Based on these observations, one prevalent hypothesis has been that RPE is already calculated in the LH and dopamine neurons merely relay this signal to other parts of the brain. To test whether the habenula is required for RPE signaling by dopamine neurons, we electrolytically lesioned the habenula bilaterally nearly completely. We then recorded the activity of optogenetically-identified dopamine neurons while mice performed a classical conditioning task. Compared to dopamine neurons in control animals, dopamine neurons in lesioned animals showed greatly diminished suppression of activity during the omission of predicted reward. In contrast, air puffs caused similar levels of inhibition in control and lesioned animals. Dopamine neurons' phasic responses to reward-predictive cues and reward itself were slightly shortened in duration, but were largely intact after lesions. Furthermore, dopamine neurons in lesioned animals showed normal suppression of reward responses by expectation. These results indicate that the habenula plays a predominant role in sending reward omission signals to dopamine neurons. Interestingly, different mechanisms appear to underlie dopamine neurons' responses to two types of negative RPEs: habenula inputs are input for reward omission but not aversive stimuli. Although the habenula plays a relatively minor role in dopamine neurons' positive RPE signaling, it appears to boost the magnitude of these responses. Together, these results indicate that dopamine neurons combine multiple mechanisms to calculate RPE.

### III-82. The interaction between learning and mood

Eran Eldar  
Yael Niv  
Princeton University

ELDAR.ERAN@GMAIL.COM  
YAEL@PRINCETON.EDU

Fluctuations of good and bad mood characterize the emotional life of many. Here, we show that fluctuations of mood may emerge from simple reinforcement learning if we assume that surprising outcomes affect mood and mood biases perception of outcomes (i.e., the same outcome seems better when in a good mood and worse when in a bad mood). Accordingly, we show empirically that a relatively large surprising outcome affects both mood and the perception of subsequent outcomes, specifically in individuals that are susceptible to mood fluctuations. Biased reward perception manifested in both choice behavior and neural response to reward, as measured using functional MRI: In participants who were characterized as susceptible to mood fluctuations (using a hypomanic personality scale), choice behavior and striatal activity reflected higher valuation of rewards received after a large mood-affecting gain, as compared to rewards received after a large mood-affecting loss. Individuals who were characterized as less susceptible to mood fluctuations were not biased by mood in either choice behavior or striatal activity. We present a reinforcement-learning model that captures the proposed interaction between learning and mood. The model outperforms standard reinforcement learning in explaining trial-to-trial choices and striatal activity of individuals whose behavior reflected a mood-related bias. Furthermore, the model correctly predicts participants' feelings (as assessed through self reports), and the mood bias that it infers from behavior correlates with individual susceptibility to mood fluctuations. In sum, our findings suggest that learning affects and is affected by mood, and that the dynamics that result from this interaction may contribute to mood instability. The findings also uncover a decision bias that can inform behavioral economics models of sequential decisions. Our theory



may inspire novel approaches to the study and treatment of psychiatric disorders such as cyclothymia and bipolar disorder.

### III-83. Striatal neuronal activity of macaque tracks local and global changes during natural habit formation

Theresa Desrochers<sup>1</sup>  
Ken-ichi Amemori<sup>2</sup>  
Ann Graybiel<sup>2</sup>

THERESA\_DESROCHERS@BROWN.EDU  
AMEMORI@MIT.EDU  
GRAYBIEL@MIT.EDU

<sup>1</sup>Brown University

<sup>2</sup>Massachusetts Institute of Technology

Habits are ubiquitous in our daily lives. In a prior study, we quantified the free scanning eye movements of naive monkeys as they acquired their own habitual patterns. Reinforcement learning (RL) algorithms have successfully provided sufficient explanatory drive for such behaviors. By introducing cost (distance the eye traveled) as the variable to be minimized in the RL framework, in behavioral work, we showed that the local and trial-by-trial changes in cost could explain the likelihood of expression of scan patterns—repeating of a pattern was predicted by small differences in the amount of work required on prior trials in which those patterns were observed. The neural correlates of such behaviors remain a mystery. We recorded neuronal activity the caudate nucleus (CN) of two macaque monkeys throughout the acquisition and performance of a free scanning eye movement task, and found that the CN activity was correlated with two learning variables operating at different time scales. First, the activity of a substantial number of CN neurons was correlated with the local changes in cost, the same measure that drove the gradual evolution of the eye movement patterns using an RL model. Second, neuronal populations in the CN gradually changed their firing patterns across months, such that their response to the end of task event (targets off) sharpened as the monkeys acquired the habitual patterns. This sharpening of responses correlated with gradual behavioral changes (increases in efficiency/decreases in entropy) that were previously found to occur continuously throughout learning. Thus, we show that the striatum is tracking both the local fluctuations in cost and the global changes in repetitiveness across learning as naive monkeys performed a naturalistic, free-viewing eye movement task. This provides a novel mechanism by which RL and its neural correlates could drive naturalistic behaviors in the freely behaving animal.

### III-84. A signal transduction model for adenosine-induced indirect corticostriatal plasticity

Naoto Yukinawa  
Kenji Doya  
Junichiro Yoshimoto

NAOTO.YUKINAWA@OIST.JP  
DOYA@OIST.JP  
JUN-Y@OIST.JP

OIST Graduate University

Adenosine is a ubiquitous signaling molecule, whose extracellular level is closely linked to metabolic energy consumption at various biological scales from cells to organs. In central nervous system, adenosine competitively inactivates D2-type dopamine receptor (D2R)-dependent signaling pathways of the striatopallidal medium spiny neurons (MSNs) expressing specifically D2Rs and A2A-type adenosine receptors (A2ARs). This antagonistic signaling property could lead to potential motor inhibition, when chronic depletion of dopamine occurs in the basal ganglia such as Parkinson's disease, via imbalanced activation of the indirect pathway over the direct pathway of the cortico-basal ganglia circuits. On the other hand, little is known about the functional role of adenosine on corticostriatal synaptic plasticity in the indirect pathway, i.e., modulation of motor control. Here we elucidate effects of the opposing biochemical interactions between adenosine and dopamine on the indirect corticostriatal learning by quantitative simulation of a detailed kinetic model for membrane trafficking of AMPA-

type glutamate receptors (AMPA), a major component of postsynaptic plasticity. Our model consists of several signaling cascades based on previous modeling studies (Nakano et al., 2010, Oliveira et al., 2012) incorporating neurophysiological evidence, including antagonistic regulation of protein kinase A mediated by two different G-protein coupling receptors of D2R and A2AR, as well as the downstream phosphoregulation of dopamine- and cAMP-regulated phosphoprotein and AMPAR. By simulating the model, we show that biologically observed levels of extracellular adenosine in the range of tens to a few hundreds nM yield broad dopamine dose-dependent phosphorylation responses of AMPAR demonstrating long-term potentiation at low dose and long-term depression at higher dose, which corresponds to exhausted states and reward-based activation of dopamine, respectively. Our results suggest that the basal level of adenosine is optimally tuned to produce the corticostriatal plasticity and may provide appropriate balance for the corticostriatal learning between the direct and indirect pathways.

### **III-85. In vivo two-photon imaging of cortico-cerebellar mossy fiber synapses: functional properties**

Daria Rylkova  
Ye Zengyou  
David Linden

DRYLKOVA@GMAIL.COM  
ZYE10@JHMI.EDU  
DLINDEN@JHMI.EDU

Johns Hopkins University

It is now appreciated that the cerebellum participates in both motor and cognitive functions, which are likely mediated by its vast interconnection with the neocortex via the pons. Using in vivo two-photon imaging in adult behaving rats, we are characterizing the morphological and functional dynamics of basal pontine mossy fibers, which form the cortico-cerebellar circuit's input stage. To this end, we are labeling mossy fibers with dsRed and GCaMP3 by injecting AAVs into the pons and performing chronic two-photon imaging through a cranial window over the cerebellar hemisphere. Dual labeling allows us to reconstruct multiple axons within a volume and subsequently characterize calcium signals in terminals with reference to their parent axon. During imaging, sensory-evoked activity is tested and animals' behavior is recorded using a high speed camera. We have observed robust activity associated with both movement and sensory stimuli. There are large overlaps between representations, with individual axons being active during movements of multiple body parts. Surprisingly, calcium signals often fail to be detected at all terminals on an 'active' axon, despite the terminal being active at other times during the imaging session. Terminals on the same axon are more highly correlated than those on different axons and there is an inverse relationship between distance in 3D space and correlated activity for all terminals. Interestingly, this does not appear to be a result of branch point failure because calcium transients can be detected in terminals distal to an inactive terminal. Moreover, there isn't a clear relationship between the number of branch points separating two terminals or the axonal distance between them and correlation, indicating a more complex interaction between presynaptic activity and the local environment. The co-activation of terminals is also dependent on behavioral state, with a larger fraction of terminals being co-active during sensorimotor experience.

### **III-86. Efficient coding of visual motion signals in the smooth pursuit system**

Leslie Osborne  
Bing Liu

OSBORNE@UCHICAGO.EDU  
LIUBING1108@GMAIL.COM

University of Chicago

Performance in sensorimotor behaviors guides our understanding of many of the key computational functions of the brain: the representation of sensory information, the translation of sensory signals to commands for movement, and the production of behavior. Eye movement behaviors have become a valuable testing ground for theories of neural computation because the circuitry has been well characterized and eye movements can be tightly coupled to cortical activity. Here we show that pursuit eye movements, and the cortical sensory signals that

mediate them, demonstrate the hallmarks of efficient coding. Barlow proposed that neurons should adapt their sensitivity as stimulus conditions change in order to maintain efficient representation of sensory inputs. Evidence for efficient coding of temporal fluctuations in visual contrast has been observed in the retina. We asked whether adaptation to stimulus variance generalizes to higher cortical areas whose neurons respond to features of visual signals that do not drive adaptation in the periphery and whether such adaptation impacts performance of visually-driven behavior. Specifically, we studied the impact of cortical adaptation to fluctuations in motion direction on pursuit. We recorded eye movements of monkeys pursuing moving targets with an added stochastic perturbation. We found that the amplitude of the linear filter relating eye to target movement rescaled in proportion to target motion variance, consistent with the efficient coding hypothesis. Steps in target motion variance created a transient decrease in the information capacity of pursuit. To test whether this behavioral adaptation arises in the visual system, we recorded unit responses in middle temporal cortical area (MT). We found that MT responses echo those of pursuit: neural filters – between firing rate and target motion scale with variance, the stimulus-response distribution rescales rapidly – within 30ms of a variance change – and some units display a rapid recovery of motion information after a step.

### III-87. Corresponding neural signatures of movement and imagined movement in human motor cortex

Chethan Pandarinath<sup>1</sup>

Vikash Gilja<sup>1</sup>

Christine Blabe<sup>1</sup>

Leigh Hochberg<sup>2</sup>

Krishna Shenoy<sup>1</sup>

Jaimie Henderson<sup>1</sup>

CHETHAN@STANFORD.EDU

GILJA@STANFORD.EDU

CBLABE@STANFORD.EDU

LHOCHBERG@PARTNERS.ORG

SHENOY@STANFORD.EDU

HENDERJ@STANFORD.EDU

<sup>1</sup>Stanford University

<sup>2</sup>Massachusetts General Hospital

Clinical studies have demonstrated that human primary motor cortex (M1) exhibits modulation during movement, as well as during imagined movement. In non-human primates (NHPs), population activity in M1 follows a well-described progression from movement planning to execution. Here we seek to understand whether population activity in human M1 follows this same progression. Further, we asked questions that are difficult to probe through NHP experiments: does the process of imagining movements follow a similar structure? How does population activity during imagined movements relate to real movements? We investigated these questions with a participant with ALS under the BrainGate2 clinical trial. The participant retains the ability to make several dexterous movements. We analyzed single channel and population spiking activity recorded from M1 as the participant performed finger flexion movements in an instructed-delay task. Single channels reliably modulated during the delay period (before the "go" cue), and then exhibited further modulation as the participant executed the movement. Using dimensionality reduction techniques, we analyzed trajectories in neural state space over the course of the task. Trajectories during planning moved from baseline to a preparatory state, and then progressed along a stereotyped trajectory during movement execution. During imagery blocks, the task followed the same structure, but the participant was asked to imagine the movements. Single channels exhibited modulation during planning, and further modulated after the "go" cue. We applied the previously learned dimensionality reduction transformation and found that neural activity during the delay period progressed to the same planning states as the movement paradigm. Trajectories following the "go" cue were qualitatively similar to the movement paradigm, but still distinguishable. These results demonstrate that the process of imagining movements follows the same structure as movement execution, and may lead to a better understanding of what aspects of M1 population activity drive motor output.

### III-88. Towards mapping the genesis of behavior in a complete nervous system

Alexander Katsov<sup>1</sup>  
Cornelia Bargmann<sup>1,2</sup>

AKATSOV@GMAIL.COM  
CORI@ROCKEFELLER.EDU

<sup>1</sup>Rockefeller University

<sup>2</sup>Howard Hughes Medical Institute

Nervous systems support coordination of activity between diverse neuron types and ensembles. Activity emerges during development, when it can influence the functional properties of the mature nervous system in both vertebrates and invertebrates. Early in development, neuronal networks show distinctive activity patterns such as global synchrony, bursts, or travelling waves, but the full spectrum of developmental activity, its effects, and its transition to mature function have not been defined. As a first step towards describing functional maturation in a complete nervous system, we are imaging developmental neuronal activity in the nematode *C. elegans* beginning in embryonic stages. Embryos move rapidly in their eggshells in the second half of embryogenesis, challenging continuous observation. To overcome this, we built a custom microscope based on selective plane illumination (Huisken et al. 2004), optimized for fast, multicolor whole-embryo imaging at cellular resolution in developing *C. elegans*. This microscope allowed us to monitor coordinated neuronal activity, reported by the Ca<sup>2+</sup> indicator GCaMP6, throughout the entire nervous system of developing embryos. We captured sub-second Ca<sup>2+</sup> modulation in developing neurons over hours of recording as the embryo moved and paused. Complex neural activity was observed throughout the nervous system during late embryogenesis, with especially robust Ca<sup>2+</sup> modulation in the developing motor neurons of the ventral nerve cord during embryonic movements. Automated neuron identification in longer recordings will be necessary to follow the entire period of embryonic neuronal activity. Current efforts focus on identifying the exact neurons that are active, and understanding whether and how these neurons regulate coordinated movements of the *C. elegans* embryo. This work aims to open the possibility of systematically mapping behavior coordination as it emerges in a complete nervous system.

### III-89. Precise temporal encoding in a vocal motor system

Claire Tang<sup>1,2</sup>  
Kyle Srivastava<sup>3,2</sup>  
Diala Chehayeb<sup>2</sup>  
Ilya Nemenman<sup>2</sup>  
Samuel Sober<sup>2</sup>

CLAIRE.TANG@UCSF.EDU  
KYLE.SRIVASTAVA@GMAIL.COM  
DIALA.CHEHAYEB@EMORY.EDU  
ILYA.NEMENMAN@EMORY.EDU  
SAMUEL.J.SOBER@EMORY.EDU

<sup>1</sup>University of California, San Francisco

<sup>2</sup>Emory University

<sup>3</sup>Georgia Institute of Technology

The sequences of action potentials produced by neurons are the medium of information transmission in the nervous system. In principle, neurons could transmit information solely through their mean firing rates or via the precise temporal patterning of their spike trains. While the importance of precise spike timing has been demonstrated in many sensory systems, the role of temporal encoding in controlling motor output is largely unknown. Using a range of experimental and computational techniques, we tested the hypothesis that significant information about motor output is represented by spike timing in the songbird vocal control system. We recorded spike trains from a motor cortex analog in songbirds and performed both metric-space analysis as well as direct computations of the mutual information between neural activity and upcoming vocal behavior. Our results indicate that premotor neurons convey information via spike timing far more often than via spike rate and that the amount of information conveyed at the millisecond timescale is far greater than that available in spike counts. Furthermore, a smaller dataset recorded from a single motor unit in a vocal muscle shows that spike timing in motor effectors, down to a few milliseconds' precision, can convey information even when none is available from a rate code. Our results therefore demonstrate that significant information about upcoming behavior can be represented by spike timing in

motor circuits and suggest that timing variations can evoke significant differences in motor output.

### III-90. Conserved and independent population activity between arm and Brain-Machine Interface control

Hagai Lalazar<sup>1</sup>  
L.F. Abbott<sup>1</sup>  
Eilon Vaadia<sup>2</sup>

HAGAI.LALAZAR@YAHOO.COM  
LFA2103@COLUMBIA.EDU  
EILON.VAADIA@ELSC.HUJI.AC.IL

<sup>1</sup>Columbia University  
<sup>2</sup>Hebrew University

Motor Brain-Machine Interface (BMI) is a therapeutic application to restore motor control to immobilized patients. Previous studies highlighted the changes in neural activity that accompanied improved performance using a BMI, while others focused on the neuronal tuning properties that were similar between BMI and arm control. However, a fundamental understanding of what components of neural activity are shared or differ when using these two very different effectors is still lacking. In our setup, during a motor task initially learned using only arm control, monkeys succeeded in using BMI control within seconds. Since learning to control a novel effector cannot occur so fast, we hypothesized that some components of the neural activity must be shared between arm and BMI control; yet since during BMI the arm doesn't move, other components must be different. In this study we focused on the relation of M1 neuronal activity and the static position of the cursor during the target hold period. During arm control, neuronal tuning functions exhibit diverse forms of complex nonlinearities, which depend on the posture of the entire arm, and underlie population activity that is high-dimensional. These nonlinear components are not residual noise, as they were informative for decoding the EMG of major muscles. By using dimensionality reduction techniques we found that there are several dimensions along which the population activity is significant only during arm control but not BMI. These are presumably the dimensions used to control the muscles, which are quiescent during BMI control when the arm is at rest. On the other hand, population activity varied as a function of target position along other dimensions, irrespective of modality. These population activity 'maps' showed a systematic deformity of the representation of hand position in neural population space, but with structure that was conserved across different arm configurations and BMI control.

### III-91. Single-trial prediction of slow or fast reaction from MEG brain activity before movement

Ryu Ohata<sup>1,2</sup>  
Kenji Ogawa<sup>3</sup>  
Hiroshi Imamizu<sup>1</sup>

OOHATA@ATR.JP  
OGAWA@LET.HOKUDAI.AC.JP  
IMAMIZU@ATR.JP

<sup>1</sup>Cognitive Mechanisms Laboratories, ATR  
<sup>2</sup>Osaka University  
<sup>3</sup>Hokkaido University

Predicting slow responses with long reaction time (RT) is of importance because these predictions can assist in preventing vehicle accidents due to delayed braking. Many studies have successfully predicted motor performance (e.g. hand-movement direction) from activity in sensorimotor area directly linked to movement initiation. However, it is difficult to predict performance before an extrinsic trigger (go signal) due to absence of activity directly related to movement execution. In this study, we attempted to predict long or short RT using activity before the go signal by extracting activation patterns from the entire brain, not limited to sensorimotor area. We combined two sparse Bayesian methods for efficient extraction of information related to RTs; 1) estimation of cortical currents using a hierarchical Bayesian method for improvement of spatial resolution from the number of magnetoencephalography (MEG) sensors, and 2) a sparse decoding algorithm for selecting critical information

for the prediction. Eight participants performed a delayed-reaching task using the right index finger. From the total number of trials, we selected 25% of the trials with the longest or shortest RT lengths. Using cortical currents in the delay period, we predicted whether one of the selected trials had a long or short RT length. Prediction accuracy was significantly above chance level ( $P < 0.05$ ) even 800 ms before the go signal onset. Spatial distribution of selected source currents for the prediction revealed that precentral area, supplementary motor area (SMA) and superior parietal area mainly contributed to the prediction around the go signal. While broad regions including prefrontal and occipital areas contributed during the early stage of the delay period. These results suggest that efficient extraction of activation patterns from broad regions provides an effective method for single-trial prediction of RT prior to an external trigger.

## III-92. Neuronal network dynamics within and between frontal and parietal cortex in the monkey

Benjamin Wellner<sup>1</sup>  
Steve Suway<sup>2</sup>  
Hans Scherberger<sup>1</sup>

BWELLNER@DPZ.EU  
SSUWAY@GMAIL.COM  
HSCHERBERGER@DPZ.EU

<sup>1</sup>German Primate Center

<sup>2</sup>University of Pittsburgh

Recently, techniques to record many neurons in parallel have greatly improved. This development has shifted research from single-spike analysis to the study of neuronal population activity and their interactions. Here, we focused on the parietal-frontal network for grasp movements of macaque monkey (area AIP and F5) that is well known for its role in sensorimotor integration. Monkeys performed a delayed grasping task consisting of a fixation period that was followed by a short cue, a memory period where the monkeys had to maintain fixation, and a movement epoch where the monkey had to grasp a handle with one of two grip types. We recorded single neuron activity of 128 channels in parallel (64 per area) and calculated cross-correlograms (CCG) of all possible pairs (following Smith and Kohn 2008, Ramalingam et al 2013). To analyze these complex interaction patterns we chose the approach of complex network analysis described in Rubinov and Sporns, 2010. As expected, different networks showed up for different grasp types. 'Motifs' were present consisting of a mixture of AIP and F5 neurons. Since one of the most important features of self-organized networks is small-worldness, we analyzed the degree of functional connections. Results were power law-distributed indicating the presence of few highly connected neurons (hubs) in addition to a majority of only sparsely connected neurons. One group of hub neurons showed strongly increased beta power in their CCG power spectrum and was located mainly in AIP, while a second group of hubs showed increased delta power and was more evenly distributed in AIP and F5. These findings suggest that communication in parietal-frontal action networks has a small-world architecture with distinct sets of hubs being functionally connected in the beta and delta frequency range. Supported by DFG research grant (SCHE 1575/1-1).

## III-93. Motor intention is marked by LFP beta amplitude suppression in a person with paralysis using a BCI

Tomislav Milekovic  
Beata Jarosiewicz  
Anish Sarma  
Leigh Hochberg  
John Donoghue  
Brown University

TOMISLAV\_MILEKOVIC@BROWN.EDU  
BEATA\_JAROSIEWICZ@BROWN.EDU  
ANISH\_SARMA@BROWN.EDU  
LEIGH\_HOCHBERG@BROWN.EDU  
JOHN\_DONOGHUE@BROWN.EDU

Local field potentials (LFPs) are electrical potentials recorded using intracortical microelectrode arrays. In motor

cortical areas, previous studies related amplitude modulations in the 'beta' frequency LFP band (BFB) to attention, movement planning or movement intention. In our study, a person with long-standing tetraplegia (S3) secondary to brainstem stroke was implanted with a 96-microelectrode array (Blackrock Microsystems) in the hand and arm area of her dominant (left) primary motor cortex (MI) as part of the BrainGate clinical trials. In 8 sessions, each recorded on a separate day, the participant controlled the movement and 'click' state of a computer cursor using a Brain-Computer Interface (BCI) based on action potentials (APs) recorded from the implanted array. During the sessions, we observed prolonged periods of suppressed activity in the BFB (here broadly defined as 9-30Hz), interrupted by large transient increases of activity (beta 'pulses'). We selected two subsets of data points from a session: (i) 'high beta', defined as all data points for which BFB amplitude surpassed a given threshold and (ii) 'low beta', the same number of points as in (i) with the smallest BFB amplitude. We then calibrated and assessed the decoding accuracy of a ridge regression decoder on these subsets. For thresholds exceeding 0.87 times the mean BFB amplitude, offline decoding of intended velocities based on APs recorded within the low beta subset was substantially more accurate than decoding within the high beta subset ( $P < 0.05$ , Wilcoxon ranksum test). Our results suggest that beta amplitude in MI of people with tetraplegia can be used as markers of engagement in voluntary movement intentions. In conjunction with AP-based BCIs, the detection of beta 'pulses' could be used as a 'pause' switch during ongoing BCI use, and periods of low BFB amplitude could be used for the automatic selection of data for decoder calibration.

### III-94. Unveiling neural code with a low dimensional model of birdsong production

Ana Amador<sup>1</sup>

Yonatan Sanz Perl<sup>1</sup>

Gabriel B. Mindlin<sup>1</sup>

Daniel Margoliash<sup>2</sup>

ANA.AMADOR@GMAIL.COM

YONISANZ@GMAIL.COM

GABO@DF.UBA.AR

DAN@BIGBIRD.UCHICAGO.EDU

<sup>1</sup>University of Buenos Aires

<sup>2</sup>University of Chicago

Songbirds are a well-studied example of vocal learning that allows integrating neural and peripheral recordings with a precisely quantifiable behavior. Although neural activity in the premotor forebrain nucleus HVC has been related to song acoustics in auditory playback experiments, it remains unresolved whether neural activity is related to song spectral structure during singing. To address this issue, we worked with a minimal physical model for bird-song production, having as an output a synthetic song. Each syllable was coded in terms of parameters related to air sac pressure and tension of the syringeal labia, defining motor 'gestures'. To validate this model, we assessed responses of HVC neurons to song playback in sleeping birds, as HVC neurons exhibit highly selective responses to the bird's own song (BOS). Remarkably, the mathematical model was able to elicit responses strikingly similar to those for BOS, with the same phasic-tonic features. These results demonstrate that a low dimensional model representing an approximation of peripheral mechanics is sufficient to capture behaviorally relevant features of song. Analyzing the HVC neurons responses to playback of each bird's own song, we observed that projection neurons were excited and interneurons were suppressed, with near-zero time lag, at the times of gesture extrema (defined as beginning, end or maxima of gestures). In this way, HVC neurons precisely encode the timing of extreme points of movement trajectories. We confirm these results with HVC recordings in singing birds. Given that HVC activity occurs with near synchrony to behavioral output, we propose that the activity of HVC neurons cannot be premotor as the accepted paradigm proposes, and therefore another model needs to be defined.

### III-95. Stereotyped, distributed networks revealed by whole-brain activity maps in behaving zebrafish

Claudia Feierstein<sup>1</sup>  
 Ruben Portugues<sup>2</sup>  
 Florian Engert  
 Michael Orger<sup>1</sup>

CLAUDIA.FEIERSTEIN@NEURO.FCHAMPALIMAUD.ORG  
 RPORTUGUES@MCB.HARVARD.EDU  
 FLORIAN@MCB.HARVARD.EDU  
 MICHAEL.ORGGER@NEURO.FCHAMPALIMAUD.ORG

<sup>1</sup>Champalimaud Neuroscience Programme

<sup>2</sup>Harvard University

All animals process sensory information to produce appropriate behaviors, but how this sensorimotor transformation is achieved by neural circuits remains a fundamental question in neuroscience. We address this question using the optokinetic behavior (OKR), a response consisting of reflexive eye movements elicited by whole-field rotational motion. Even such a simple response can involve coordinated activity in hundreds of neurons distributed in areas throughout the brain, making it essential to record simultaneously the activity of many neurons. Advances in imaging techniques have made it possible to record the activity of tens of thousands of neurons, while raising fundamental challenges in how to extract neural circuit information from these large, complex datasets. Using two-photon imaging of zebrafish expressing GCaMP5G in all neurons, we recorded neuronal activity throughout the whole brain of zebrafish larvae with single-cell resolution, while they were engaged in behavior. To be able to relate neuronal responses to behavior, we tracked the fish's eye and tail movements. We identified a sparse, broadly distributed network of neurons active during smooth oscillations of the eyes elicited by periodic motion of whole-field visual stimuli. The active network has a highly symmetric structure, with areas on opposite sides of the brain active at opposite stimulus phases. Neuronal activity fell into distinct clusters that reflected sensory and motor processing. We designed a set of stimuli to dissociate sensory and motor aspects of the behavior, and using correlation analysis, we systematically mapped which signals are represented in different brain areas. We used image registration techniques to align brains of different fish, and comparing activity across multiple fish we showed that the network active during OKR is highly stereotyped. These whole-brain functional maps provide a comprehensive view of brain dynamics during sensorimotor integration and enable the characterization of a complete vertebrate neural circuit controlling behavior.

### III-96. How do rats and humans use sensory information to time their actions?

Adina Drumea  
 Mathew Diamond

ADRUMEA@SISSA.IT  
 DIAMOND@SISSA.IT

International School for Advanced Studies

Computational and systems neuroscience has focused on how the brain weighs sensory evidence to select between possible actions, while there has been less emphasis on how the brain decides when to respond. We present a paradigm where rats and humans integrate a continuous stream of tactile signals in order to select the correct time to respond. Rats receive vibrations applied to their whiskers and respond by withdrawing from the nose-poke. Each stimulus is a stochastic train of velocity values, filtered at 150 Hz, modulated by a sinusoidal wave of lower frequency (0.35-1 Hz): the final form is a noisy vibration whose intensity envelope oscillates over time. Rats are rewarded for responding during the high intensity intervals of the train (the first or any successive peak in the sinusoidal envelope). The amplitude and frequency of the sinusoid vary across trials, so selecting response time requires integrating the ongoing stream of velocities. The same experiments are performed with humans, who receive stimuli on their fingertip and respond by button press. Rats and humans perform significantly above chance: response times are bound to peak intervals of the sinusoid. On high frequency or low amplitude trials, rats and humans wait more cycles before responding. When sessions include a fixed number of trials, humans wait more cycles per trial than do rats. However, when allowed to do as many trials as possible in time-limited sessions, humans sacrifice accuracy for speed. We thus find that when the decision criterion shifts, human behavior resembles that of rats, suggesting that species differences in this task depend largely on the set



point along the speed-accuracy axis. Work is underway to obtain neuronal recordings in behaving rats in order to unravel the computations performed at the critical stages in the sensation-decision pathway.

### III-97. Dynamic motor-related input to visual interneurons in *Drosophila*

Anmo Kim  
Gaby Maimon  
Rockefeller University

AKIM03@ROCKEFELLER.EDU  
MAIMON@ROCKEFELLER.EDU

Active locomotion, like flight or walking, leads to an increase in the response gain of visual neurons in flies and mice. In addition to such long-term modulations in sensory processing, can motor actions influence visual processing on faster timescales in these model organisms, akin to how saccadic eye movements modulate primate vision every few hundred milliseconds? In this study, we describe a dynamic motor-related input to visual neurons in *Drosophila*. We glued flies to a custom platform and performed whole-cell patch clamp recordings during tethered flight from two specific classes of fourth-order visual neurons: horizontal system (HS) cells and vertical system (VS) cells. We found that in addition to their visual responses, HS cells show transient motor-related potentials during rapid left or right turns of the wings. These potentials were clear even in the absence of visual input, while flies viewed a blank screen and made spontaneous turns. The potentials appeared at the early phase of each wing turn event but did not persist during maintained turns. The amplitudes of the transient potentials were positively correlated with the magnitudes of the turns. The transient potentials were cell type specific — strongest in one HS cell, moderate in the other two, and very weak in VS neurons. We developed an algorithm to cull hundreds of individual rapid wing-turn events and motor-related potentials, and we measured their relative timing. Transient potentials were essentially synchronous with wing turns, but many individual potentials clearly initiated prior to detectable changes in wingbeat amplitude, suggesting an internal origin for the signal. In sum, we describe a cell-type-specific motor-related input to the fly visual system, which serves to modulate vision during rapid flight turns. This efference-copy-like signal in a genetic model organism provides an attractive platform to study the functions of dynamic motor signals in sensory processing.

### III-98. Motor cortical activity predicts behavioral corrections following last-minute goal changes

Katherine Ames  
Krishna Shenoy  
Stanford University

KCAMES@STANFORD.EDU  
SHENOY@STANFORD.EDU

Studies of motor preparation and execution often separate these processes using delayed reaches. However, moving in an unpredictable environment requires behavioral flexibility in order to respond to new information. The motor system is therefore presumably able to perform computations typically associated with motor preparation (such as target selection) even while executing a movement. To probe the the motor system's response to task goal changes, we had monkeys (K, S) perform a delayed-reaching task in which on 20% of trials the target changed locations after the go cue but before movement. Both monkeys exhibited a variety of reaches on switch trials: some were initiated toward the new target, while others traveled in the direction of the original target before diverting to the new target. The time between the target switch and movement onset was a strong predictor of whether the motor system would need to correct the reach online: the less time between the target switch and movement onset, the further the arm would travel in the wrong direction. We examined whether neural activity provides a signature of commitment to a given reach. We recorded neural activity in M1 and PMd. We generated a vector connecting the average preparatory activity and the average movement-onset activity for a given non-switch reach. The position of the neural state along this vector shows how close neural activity has come to initiating this reach. Projecting neural activity at the time of the target switch onto this vector explains an average

of 27% (22%) of the variance in reaching path length in Monkey K (S). This indicates that the degree to which monkeys can correct their reach before moving can be partly explained simply by how far motor cortical activity has progressed toward generating that movement before the target changes.

### **III-99. Dynamic range adaptation in motor cortical neurons**

Robert Rasmussen<sup>1</sup>  
Andrew Schwartz<sup>1</sup>  
Steven Chase<sup>2</sup>

ROB@RGRASMUS.COM  
ABS21@PITT.EDU  
SCHASE@CMU.EDU

<sup>1</sup>University of Pittsburgh

<sup>2</sup>Carnegie Mellon University

Neurons from various sensory regions have been shown to demonstrate gain control or dynamic range adaptation in response to changes in the statistical distribution of their input stimuli. Although a similar phenomenon has been observed in single neurons in primary motor cortex (M1) during an isometric force task (Hepp-Raymond, et. al. *Exp Brain Res* 1999 128:123-133), the mechanisms behind this dynamic range adaptation are not well understood. We designed an experiment to test whether populations of neurons in M1 reconfigure their activity to adapt to changes in the statistics of motor output. We trained two monkeys to control a computer cursor using a brain-computer interface (BCI) in both two- (2D) and three- (3D) dimensions. During each recording session, the subject performed blocks of trials under 2D and 3D control with the same neurons. These tasks were designed to share 8 identical targets in the xy-plane. We computed the dynamic range of individual neurons as the difference between the trial-averaged maximum and minimum firing rates to these 8 identical targets, and compared the dynamic range when the targets were presented in a 2D context or a 3D context. The majority of neurons in M1 exhibited an increased dynamic range to the targets presented in the 2D context relative to their dynamic range during 3D. To identify the mechanism behind this finding, we investigated whether these dynamic range increases were correlated across the population in a manner that might reflect changes in the intended movement. Instead, our data are better explained by individual, uncorrelated changes in dynamic range within single neurons. These contextual changes in tuning provide evidence that the activity of M1 neurons operates as a high-level control signal, encoding context-relevant information within a limited dynamic range.

### **III-100. Pushing in the wrong direction: optogenetic perturbation misaligns with motor cortical dynamics**

Daniel O'Shea<sup>1</sup>  
Werapong Goo<sup>1</sup>  
Paul Kalanithi<sup>1</sup>  
Ilka Diester<sup>2</sup>  
Karl Deisseroth<sup>1</sup>  
Krishna Shenoy<sup>1</sup>

DJOSHEA@STANFORD.EDU  
WGOO@STANFORD.EDU  
PKALANITHI@STANFORD.EDU  
ILKA.DIESTER@ESI-FRANKFURT.DE  
DEISSERO@STANFORD.EDU  
SHENOY@STANFORD.EDU

<sup>1</sup>Stanford University

<sup>2</sup>Ernst Strungmann Institute

Several reports have recently demonstrated optogenetic modulation of behavior in primates; however, these behavioral effects are modest relative to electrical stimulation. We asked whether the contrast between strong neural modulation and weak behavioral modulation results from a misalignment between the optogenetic perturbation and the patterns of neural activity relevant to the behavioral task. Optogenetic excitation of dorsal premotor (PMd) cortex slows reaction times in an instructed-delay reaching task when delivered at the go cue, but leaves movement kinematics essentially unaltered. We report here that optogenetic stimulation of PMd and primary motor cortex (M1) delivered during movement similarly leaves unaltered reach kinematics. We recorded neuronal pop-

ulation responses to stimulation and found that many units displayed firing rates increased by 25-200 Hz out to 3 mm from the light source. The variance of these stimulation-induced changes in neural firing rates was 6.4-13.2 times the variance exhibited normally during non-stimulated trials. We applied PCA to non-stimulated firing rates to find a low-dimensional subspace explored by task-activity, consistent with previous reports. We found that the majority (83.1-91.6%) of stimulation-induced variance affected patterns of neural activity orthogonal to this task subspace, providing direct evidence that unexplored directions of neural activity may indeed be irrelevant to task performance. Furthermore, we used demixing PCA to demonstrate that the remaining stimulation variance primarily affects reach-condition independent modes of task-related activity. Consequently, only 4.4-6.2% of overall stimulation projects onto condition-dependent modes of task-related activity, which likely mitigates the behavioral impact. These results suggest that better aligning the vector of optogenetic perturbation with behaviorally-relevant axes of neural activity, e.g. via genetic and anatomical opsin targeting or spatial light patterning, could enhance behavioral modulation and more informatively probe the relationship between neural activity and behavior.

## Author Index

- Abbott L., 33, 123, 186, 190, 209  
Abouzari M., 109  
Acerbi L., 67  
Adam V., 127  
Advani M., 176  
Agarwal G., 198  
Agnes E. J., 97  
Ahmammad P., 38  
Ahissar E., 135  
Ahissar M., 157  
Ahmadian Y., 88  
Ahn D., 134  
Ahrens M., 32, 154  
Aitchison L., 94, 119  
Akaishi R., 73  
Akrami A., 77  
Albarran E., 55  
Albeanu D., 125  
Albers C., 180  
Alemi-Neissi A., 191  
Aljadeff J., 82  
Alkema M., 49  
Alonso J., 150  
Alonso L., 89  
Aluisio L., 155  
Amador A., 211  
Amari S., 109  
Amemori K., 205  
Ames K., 213  
Anastassiou C. A., 104  
Angelucci A., 150  
Anselmi F., 125  
Archer E., 162  
Arkadir D., 202  
Averbeck B., 71  
Avermann M., 93  
Avery M., 49  
Awwad B., 166  
Axel R., 123
- Baccus S. A., 172  
Bagdasarian K., 135  
Baker A., 130  
Baker B., 130  
Balasubramanian V., 117, 124  
Baldassi C., 191  
Baldauf D., 56
- Balle J., 89  
Bandyopadhyay A., 125  
Bao S., 128  
Barack D., 62  
Barahona M., 104  
Barak O., 63, 91  
Bargmann C., 208  
Barnes S., 110  
Barral J., 101  
Barrett D. G., 86, 170  
Barron H., 75  
Barthelemy F., 141  
Bartumeus F., 37  
Batista A., 39  
Bauer F., 100  
Bauza M., 148  
Bazhenov M., 181  
Beck J., 169  
Bednar J., 148  
Beer R. D., 52  
Behabadi B., 79  
Behnia R., 31  
Behrens T., 64, 75, 201  
Bellay T., 153, 160  
Ben-Shahar O., 142  
Ben-Tov M., 142  
Bennett D. V., 32  
Bensmaia S., 47  
Berdonini L., 119  
Berdyeva T., 155  
Berens P., 134  
Berenyi A., 198  
Berman G., 37  
Bermudez M., 141  
Bernacchia A., 76  
Berry M., 114, 117, 118, 164  
Berteau S., 88  
Besserve M., 199  
Bethge M., 116, 134, 179  
Bhalla U., 119, 185, 190  
Bhardwaj M., 68  
Bharioke A., 167  
Bialek W., 37, 118  
Bijanzadeh M., 150  
Billeh Y. N., 104  
Billings S., 139

- Bittner K., 123  
Blabe C., 207  
Blattler F., 203  
Bogacz R., 70  
Bonaventure P., 155  
Bonnen K., 158  
Boorman E., 201  
Born R., 43  
Borzello M., 138  
Bosman C., 146  
Botvinick M., 200  
Bouchard K., 51  
Brague J. C., 183  
Brainard D. H. H., 116  
Brainard M., 51  
Branson K., 60  
Brendel W., 80  
Bressman S. B., 202  
Brito C. S. N., 30  
Brody C., 60, 200  
Bromberg-Martin E., 62  
Brown E., 46  
Brunel N., 92, 191  
Brunton B., 172  
Brunton S., 172  
Buckley C., 173  
Buesing L., 84, 178  
Buice M., 100  
Burbank K., 187  
Burge J., 132, 158  
Burns L., 155  
Buzsaki G., 198
- C. DeAngelis G., 58, 131  
Cafaro J., 106  
Cahill C., 131  
Cain N., 85, 100  
Calhoun A., 61  
Callaway E. M., 109  
Cannon J., 193  
Canopoli A., 203  
Carandini M., 165  
Carlin M., 194  
Carnevale F., 63  
Carroll S., 68, 76  
Carruthers I., 129  
Cecchi G., 89  
Chae H. G., 125  
Chaillet A., 82  
Chalasan S., 61  
Chambers B., 154  
Chao C. M., 143  
Chartrand T., 197  
Chase S., 214  
Chaudhuri R., 90  
Chehayeb D., 208  
Chen J., 75  
Chen Y., 143  
Chen Z., 80  
Chichilnisky E. J., 134, 186  
Chitour Y., 82  
Chklovskii D. B., 139, 167, 184, 192  
Choi D., 37  
Choi J. H., 110  
Chou S. S., 160  
Christopoulos V., 67  
Churchland A., 28  
Churchland M., 40, 190  
Clandinin T., 31  
Clark C., 49  
Clark D., 31, 137  
Clemens J., 51  
Clopath C., 182  
Coca D., 139  
Cockburn J., 66  
Coen P., 51  
Coen-Cagli R., 104, 168  
Cohen M., 116, 160  
Cohen Y., 57  
Collins A., 201  
Constantinidis C., 76  
Cormack L. K., 158  
Corneil D., 177  
Corradi N., 119  
Costa R., 27  
Cotton R. J., 134, 145  
Cowley B., 81  
Crutchfield J., 84  
Cueva C. J., 40  
Curto C., 105, 173
- Dal Monte O., 71  
Dagleish H., 84  
Dan Y., 44, 53  
Daniel T., 172  
Danielson N., 36  
Davis M. B., 125  
de Lafuente V., 63  
de Silva N., 55  
Deger M., 93  
Dehaene G., 169  
Deisseroth K., 214  
Dekleva B., 174  
Deneux T., 141  
Deneve S., 34, 86, 113, 126, 170  
DePasquale B., 190  
Desimone R., 56, 146  
Desplan C., 31  
Desrochers T., 205  
DeWeese M., 84

- Dhawale A., 190  
 Diamond M., 77, 132, 212  
 DiCarlo J., 140, 144, 167  
 Dichter B., 30  
 Dickerson B., 172  
 Dickson B., 51  
 Diester I., 214  
 Doiron B., 160, 184, 196  
 Dolan R., 64, 75  
 Donchin O., 142  
 Donoghue J., 112, 210  
 Dotson N., 74  
 Doya K., 205  
 Druckmann S., 38  
 Drugowitsch J., 69  
 Drumea A., 212  
 Duan C. A., 60  
 Dudman J. T., 122, 159  
 Dugovic C., 155  
 Durstewitz D., 189
- E. Angelaki D., 58, 131  
 Eberle A., 172  
 Ebitz R. B., 55  
 Ecker A. S., 116, 134, 145  
 Effenberger F., 189  
 Eglen S. J., 149  
 Eldar E., 204  
 Elhilali M., 194  
 Elie J., 128, 165  
 Engelken R., 115  
 Engert F., 212  
 Erlich J., 60, 200  
 Escola S., 33  
 Eshel N., 41  
 Esmaeili V., 77  
 Evans J., 159
- Fairhall A. L., 35, 74, 113  
 Faisal A. A., 57, 78, 169  
 Faraji M., 56  
 Farkhooi F., 86  
 Fassih A., 77  
 Feierstein C., 80, 212  
 Fernandes H., 174  
 Ferrera V., 157  
 Festa D., 194  
 Field G., 134  
 Fischer B., 126  
 Fiser J., 68, 102  
 Fisher D., 187  
 Fitch K., 60  
 Fitzgerald J. E., 137  
 Fleischer J., 181  
 Fletcher A., 175
- Florence T., 122  
 Flores F., 46  
 Fox R., 158  
 Frady E., 153  
 Francis J., 87  
 Frank M., 64, 66, 201  
 Franklin N., 64  
 Franks K. M., 186  
 Freeman J., 32, 154  
 Freiwald W., 138  
 Fries P., 146  
 Froudarakis E., 134, 145  
 Fukai T., 97
- Gale J. T., 163  
 Gallant J., 144  
 Ganguli S., 31, 39, 95, 96, 176  
 Gao P., 31  
 Gardner J., 144  
 Gardner T., 41, 193  
 Garipey J., 62  
 Gatys L., 116  
 Geffen M., 45, 129  
 Geisler W., 132  
 Gekas N., 200  
 Gerstner W., 30, 56, 93, 97, 152, 183  
 Gifford A., 57  
 Gilja V., 207  
 Gilra A., 185  
 Girardin C., 51  
 Giusti C., 105  
 Gjorgjieva J., 35, 170  
 Glimcher P., 65  
 Gold J., 29, 71, 201  
 Golden J. R., 141  
 Goldman M., 35, 197  
 Golovko T., 189  
 Golub M., 39  
 Gomez-Marin A., 38  
 Goo W., 214  
 Goris R. L., 65, 89  
 Gosh K., 155  
 Gouvea T. S., 58  
 Grabska-Barwinska A., 179  
 Gray C., 74  
 Graybiel A., 205  
 Greschner M., 134  
 Gu Y., 58  
 Guetig R., 96, 99  
 Guitchounts G., 41  
 Gupta N., 181  
 Gupta P., 125
- Haas J., 183  
 Haefner R., 68

- Hagan M., 156  
Hagens O., 183  
Hahnloser R. H., 203  
Halassa M., 46  
Hamilton L., 128  
Hangya B., 42, 73  
Harnack D., 82  
Haroush K., 42  
Harris K., 165  
Haruno M., 59  
Harvey C., 120  
Hass J., 189  
Hasson U., 75  
Hattori D., 123  
Hausser M., 84  
Hawellek D., 72  
Hawrylycz M., 100  
He B., 90  
Heitman A., 134  
Hellgren Kotaleski J., 92  
Hemmelder V., 53  
Henderson J., 207  
Hennequin G., 94, 99, 194  
Hennig M. H., 92, 119  
Heo R., 110  
Hermundstad A., 117, 124  
Hershenhoren I., 166  
Hertaeg L., 189  
Hikosaka O., 62  
Hilgen G., 119  
Hillar C., 98  
Hillmann J., 99  
Hiratani N., 97  
Ho D., 133  
Hochberg L., 112, 207, 210  
Hoekstra H., 27  
Hofer S., 147, 148  
Honey C., 75  
Hong H., 140, 144, 167  
Horn M., 54  
Hosoya H., 145  
Houtman D., 166  
Hoy R., 141  
Hsu Y., 143  
Hu T., 139, 184, 192  
Hu Y., 103  
Huang W., 108  
Huk A., 43  
Hung C., 143  
Hunt L., 64  
Hyvarinen A., 145  
  
Ichida J., 150  
Imamizu H., 209  
Indiveri G., 177  
  
Ioffe M., 118, 164  
Iordanou J., 87  
Itskov V., 105  
Iyer R., 85, 100  
Izhikevich E., 187  
Izquierdo E., 52  
  
Jaegle A., 129  
Jarosiewicz B., 210  
Jayaraman V., 38  
Jeanne J., 50  
Jenvay S., 44  
Jessell T., 25  
Jin D., 51, 127  
Jin J., 150  
Jin L., 79, 161  
Joesch M., 135  
Josic K., 68, 76, 95, 145  
Jost J., 189  
Juusola M., 139  
  
Kabra M., 60  
Kahana M., 159  
Kaifosh P., 36  
Kalanithi P., 214  
Kamigaki T., 44  
Kang Y., 61  
Kanitscheider I., 104  
Kanwisher N., 130  
Kaplan H., 33  
Kapoor V., 53  
Karpova A., 60  
Kaschube M., 100  
Kastner D., 172  
Kato S., 33  
Katsov A., 208  
Katz L., 43  
Kaufman M., 40  
Kawashima T., 32  
Kaye A., 109  
Keck T., 110  
Kee T., 181  
Keil W., 149  
Keinath A., 122  
Keller G., 147, 148  
Kennedy A., 123  
Kepecs A., 42, 66, 73  
Kerr C. C., 87  
Khan A., 147, 148  
Khanna S., 107  
Khouri L., 166  
Kiani R., 65  
Kilpatrick Z., 76, 95  
Kim A., 213  
Kim H. R., 131

- Kim J. H., 110  
Kinney J., 66  
Klaus A., 101, 153, 188  
Klibaite U., 37  
Klier E. M., 58  
Knill D., 136  
Ko Y., 70  
Kobak D., 80  
Koch C., 104, 121  
Koechlin E., 69  
Koeppell K., 98  
Koerding K., 174  
Koester U., 98, 168  
Koh M., 125  
Kohn A., 81  
Kolling N., 73  
Komban S. J., 150  
Kopeck C., 200  
Kopell N., 193  
Koren V., 86  
Krakauer J., 28  
Kreiman G., 192  
Kremkow J., 150  
Krishnamurthy K., 124  
Kristan W., 153  
Krupic J., 148  
Kubanek J., 72  
Kuehn C., 188  
Kummerer M., 179  
Kunert J., 78  
Kutz J. N., 78, 172  
Kwon O., 136  
Kyweriga M., 131
- La Camera G., 202  
Lahiri S., 39  
Lajoie G., 83  
Lakshminarasimhan K., 58, 91  
Lalazar H., 209  
Laplagne D., 129  
Larsson J., 118  
Lashgari R., 150  
Latham P., 119, 179  
Latimer K., 162, 186  
Lau H., 70  
Lee D., 50, 108, 171  
Lee E., 71  
Lee H., 39  
Lee K., 110  
Lee S., 151  
Lee S. W., 195  
Lee T., 128  
Lee T. S., 151  
Leeds D., 83  
Leifer A., 49
- Lengyel M., 46, 94, 99, 194, 197  
Lesica N., 79, 110  
Levina A., 189  
Lewis C., 146  
Lewis L., 46  
Li P., 134  
Li Y., 71, 159  
Lieber J., 47  
Lim S., 35  
Lim Y., 41  
Lin C., 143  
Lin I., 165  
Linden D., 206  
Linderman S., 80  
Lindner B., 163  
Litke A., 134  
Litwin-Kumar A., 196  
Liu B., 206  
Liu C., 48  
Liu S., 58  
Loewenstein Y., 157  
Logothetis N., 199  
Lomber S. G., 43  
Longtin A., 166  
Looger L., 32  
Losonczy A., 36  
Lott G., 38  
Louie K., 65  
Louis M., 38  
Lovenberg T., 155  
Lucas T., 159  
Luna Ortiz C., 139  
Luongo F., 54  
Lyamzin D., 110  
Lytton W., 87
- Ma W. J., 54, 55, 68, 115  
MaBouDi H., 109  
Maccione A., 119  
Machens C., 80, 170  
Macke J., 84, 162  
MacLean J., 105, 154  
Magee J., 27, 123  
Magnasco M., 89  
Maheswaranathan N., 95  
Maimon G., 213  
Mainen Z., 80  
Majaj N., 140  
Makin J., 30  
Malik W., 112  
Manakov M., 60  
Manrique J., 102  
Marbach F., 125  
Marder E., 98  
Margoliash D., 211



- Markowitz J., 41, 193  
Markram H., 152  
Marre O., 114, 117, 164  
Marti D., 92  
Marzen S., 84, 98  
Masset P., 66  
Masson G. S., 141  
Mazzoni P., 202  
McCann B., 132  
McDermott J., 130  
Mease R., 35  
Mehraban Pour Behbahani F., 57  
Mehta P., 114  
Meier P., 187  
Meisel C., 188  
Meister M., 135, 170  
Mel B. W., 79, 161  
Menda G., 141  
Mensi S., 183  
Merel J., 158, 180  
Merlin S., 150  
Meyers E., 138  
Miconi T., 192  
Mihalas S., 85, 100, 121  
Milekovic T., 210  
Millard D., 47, 193  
Miller K., 88, 150, 200  
Miller L. E., 174  
Mindlin G. B., 211  
Miner D., 94  
Mishchenko Y., 175  
Mizuseki K., 198  
Mlynarski W., 126  
Modi M., 190  
Monteforte M., 115  
Monteiro T., 58  
Montgomery E. B., 163  
Moody W., 35  
Mooney R., 44  
Moore A., 48, 156  
Moore T., 55  
Moorkanikara Nageswaran J., 187  
Mora T., 124  
Motiwala A., 58  
Mrsic-Flogel T., 26, 147, 148  
Mu Y., 32  
Mulliken G., 146  
Murray J., 76  
Murthy M., 51  
Murthy V. N., 53, 124  
Musy M., 38  
Muthman O., 119  
Muzzio I., 122  
Mwilambwe-Tshilobo L., 45, 129  
Nassar M., 71  
Nassi J., 49  
Natan R., 45, 129  
Naud R., 93, 166  
Neftci E., 177  
Neishabouri A., 169  
Nelken I., 166  
Nelson A., 44  
Nemati S., 80  
Nemenman I., 37, 114, 208  
Nemri A., 148  
Nguyen-Vu T. B., 39  
Nihonsugi T., 59  
Nikbakht N., 132  
Nitzany E., 141  
Niv Y., 28, 202, 204  
Norman-Haignere S., 130  
Nowack A., 138  
Nurminen L., 150  
Nygren E., 182  
O'Donnell C., 196  
O'Leary T., 98  
O'Shea D., 214  
O'Doherty J., 195  
Ocker G. K., 196  
Ogawa K., 209  
Ohata R., 209  
Okun M., 165  
Oldenburg I., 121  
Oleskiw T., 138  
Oliver M., 144  
Ollershaw D., 193  
Olshausen B., 168  
Orchard J., 70  
Orger M., 212  
Orhan E., 115  
Orsolic I., 147, 148  
Osborne L., 206  
Ostojic S., 34, 92  
Otazu G., 125  
Otte S., 155  
Oyibo H. K., 125  
Pachitariu M., 79, 87, 147, 148  
Packer A., 84  
Pagan M., 112  
Palanski K., 37  
Palmer S. E., 105, 117, 133  
Pandarinath C., 207  
Pang R., 74  
Paninski L., 40, 154, 158, 180  
Parga N., 63, 102  
Park I. M., 43, 162  
Park M., 142

- Parmelee C., 173  
Pastalkova E., 105  
Pasupathy A., 138  
Paton J., 58  
Pawelzik K., 180  
Pehlevan C., 192  
Pelko M., 82  
Perrinet L., 148  
Pesaran B., 72, 156  
Peters M. A. K., 70, 178  
Petersen C., 93  
Pettit N., 84  
Pfau D., 154  
Pfeiffer M., 177  
Pfister J., 177  
Pillow J. W., 43, 142, 158, 162, 186  
Pirmoradian S., 119  
Pitkow X., 58, 91, 131  
Pizzorni-Ferrarese F., 118  
Platt M., 62  
Plenz D., 101, 153, 160, 188  
Pnevmatikakis E., 180  
Podlaski W., 152  
Poole B., 95, 96  
Poort J., 147, 148  
Portugues R., 212  
Potetz B., 151  
Pouget A., 91, 104, 169  
Pozzorini C., 183  
Prentice J., 164  
Preuschoff K., 56  
Proekt A., 89  
Proskurin M., 60  
Pyles J., 83
- Qian A., 158  
Quick K., 39  
Quiroga R. Q., 132
- Rabinowitz N. C., 89  
Radulescu A., 202  
Rajan K., 120  
Rajendran V., 38, 201  
Ramayya A., 159  
Ramirez-Villegas J., 199  
Ramkumar P., 174  
Ranade S., 42  
Rangan S., 175  
Ranjan R., 152  
Rasmussen R., 214  
Raviv O., 157  
Raymond D., 202  
Raymond J. L., 39, 197  
Reichert D. P., 136  
Reiser M., 122
- Reyes A., 101  
Reynolds J., 49  
Rieke F., 106, 186  
Rokni D., 53  
Romo R., 63, 76, 80  
Rose G., 166  
Rosen J., 32  
Rosenbaum R., 184  
Rowekamp R., 82  
Roy N., 76  
Rozell C., 47  
Rubin J., 184  
Rudiger P., 148  
Ruff D. A., 116  
Rushworth M., 73  
Rust N., 112  
Rylkova D., 206  
Ryu J., 151  
Ryu S., 31, 39  
Ryu W., 37
- Saal H., 47  
Sabatini B., 121  
Sabes P., 30  
Sadovsky A., 154  
Sadtler P., 39  
Saggau P., 134  
Sahani M., 79, 87, 127, 147, 148  
Salazar R., 74  
Salisbury J., 133  
Samonds J. M., 151  
Sanda P., 181  
Sanders J., 66, 73  
Santaniello S., 163  
Santhanam G., 31  
Sanz Perl Y., 211  
Sarma A., 210  
Sarma S., 87, 163  
Sasaki R., 131  
Savin C., 34  
Schachter M., 128  
Schafer P., 127  
Schaub M. T., 104  
Scherberger H., 210  
Schneider D., 44  
Schneidman E., 26  
Schnitzer M. J., 155  
Schottdorf M., 149  
Schrater P., 67, 176  
Schroedel T., 33  
Schulze A., 38  
Schwab D., 114, 133  
Schwalger T., 93, 163  
Schwartz A., 214  
Schwartz O., 168

- Schwarz C., 111  
Schwemmer M., 113  
Sederberg A., 105  
Seeholzer A., 152  
Seely J., 40  
Segev R., 142  
Seibert D., 144, 167  
Seitz A., 200  
Sejnowski T., 196  
Senn W., 182, 202  
Seo H., 50  
Seo M., 71  
Series P., 200  
Sernagor E., 119  
Serre T., 136  
Seshadri S., 153, 160  
Sevetson J., 183  
Shadlen M., 61  
Shaevitz J., 37  
Shamble P., 141  
Shams L., 178  
Sharan A., 159  
Sharpe J., 38  
Sharpee T., 61, 82, 109, 172  
Shatz C. J., 39  
Shea-Brown E., 83, 103, 106, 113  
Shen S., 54  
Shenoy K., 31, 40, 207, 213, 214  
Shephard C., 111  
Sher A., 134  
Shimazaki H., 109  
Shimojo S., 195  
Shiple F., 49  
Shlizerman E., 78  
Simoncelli E., 89, 112  
Simony E., 75  
Singer J., 192  
Sinz F. H., 134  
Smith M., 81, 107  
Smolyanskaya A., 43  
Snow M., 168  
Snyder A., 107  
Snyder L., 72  
Sober S., 208  
Sohal V., 54  
Sohl-Dickstein J., 95, 96, 128  
Soltani A., 55  
Soltanian-Zadeh H., 109  
Sommer F., 128, 198  
Sompolinsky H., 108, 170  
Sperling M., 159  
Sritharan D., 87  
Srivastava K., 208  
Srivastava N., 176  
Stanley G. B., 47, 111, 193  
Stepien A. E., 203  
Stern M., 82, 186  
Stevens C., 124  
Stevens J., 148  
Stewart W., 131  
Stocker A., 57, 171  
Stopfer M., 181  
Surace S. C., 177  
Sussillo D., 91  
Suvrathan A., 39  
Suway S., 210  
Sweeney Y., 92  
Szatmary B., 187  
  
Tadin D., 136  
Tafreshiha A., 132  
Tajima S., 85  
Talluri B. C., 148  
Tang C., 208  
Tank D., 120  
Tarr M., 83  
Tchumatchenko T., 101, 116, 182  
Teeter C., 121  
Teichert T., 157  
Tervo G., 60  
Theunissen F., 128, 165  
Thivierge J., 83  
Tian J., 41, 204  
Ticchi A., 78  
Tkacik G., 114, 117, 118, 164  
Tolias A. S., 134, 145  
Tomm C., 93  
Torrence H., 133  
Towfic Z. J., 192  
Toyoizumi T., 85, 173  
Trautmann E., 31  
Triesch J., 94  
Tsao D., 29, 138  
Tubiana J., 108  
Turaga S. C., 84  
Turner M., 106  
Tyler C., 151  
  
Uchida N., 41, 204  
Ujfalussy B., 197  
Ulanovsky N., 26  
Urbanczik R., 182, 202  
  
Vaadia E., 209  
van Opheusden B., 114  
van Rossum M., 82  
van Vreeswijk C., 86  
Vanzetta I., 141  
Vaughan A., 130  
Veliz-Cuba A., 95

- Venkadesan M., 38  
Verduzco-Flores S., 198  
Victor J., 141  
Vijayakumar S., 67  
Vladimirov N., 32  
Vogels T., 93, 152  
Voigts J., 46  
Vollmer A., 60
- Waiblinger C., 111  
Walczak A., 124  
Wallach A., 135  
Waltz J., 201  
Wanda P., 174  
Wang P., 123  
Wang Q., 193  
Wang X., 50, 76, 90  
Wang Y., 150  
Wang Z., 108, 171  
Webb R., 65  
Weber A., 47  
Wehr M., 48, 131, 156  
Weible A., 48  
Wellner B., 210  
Westkott M., 180  
White B., 102  
Wiese T., 160  
Williams P. L., 52  
Williams Z., 42  
Willis C., 107  
Wilson M., 46  
Wilson R., 50, 130  
Wittenbach J., 51  
Wolf F., 115, 149  
Wolpert D., 46, 67  
Womelsdorf T., 146  
Wong Y., 72  
Wyart V., 69
- Xu M., 44, 53
- Yamins D., 140, 144, 167  
Yang C., 32  
Yang G. R., 50  
Yang H., 101  
Yang S. C., 46  
Yarden T., 166  
Yaron A., 166  
Yates J., 43  
Yates J. L., 158  
Yatsenko D., 134, 145  
Yee H., 133  
Yeh L., 143  
Yildiz I., 126  
Yoshimoto J., 205
- Yu B., 31, 39, 81  
Yu S., 101  
Yukinawa N., 205
- Zai A. T., 203  
Zaremba J., 36  
Zecchina R., 191  
Zengyou Y., 206  
Zenke F., 97  
Zhang K., 108  
Zhang S., 44, 53  
Zhao G. Q., 39, 197  
Zheng H., 193  
Zimmer M., 33  
Ziv Y., 155  
Zoccolan D., 132  
Zuker C., 122  
Zylberberg J., 103, 106