COSYNE 2015

cosyne.org

MAIN MEETING

Salt Lake City, UT

Mar 5–8

Program Summary

Thursday, 5 March

4.00p	**Registration opens**
5.00p	**Welcome reception**
5:45p	**Opening remarks**
6.00p	**Session 1:** Invited speakers: Florian Engert, Nicole Rust
7.30p	**Poster Session I**

Friday, 6 March

7.30a	**Breakfast**
8.30a	**Session 2:** Invited speaker: Marla Feller; 4 accepted talks
10:45a	**Session 3:** Invited speaker: Shawn Lockery; 2 accepted talks
12.00p	**Lunch break**
2.00p	**Session 4:** Invited speaker: Liam Paninski; 3 accepted talks
3:45p	**Session 5:** Invited speaker: Emo Todorov; 3 accepted talks
5.30p	**Dinner break**
7.30p	**Poster Session II**

Saturday, 7 March

7.30a	**Breakfast**
8.30a	**Session 6:** Invited speaker: Matteo Carandini; 3 accepted talks
10.30a	**Session 7:** Invited speaker: Tatyana Sharpee; 3 accepted talks
12.00p	**Lunch break**
2.30p	**Session 8:** Invited speaker: Eitan Globerson; 4 accepted talks
5.30p	**Dinner break**
7.30p	**Poster Session III**

Sunday, 8 March

7.30a	**Breakfast**
8.30a	**Session 9:**
	Invited speaker: Amy Bastian; 3 accepted talks
10.30a	**Session 10:**
	Invited speaker: Wulfram Gerstner; 3 accepted talks
12.00p	**Lunch break**
2.00p	**Session 11:**
	Invited speaker: Sophie Deneve; 2 accepted talks

Inventors. Dreamers. Doers.

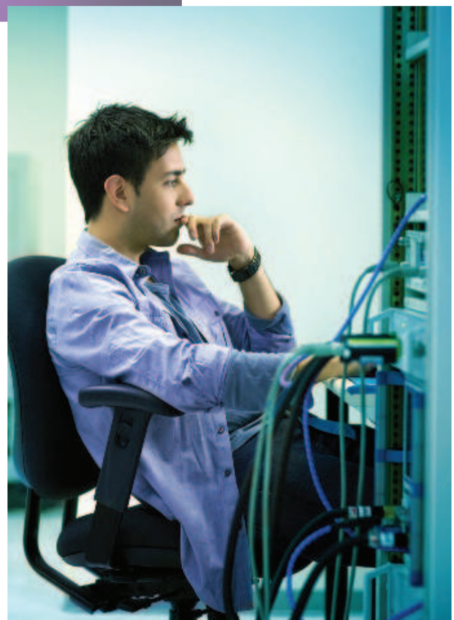

Qualcomm is a world leader in providing mobile communications, computing and network solutions that are transforming the way we live, learn, work, and play. We focus on a single goal—invent mobile technology breakthroughs.

As a pioneering technology innovator, we believe in the power of innovation which can transform the way people live and communicate. We are committed to continuing to foster invention both within Qualcomm and in the community. Our researchers and computational scientists engage in a wide variety of exciting and technically challenging projects—including exploring applications of systems neuroscience research to machine learning, to enable "smarter" and more efficient 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

QUALCOMM®

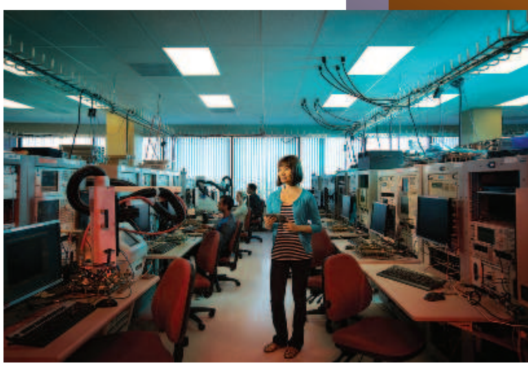

$$\sum_i \omega_i (x_i - \alpha)$$

SCiNDU: Systems & Computational Neuroscience Down Under

Tuesday 15th-Thursday 17th December, 2015

Queensland Brain Institute, The University of Queensland, Brisbane, Australia

Further information and Registration at: www.qbi.uq.edu.au/scindu

This conference brings together international leaders in understanding the computational principles underlying how neural circuits decode sensory information, make decisions, and learn from experience.

Speakers include:

Ehsan Arabzadeh (ANU)
Mark Bear (MIT)
Michael Breakspear (QIMR)
Allen Cheung (QBI)
Yang Dan (UC Berkeley)
Peter Dayan (UCL)
Geoffrey Goodhill (QBI)
Zach Mainen (Champalimaud)
Jason Mattingley (QBI)
Linda Richards (QBI)
Peter Robinson (Sydney)
Marcello Rosa (Monash)
Mandyam Srinivasan (QBI)
Greg Stuart (ANU)
Stephen Williams (QBI)
Li Zhaoping (UCL)

On December 15th the conference will be preceded by tutorials including:

Mark Bear: Experience-dependent synaptic plasticity
Peter Dayan: Neural reinforcement learning
Jason Mattingley: Brain stimulation, attention and plasticity
Li Zhaoping: Vision, efficient coding and salience

Australian Research Council
Centre of Excellence for
Integrative Brain Function

The University of Queensland
Australia
QBI
Queensland Brain Institute

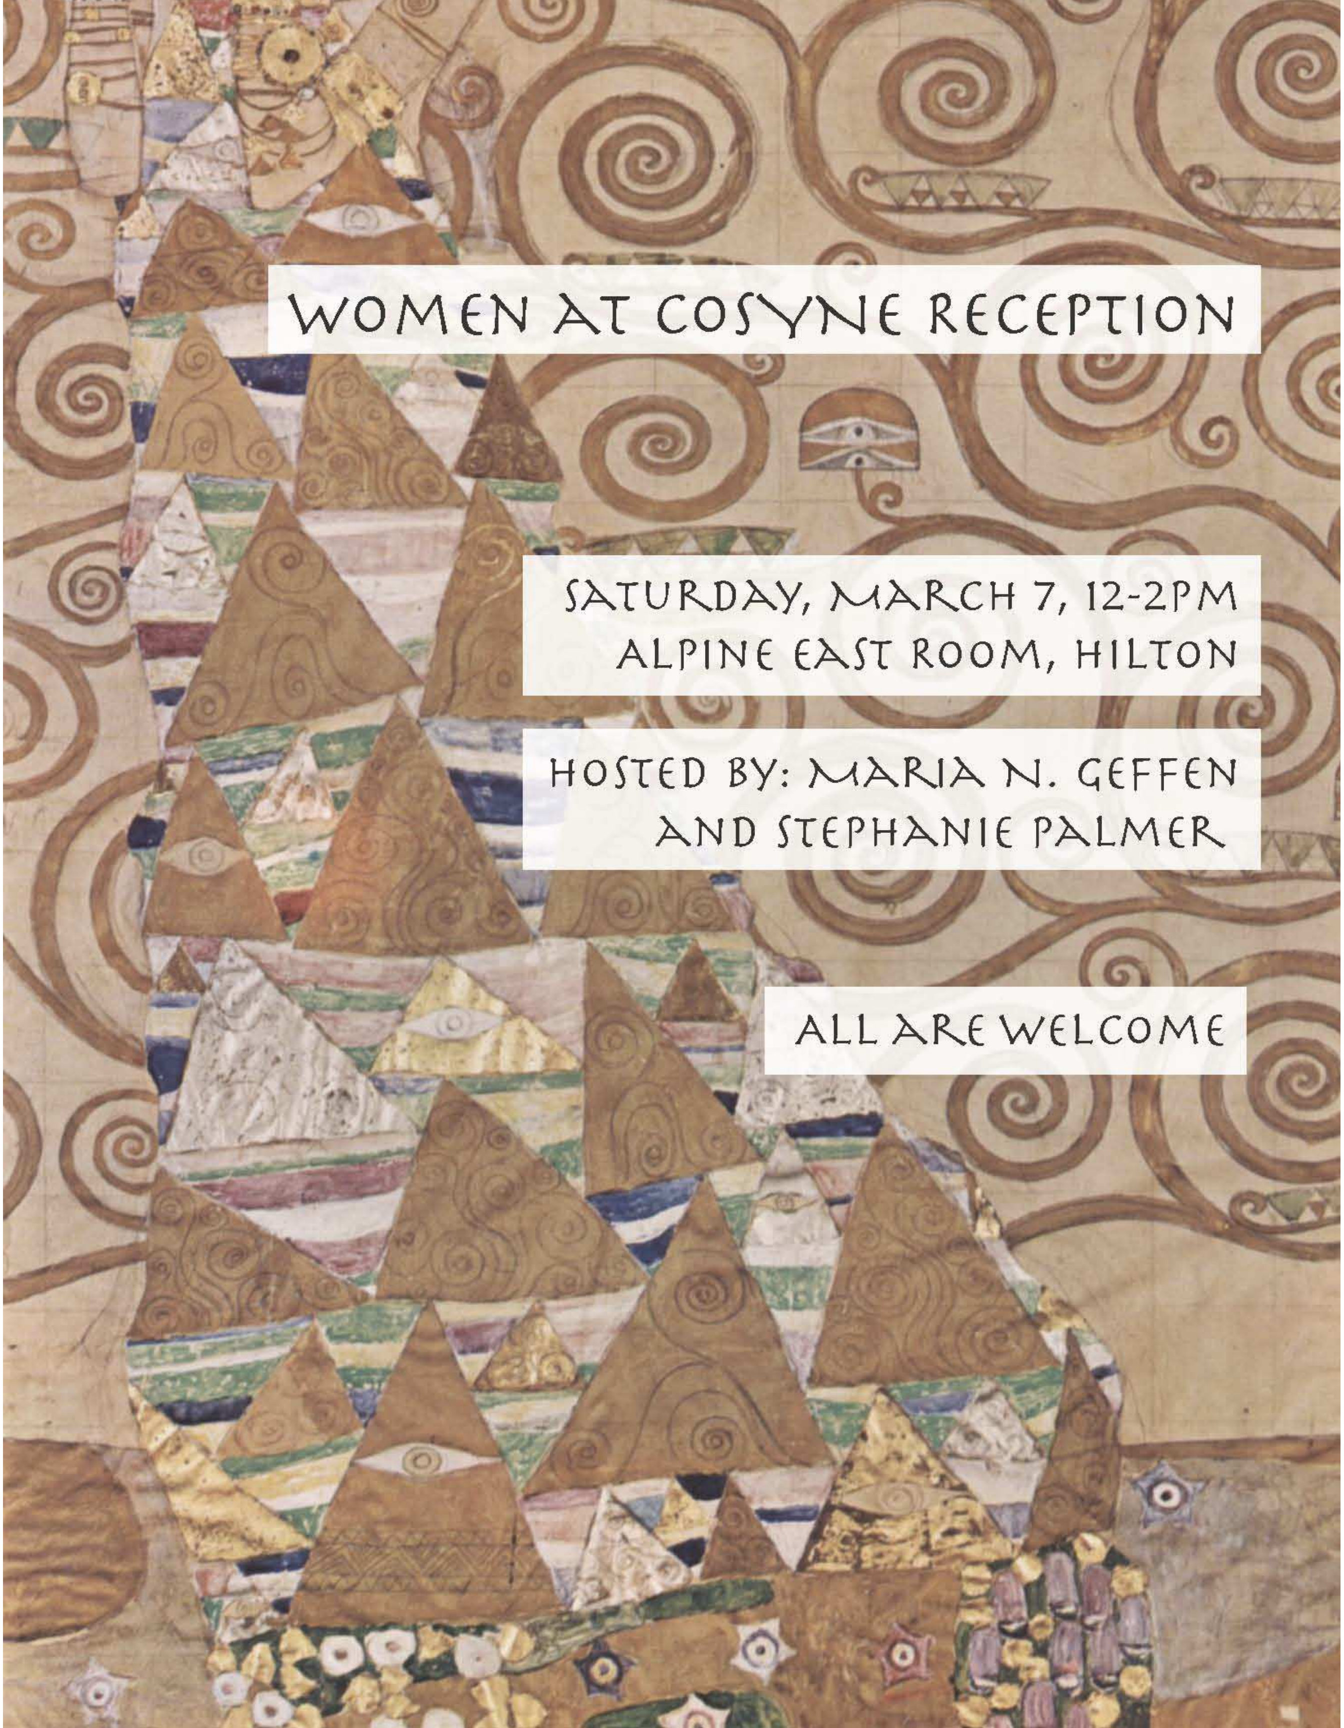
WOMEN AT COSYNE RECEPTION
SATURDAY, MARCH 7, 12-2PM
ALPINE EAST ROOM, HILTON
HOSTED BY: MARIA N. GEFFEN
AND STEPHANIE PALMER
ALL ARE WELCOME

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 2015 Leadership

Organizing Committee

General Chairs

Stephanie Palmer (University of Chicago), Michael Long (New York University)

Program Chairs

Maria Geffen (University of Pennsylvania), Konrad Kording (Northwestern University)

Workshop Chairs

Claudia Clopath (Imperial College), Robert Froemke (New York University)

Communications Chair

Xaq Pitkow (Rice University)

Executive Committee

Anne Churchland (Cold Spring Harbor Laboratory)

Zachary Mainen (Champalimaud Neuroscience Programme)

Alexandre Pouget (University of Geneva)

Anthony Zador (Cold Spring Harbor Laboratory)

Program Committee

Maria Geffen (University of Pennsylvania), co-chair
Konrad Kording (Northwestern University), co-chair
Misha Ahrens (Janelia Farm Research Campus)
Matthias Bethge (University of Tuebingen)
Gunnar Blohm (Queen's University)
Johannes Burge (University of Pennsylvania)
Daniel Butts (University of Maryland)
Damon Clark (Yale University)
Long Ding (University of Pennsylvania)
Brent Doiron (University of Pittsburgh)
Jay Gottfried (Northwestern University)
Richard Hahnloser (ETH Zurich)
Mehrdad Jazayeri (Massachusetts Institute of Technology)
Alla Karpova (Janelia Farm Research Campus)
Adam Kohn (Albert Einstein College of Medicine, Yeshiva University)
Kenway Louie (New York University)
Marta Moita (Champalimaud Centre for the Unknown)
Kanaka Rajan (Princeton University)
David Schneider (Duke University)
Tatjana Tchumatchenko (Max Planck Institute)
Srini Turaga (Janelia Farm Research Campus)
Ilana Witten (Princeton University)

Reviewers

Arash Afraz, Alaa Ahmed, Jonathan Amazon, Aaron Andalman, Wael Asaad, Behtash Babadi, Omri Barak, Andrea Benucci, Philipp Berens, Johanni Brea, Ethan Bromberg-Martin, Lars Buesing, Tim Buschman, Eugenia Chiappe, Jan Clemens, Ruben Coen-Cagli, Matthew Cook, Frederic Crevecoeur, Stephen David, Nai Ding, Takahiro Doi, Shaul Druckmann, Josh Dudman, Alexander Ecker, Steven Eliades, Bernhard Englitz, Jeffrey Erlich, Sean Escola, Kevin M Franks, Jeremy Freeman, Arko Ghosh, Julijana Gjorgjieva, Robbe Goris, Diego Gutnisky, Ralf Haefner, Timothy D Hanks, Benjamin Hayden, Alex Huk, Viren Jain, James Jeanne, Joseph Kable, Patrick Kanold, Matthias Kaschube, Christoph Kayser, Ann Kennedy, Ian Krajbich, Brian Lau, Arthur Leblois, Albert Lee, Andrew Leifer, Mate Lengyel, Talia Lerner, Nicholas Lesica, Ashok Litwin-Kumar, Matthieu Louis, Cheng Ly, Valerio Mante, Alexander Mathis, James McFarland, Melchi Michel, Jason Middleton, Ruben Moreno-Bote, Sarah Muldoon, Matt Nassar, Jean-Jacques Orban de Xivry, Michael Orger, Srdjan Ostojic, Anne-Marie Oswald, Adam Packer, Il Memming Park, Joseph Paton, Bijan Pesaran, Franco Pestilli, Jean-Pascal Pfister, Braden Purcell, Miguel Remondes, Alfonso Renart, Sandro Romani, Robert Rosenbaum, Arnd Roth, Herve Rouault, Simon Rumpel, David Schwab, Stephen Scott, Sam Solomon, Dominic Standage, Marcel Stimberg, Fatuel Tecuapetla, Dougal Tervo, Joji Tsunada, Glenn Turner, John Tuthill, Marcel van Gerven, Jonathan Victor, Iris Vilares, Melissa Warden, Andrew Welchman, Felix Wichmann, Ross Williamson, Jason Wittenbach, Byron Yu

Conference Support

Administrative Support, Registration, Hotels

Denise Acton, 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. Four award granting programs were available for Cosyne 2015.

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

- Burroughs Wellcome Fund
- Google
- National Science Foundation (NSF)
- The Gatsby Charitable Foundation
- Qualcomm Incorporated
- Brain Corporation

Cosyne Presenters Travel Grant Program

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

The 2015 recipients are:

Luigi Acerbi, Johnatan Aljadeff, Helen Barron, Philipp Berens, Matthew Chalk, S. Thomas Christie, Radoslaw Cichy, Ruben Coen-Cagli, Gamaleldin Elsayed, James Fransen, Vikram Gadagkar, Mona Garvert, Thiago Gouvea, Christopher Henry, Cynthia Hsu, Lacey Kitch, Laura Lewis, David Moses, Alon Rubin, James Sturgill, Jakob Voigts, Niklas Wilming, Yan Wu, Kelly Zalocusky

Cosyne New Attendees Travel Grant Program

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

The 2015 recipients are:

Vikas Bhandawat, Stephanie Chan, Yedidyah Dordek, Rodrigo Echeveste, Arseny Finkelstein, Grant Gillary, Jennifer Hobbs, Kishore Kuchibhotla, Baohua Liu, Tamas Madarasz, Hiroshi Makino, Jesse Marshall, Luca Mazzucato, Sunny Nigam, Dina Popovkina, Alexander Rivkind, Madineh Sedigh-Sarvestani, Abdul-Saboor Sheikh, Ryan Williamson

Cosyne Mentorship Travel Grant Program

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

The 2015 Cosyne Mentors are listed below, each followed by their mentee:

Christopher Honey and Kevin Himberger, Dori Derdikman and Gilad Tocker, Tatjana Tchumatchenko and Sara Konrad, Damon Clark and Emilio Salazar Cardozo

Cosyne Undergraduate Travel Grant Program

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

The 2015 recipients are:

Brian Aguirre Parada, Justin Chavez, Javier Cusicanqui, Monica Gates, Kayla Holston, Laetitia Jubin, Rachel Lee, Linhchi Nguyen, Gladynel Saavedra Pena, Milena Radoman, Eudorah Vital, Matthew Webster

Program

Note: Printed copies of this document do not contain the abstracts; they can be downloaded at:
http://cosyne.org/c/index.php?title=Cosyne2015_Program
Institutions listed in the program are the primary affiliation of the first author. For the complete list, please consult the abstracts.

Thursday, 5 March

4.00p **Registration opens**

5.00p **Welcome reception**

5:45p **Opening remarks**

Session 1:

(Chair: Maria Geffen, Konrad Kording)

6.00p A sensory motor circuit for binocular motion integration in larval zebrafish
Florian Engert, Harvard University **(invited)** 27

6.45p High-level representations arise from low-level computations during target search
Nicole Rust, University of Pennsylvania **(invited)** 27

7.30p **Poster Session I**

Friday, 6 March

7.30a **Continental breakfast**

Session 2:

(Chair: Xaq Pitkow)

8.30a The development and function of direction selective circuits in the retina
Marla Feller, University of California, Berkeley **(invited)** 28

9.15a Spatial decisions in the hippocampus
A. B. Saleem, M. Carandini, K. Harris, University College London 32

9.30a Calcium imaging in behaving mice reveals changes in hippocampal codes accompanying spatial learning
L. Kitch, Y. Ziv, M. Schnitzer, Stanford University 32

9.45a Grid cells reflect the locus of attention, even in the absence of movement
N. Wilming, P. Koenig, E. Buffalo, University of Osnabrueck 33

10.00a	An efficient grid cell decoder: are super-polynomial codes neurally plausible? N. Tran, I. Fiete, University of Texas at Austin	33
10.15a	**Coffee break**	

Session 3:

(Chair: Damon Clark)

10.45a	Neuronal and theoretical analysis of random walks in C. elegans foraging behaviors Shawn Lockery, University of Oregon **(invited)**	28
11.30a	Whole-brain imaging yields an embedding of the behavioral graph in neural activity space S. Kato, T. Schroedel, H. Kaplan, M. Zimmer, Institute of Molecular Pathology	34
11.45a	A neural basis for the patterning of spontaneous behavior in larval zebrafish T. Dunn, Y. Mu, S. Narayan, E. Naumann, C. Yang, O. Randlett, A. Schier, J. Freeman, F. Engert, M. Ahrens, Harvard University	35
12.00p	**Lunch break**	

Session 4:

(Chair: Srini Turaga)

2.00p	Challenges and opportunities in statistical neuroscience Liam Paninski, Columbia University **(invited)**	28
2.45p	The limits of Bayesian causal inference in multisensory perception L. Acerbi, T. Holland, W. J. Ma, New York University	35
3.00p	Cognitive cost as optimal control of metabolic resources S. T. Christie, P. Schrater, University of Minnesota	36
3.15p	Using speech-optimized convolutional neural networks to understand auditory cortex D. Yamins, A. Kell, S. Norman-Haignere, J. McDermott, Massachusetts Institute of Technology	36
3.30p	**Coffee break**	

Session 5:

(Chair: Kenway Louie)

4.00p	Synthesis of contact-rich behaviors with optimal control Emo Todorov, University of Washington **(invited)**	29
4.45p	Striatal dynamics explain duration judgments T. Gouvea, T. Monteiro, A. Motiwala, S. Soares, C. Machens, J. Paton, Champalimaud Centre for the Unknown	37
5.00p	Human visual representations are predicted in space and time by convolutional neural networks R. Cichy, A. Khosla, D. Pantazis, A. Torralba, A. Oliva, Massachusetts Institute of Technology	38
5.15p	Evidence that the ventral stream uses gradient coding to perform hierarchical inference E. Issa, C. Cadieu, J. DiCarlo, Massachusetts Institute of Technology	38
5.30p	**Dinner break**	
5.30p–7.00p	**Reception: Simons Collaboration on the Global Brain**	
7.30p	**Poster Session II**	

Saturday, 7 March

7.30a **Continental breakfast**

Session 6:

(Chair: Johannes Burge)

8.30a A canonical neural computation
Matteo Carandini, University College London **(invited)** 29

9.15a Efficient receptive field tiling in primate V1
I. Nauhaus, K. Nielsen, E. Callaway, University of Texas at Austin 39

9.30a Addition by division: a recurrent circuit explains cortical odor response regulation by SOM cells
F. Sturgill, P. Frady, J. Isaacson, University of California San Diego 39

9.45a Correlative and causal evidence that attention improves communication between cortical areas
D. Ruff, M. Cohen, University of Pittsburgh 40

10.00a **Coffee break**

Session 7:

(Chair: Brent Doiron)

10.30a Edge of stability in the hierarchy of neuronal types
Tatyana Sharpee, Salk Institute for Biological Studies **(invited)** 30

11.15a Stimulus driven inhibition by layer 6 underlies the neocortical processing of sensory change
J. Voigts, C. Deister, C. Moore, Massachusetts Institute of Technology 40

11.30a Interdigitated functional subnetworks in somatosensory cortex during active tactile behavior
S. Peron, F. Olafsdottir, J. Freeman, K. Svoboda, Janelia Research Campus 41

11.45a Optogenetic investigation of dopamine D2 receptor signaling and loss-sensitivity in a model of problem gambling
K. Zalocusky, T. Davidson, C. Ramakrishnan, T. N. Lerner, B. Knutson, K. Deisseroth, Stanford University 41

12.00p **Lunch break**

12.00p–2.00p **Women at Cosyne reception, Alpine East room**

Session 8:

(Chair: Long Ding)

2.30p Characteristic dynamics of value coding in normalizing decision circuits
K. Louie, T. LoFaro, R. Webb, P. Glimcher, New York University 42

2.45p A multi-layered olfactory population code enhances odor detection speed
J. Jeanne, R. Wilson, Harvard Medical School 43

3.00p A common mechanism underlies changes of mind about decisions and confidence
R. van den Berg, K. Anandalingam, A. Zylberberg, L. Woloszyn, R. Kiani, M. Shadlen, D. Wolpert, University of Cambridge 43

3.15p Prefrontal and midbrain contributions to fast executive control in the rat
C. A. Duan, J. Erlich, C. Kopec, C. Brody, Princeton University 44

3:30p	**Coffee break**	
4.15p	Google lecture Eitan Globerson, Jerusalem Academy of Music and Dance **(invited)**	30
5.30p	**Dinner break**	
7.30p	**Poster Session III**	

Sunday, 8 March

7.30a	**Continental breakfast**	

Session 9:

(Chair: David Schneider)

8.30a	Learning and relearning movement Amy Bastian, Johns Hopkins University **(invited)**	31
9.15a	Single-trial motor and cortical variability are tightly and bi-directionally coupled M. Pachitariu, C. Ames, B. Yu, G. Santhanam, S. Ryu, K. Shenoy, M. Sahani, Gatsby Computational Neuroscience Unit, UCL	44
9.30a	Neural network model of 3D head-direction tuning in bats A. Rubin, N. Ulanovsky, M. Tsodyks, Weizmann Institute of Science	45
9.45a	Zebra finch ventral tegmental area neurons encode song prediction error V. Gadagkar, E. Baird-Daniel, A. Farhang, J. Goldberg, Cornell University	45
10.00a	**Coffee break**	

Session 10:

(Chair: Tatjana Tchumatchenko)

10.30a	Turning the table on population coding: The balance is the key. Sophie Deneve, ENS **(invited)**	31
11.15a	Synapses represent and exploit estimates of uncertainty in their synaptic weight P. Latham, L. Aitchison, A. Pouget, Gatsby Computational Neuroscience Unit, UCL	46
11.30a	Complex synapses as efficient memory systems M. K. Benna, S. Fusi, Columbia University	46
11.45a	Selective rebalancing of cortical associations in humans via inhibitory plasticity H. Barron, T. Vogels, U. Emir, T. Makin, J. OŚhea, R. Dolan, T. E. J. Behrens, University College London	47
12.00p	**Lunch break**	

Session 11:

(Chair: Daniel Butts)

3.15p **Closing remarks**

Poster Session I

7:30 pm Thursday 5 March

Poster Session II 7:30 pm Friday 6 March

Poster Session III — 7:30 pm Saturday 7 March

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

This PDF can be downloaded at: http://cosyne.org/cosyne15/Cosyne2015_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. A sensory motor circuit for binocular motion integration in larval zebrafish

Florian Engert FLORIAN@MCB.HARVARD.EDU
Harvard University

Zebrafish process whole field visual motion, extract the net direction of such stimuli and use this information to guide their swimming behavior to match the direction and speed of these external cues. This innate behavior, called the optomotor reflex (OMR) is ubiquitous in the animal kingdom and presumably serves to stabilize an animal's position in the world when it is being moved by external forces. Here we use a closed loop behavioral assay in freely swimming fish that allows specific and independent stimulation of the two eyes - with coherent as well as conflicting motion signals. We can then answer questions of how the two eyes interact to combine, suppress and filter the various permutations of motion stimuli. We subsequently use whole brain imaging in tethered larvae to identify the complete neural circuitry underlying these various sensory motor transformations. Specifically we provide a working model of the complete circuit, that quantitatively captures the complete behavioral output as well as the response characteristics of the majority of the active neurons identified by independent cluster analysis. This rate based computational model makes very specific predictions about connectivity and synaptic polarity of the functionally identified neurons, is easy to test and falsify and serves as an ideal platform and hypothesis generator for a whole range of future experiment.

T-2. High-level representations arise from low-level computations during target search

Nicole Rust NRUST@PSYCH.UPENN.EDU
University of Pennsylvania

In its extreme, the proposal of "canonical cortical computation" predicts that every cortical brain area transforms incoming information in the same way and the only thing that differs between brain areas is the input that they receive. While this framework forms the backbone of many hierarchical models of neural processing (e.g. for visual object recognition), it conflicts with notions that neural processing gradually transitions from "simple" and "machine-like" at low levels to "complex" and "cognitive" at higher stages. In recent work, we have found that one high-level brain area, perirhinal cortex, reflects many of the signatures of high-level complexity: perirhinal responses exhibit "mixed selectivity", its computations contribute to the solution for a task that requires cognitive flexibility, and it produces this solution dynamically. Because perirhinal cortex is a high-level brain area, it is tempting to assume that these complex response properties largely arise from complex perirhinal processing; however, it is crucial to also consider the degree to which perirhinal responses reflect the many stages of processing leading up to its inputs. To describe how the inputs to perirhinal cortex - which largely arrive from inferotemporal cortex (IT) - are transformed by perirhinal cortex, we developed new techniques to fit LN models to population data recorded in IT and we compared the resulting model populations with our recorded perirhinal population. Remarkably, we found that a model similar to those used to describe V1 captured the most notable aspects of perirhinal computation, including changes in the amount and format of different types of signals relative to IT as well the timecourses with which these signals evolved. These results demonstrate that the complex responses of at least one high-level brain area result from simple, low-level computations in that structure, and they highlight the importance of using quantitative approaches to "demystify" high-level neural computation.

T-3. The development and function of direction selective circuits in the retina

Marla Feller MFELLER@BERKELEY.EDU
University of California, Berkeley

Does structure predict function? We address this question in the retina, which is comprised of multiple circuits that encode different features of the visual scene, culminating in the roughly 15 different types of retinal ganglion cells. Direction-selective ganglion cells (DSGCs) respond strongly to an image moving in the preferred direction and weakly to an image moving in the opposite, or null direction. I will present recent results from my laboratory regarding the organization and development of the retinal wiring diagrams that mediate this robust neural computation.

T-4. Neuronal and theoretical analysis of random walks in C. elegans foraging behaviors

Shawn Lockery SHAWN@UONEURO.UOREGON.EDU
University of Oregon

Random search is a behavioral strategy used by organisms from bacteria to humans to locate food that is randomly distributed and undetectable at a distance. We have been investigating this behavior in the nematode Caenorhabditis elegans, an organism with a simple, well-described nervous system. Here we formulate a mathematical model of random search abstracted from the C. elegans connectome and? fit to the first large-scale kinematic analysis of C. elegans behavior at submicron resolution. The model correctly predicts the electrophysiological sign and strength of key synaptic connections in the biological network. It also predicts the unexpected behavioral effects of neuronal ablations and genetic perturbations of the network's electrophysiological state. We propose that random search in C. elegans is controlled by a neuronal flip-flop circuit involving reciprocal inhibition between two populations of stochastic neurons. Our findings identify a new function for reciprocal inhibition and provide a testable neuronal model of random search in any organism.

T-5. Challenges and opportunities in statistical neuroscience

Liam Paninski LIAM@STAT.COLUMBIA.EDU
Columbia University

Systems and circuit-level neuroscience has entered a golden age: with modern fast computers, machine learning methods, and large-scale multineuronal recording and high-resolution imaging techniques, we can analyze neural activity at scales that were impossible even five years ago. One can now argue that the major bottlenecks in systems neuroscience no longer lie just in collecting data from large neural populations, but rather in understanding this data. I'll discuss several cases where basic neuroscience problems can be usefully recast in statistical language; examples include applications in calcium imaging of multiple neurons, inference of cell-type information from incomplete molecular data, and inference of network connectivity and low-dimensional dynamical structure from multineuronal spiking data.

T-6. Synthesis of contact-rich behaviors with optimal control

Emo Todorov ETODOROV@GMAIL.COM

University of Washington

Animals and machines interact with their environment mainly through physical contact. Yet the discontinuous nature of contact dynamics complicates planning and control, especially when combined with uncertainty. We have recently made progress in terms of optimizing complex trajectories that involve many contact events. These events do not need to be specified in advance, but instead are discovered fully automatically. Key to our success is the development of new models of contact dynamics, which enable continuation methods that in turn help the optimizer avoid a combinatorial search over contact configurations. We can presently synthesize humanoid trajectories in tasks such as getting up from the floor, walking and running, turning, riding a unicycle, as well as a variety of dexterous hand manipulation tasks. When augmented with warm-starts in the context of model-predictive control, our optimizers can run in real-time and be used as approximately-optimal feedback controllers. Some of these controllers have already been transferred to physical robots, via ensemble optimization methods that increase robustness to modeling errors. The resulting trajectory libraries are also used to train recurrent neural networks. After training the networks can control the body autonomously, without further help from the trajectory optimizer.

T-7. A canonical neural computation

Matteo Carandini MATTEO@CORTEXLAB.NET

University College London

There is increasing evidence that the brain relies on a set of canonical neural computations, repeating them across brain regions and modalities to apply similar operations to different problems. One of these computations is normalization, where the responses of neurons are divided by a common factor, which typically includes the summed activity of a pool of neurons. Normalization is thought to underlie operations as diverse as the representation of odors, the modulation of visual attention, the encoding of value, and the integration of multisensory information. In the visual system, normalization is thought to operate at multiple stages, and has been extensively studied in the primary visual cortex (V1). Normalization accounts for multiple nonlinear properties of V1 neurons. It governs the activity of V1 populations, making these populations operate in a summation regime or a winner-take-all regime depending on overall stimulus contrast.To probe the causal role of cortical circuits in normalization we used antidromic optogenetic stimulation to trigger spikes in a local region of mouse V1. This local activity caused two effects at distal V1 locations: summation and division. The balance between the two depended on visual contrast exactly as predicted by normalization. Intracellular recordings indicate that these effects are due to differential effects on synaptic excitation and inhibition. When visual contrast is zero, distal activation increases excitation more than inhibition. At high contrast, instead, distal activation increases inhibition more than excitation. We have thus established the causal synaptic basis for normalization in area V1, and possibly for the whole cortex. However, normalization in other systems such as the fly olfactory system is likely to involve different mechanisms. What is canonical about normalization and other putative canonical computations, are not the circuits but rather the computations.

T-8. Edge of stability in the hierarchy of neuronal types

Tatyana Sharpee SHARPEE@SALK.EDU
Salk Institute for Biological Studies

Computation in the brain involves multiple types of neurons, yet the organizing principles for how these neurons work together remain unclear. Recent experiments indicate that separate neuronal types exist that encode the same stimulus features, but do so with different thresholds. In this talk, I will show that the emergence of these types of neurons can be quantitatively described by the theory of transitions between different phases of matter. The theory accounts, without any adjustable parameters, for the coordination between two recently discovered types of salamander OFF retinal ganglion cells, as well as the absence of multiple types of ON cells. Furthermore, under changing environment conditions, the dynamic coordination between such cell types tracks the boundary of stability in the phase transition region. Such dynamic coordination between cell types can serve to maximize information transmission in a given environment while retaining the ability to quickly adapt to a new environment.

T-9. The construction of confidence in perceptual decisions

Mariano Sigman MARIUCHU@GMAIL.COM
University of Buenos Aires

Subjective confidence is used ubiquitously in approximate everyday expressions such as - “I think. . .”, “Maybe. . .”, “I am sure that”. In signal detection theory, confidence has a precise mathematical definition, indexing the probability that the decision is actually correct Neuroscience research has consistently shown that the brain can be close to optimal when performing perceptual inferences under uncertainty, integrating multiple sources of evidence weighted by their reliability. The emerging picture of optimal probabilistic inference in neuroscience, however, is in tension with the principles of behavioral economics. Decades of experimentation with “real-life” decision problems have shown that human confidence judgments exhibit reliable inconsistencies: they rely on sub-samples of the data, focus on tokens (representative exemplars), ignore the variance (or reliability) of the distribution, and over-weight evidence confirming their previous commitments and choices. In this talk I will present research in human psychophysics, neuronal data in monkeys in sequential decision making and computational models that aim to reconcile these two views.

T-10.

Eitan Globerson[1,2] GLEITAN@ZAHAV.NET.IL
[1]Jerusalem Academy of Music and Dance
[2]Bar Ilan University

T-11. Learning and relearning movement

Amy Bastian BASTIAN@KENNEDYKRIEGER.ORG
Johns Hopkins University

Human motor control and learning depend on a suite of brain mechanisms that are driven by different signals and operate on timescales ranging from minutes to years. Understanding these processes requires identifying how new movement patterns are normally acquired, retained, and generalized, as well as the effects of distinct brain lesions. This lecture will focus on studies of human motor behavior and learning during reaching and walking movements. It will highlight how our current understanding the effects of cerebellar damage on movement control has lead to a more general understanding of error-based learning that can be applied to treat individuals with cerebral stroke. Specifically, the lecture will cover: i) why cerebellar damage causes reaching incoordination or ataxia; ii) the effects of cerebellar damage on motor learning in reaching and walking; iii) normal acquisition, retention, and generalization of cerebellum-dependent motor learning; iv) the application of cerebellum-dependent mechanisms to cerebral stroke rehabilitation. The overall goal of the lecture is to highlight features of normal and abnormal motor learning, and how we can use this information to improve rehabilitation for individuals with neurological damage.

T-12. Turning the table on population coding: The balance is the key.

Sophie Deneve SOPHIE.DENEVE@ENS.FR
Ecole Normale Superieure

Cortical responses are hugely variable and heterogeneous. Single neurons adapt on multiple time scales and change their response properties according to brain states. Spontaneous activity does not look that different from stimulus evoked. Finally, the brain is amazingly robust, able to withstand multiple lesions and neural deaths without impacting behavior. But does that really matter, given than any stimulus recruits millions of cells? According to population coding approaches, many unreliable units vote for their preferred stimulus, achieving accuracy thought numbers. However, how does that fit with emerging data on exquisitely tuned microcircuits and synaptic plasticity rule dependent on precise spike timing? Neurons are embedded in highly recurrent networks and their activity is shaped by other neurons more that it is by the sensory input. This suggests that they collaborate to represent and compute, rather than voting independently. We will thus turn the table on population coding, and consider that spiking neural networks represent (collectively) their inputs as efficiently and robustly as possible. We will show that the excitatory-inhibitory balance, ubiquitous in cortical circuits, is a signature of this efficiency and robustness. Meanwhile, neural variability is not compensated by redundancy, but caused by degeneracy, e.g. multiple patterns of responses code for the same stimulus. The resulting population code is orders of magnitude more precise, and high capacity, than if neurons were spiking independently. As a consequence, neural tuning curves are not fixed properties of single neurons, but temporary network solutions to an optimization problem. Neural and behavioral adaptation can be understood as ways to limit metabolic costs without changing the code. And finally, oscillations, spontaneous activity and brain states as observed in cortex could be direct consequences of coding efficiently with noise.

T-13. Modeling synaptic plasticity: from synapses to function

Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH
Ecole Polytechnique Federale de Lausanne

Modeling synaptic plasticity has two major challenges: First, numerous electrophysiological experiments have shown that changes of synaptic connections depend on voltage, spike timing, firing rate as well as other factors and evolve on the time scale of milliseconds to hours. It would be desirable to have a model that covers all these aspects and accounts for synaptic plasticity as measured in slices. Second, the field of experimental and theoretical neuroscience is interested in synaptic plasticity because it has been linked to learning and memory formation. Therefore, a synapse model should also be functionally useful in large circuits, for example as a working memory. Typically, the two challenges are addressed separately: For example there are numerous models of induction of plasticity or models of working memory with fixed connectivity, but only few models that try to combine the plasticity with functional memory. In this talk, I will sketch how the two lines of research can be brought together.

T-14. Spatial decisions in the hippocampus

Aman B Saleem AMAN.SALEEM@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK

University College London

Place cells in the hippocampus encode an animal's position in space based on environmental cues [1]. How does their activity depend on the reliability of the cues? In an uncertain environment, does the hippocampus represent an estimate of position that guides spatial decisions? We addressed these questions by recording populations of CA1 neurons from mice that were trained to navigate a virtual corridor [2] and lick at a fixed location for a water reward. The only sensory cues to position were visual, and we manipulated their reliability by presenting them at one of three contrast levels (18, 60, and 72%). Animals performed a large fraction of trials ($>$80%) correctly at higher contrasts, and their performance dropped at low contrast. Place cells maintained the same preferred location and did not change their mean firing rate between contrast conditions. However, their place fields were narrower at higher contrast. We used an independent Bayes decoder to predict the position of animal at each moment from simultaneously recorded Ca1 neurons, and found predictions to accurate, especially at high contrast. Animal behavior correlated with hippocampal activity on a trial-by-trial basis. On error trials, animals licked in wrong positions of the corridor: too early or too late. However, when we decoded position based on hippocampal activity, we found that licks tended to occur when the hippocampus signaled the presence of the animal in the reward position. We could thus define a decision variable based on hippocampal population activity, which predicted whether the animal would lick early, correctly, or late. We conclude that the hippocampal place fields represent an animal's cognition of their position even in an uncertain environment. Its signals support spatial decisions, whether correct or incorrect.

T-15. Calcium imaging in behaving mice reveals changes in hippocampal codes accompanying spatial learning

Lacey Kitch[1] LJKITCH@STANFORD.EDU
Yaniv Ziv[1,2] YANIV.ZIV@WEIZMANN.AC.IL
Mark Schnitzer[1] MSCHNITZ@STANFORD.EDU

[1]Stanford University
[2]Weizmann Institute of Science

The mammalian hippocampus and its neuronal representations of spatial environments are thought to be crucial for spatial learning and memory. However, it remains unclear how changes in hippocampal neural codes relate to spatial learning over time scales of multiple days. Specifically, what changes in place cell firing patterns might underlie progressive improvements in spatial navigation? To address this question, we imaged the calcium dynamics of CA1 hippocampal place cells in freely-behaving mice as the animals learned to navigate a radial arm maze. To do this, we used a miniature fluorescence microscope, a chronic mouse preparation for long-term imaging of hippocampus, and the genetically encoded calcium indicator GCaMP6 targeted to CA1 pyramidal cells. This approach allowed us to observe the concurrent dynamics of hundreds of individual CA1 neurons during active mouse behavior. As mice learned to navigate the radial arm maze, we examined the accompanying changes in the CA1 neural ensemble representation of space. As the animals' performance improved over the course of five days, our analysis revealed a refinement in the ensemble representation of space, such that a progressively smaller subset of cells were active during maze running. Nevertheless, a greater percentage of active cells conveyed statistically significant spatial information in their activity patterns. Early in learning, there were typically multiple place fields per neuron. As behavioral performance improved, the number of place fields per cell declined, and cells were more reliably active as the mouse passed through the place field. A Bayesian decoding analysis revealed that errors in reconstructing the animal's trajectory from the calcium imaging data declined as learning advanced. This suggests that CA1 neural representations gradually increased in spatial accuracy as behavioral performance improved. Altogether, our findings support the notion that the representation of space becomes more reliable and spatially selective over the course of spatial learning.

T-16. Grid cells reflect the locus of attention, even in the absence of movement

Niklas Wilming[1,2] NIKLAS.WILMING@GMAIL.COM
Peter Koenig PKOENIG@UOS.DE
Elizabeth Buffalo[2] BETH.BUFFALO@WANPRC.ORG

[1]University of Osnabrueck
[2]University of Washington

Entorhinal grid cells allow for the precise decoding of the position of freely moving animals and have been implicated in spatial navigation (Hafting, Fyhn, Molden, Moser, & Moser, 2005; Sreenivasan & Fiete, 2011). However, in macaque monkeys, grid cells have been identified that reflect the location of eye movements, independent of locomotion through space (Killian, Jutras, & Buffalo, 2012). It has recently been suggested that grid cells might support cognitive functions that do not necessarily involve physical movement, and that the neuronal algorithms underlying navigation in real and mental space are fundamentally the same (Buzsáki & Moser, 2013). However, to date, the grid cell system has only been studied with tasks that involve spatial navigation or visual exploration. Here, we identify entorhinal grid cells in monkeys engaged in a covert attention task that requires no physical movement. We examined spatial representations of cells in the entorhinal cortex in monkeys covertly tracking a moving dot, whose movement covers space evenly. We found a significant proportion of recorded cells that exhibited spatial modulations in their firing rate that resemble grid-like patterns. The average gridscore of these cells was 1.1 and their average firing field modulation index was 0.12. In addition we find some cells that do not show grid-like firing fields, but are nevertheless spatially modulated. These cells might encode other spatial representations. The existence of grid-like firing fields during movement of covert attention suggests that the grid cell network in macaque monkeys does not necessarily rely on physical movement. These results support the notion that grid cells in the entorhinal cortex are capable of serving a variety of different cognitive functions and suggests that grid cells may represent a versatile component of many neural algorithms.

T-17. An efficient grid cell decoder: are super-polynomial codes neurally plausible?

Ngoc Tran TRAN.MAI.NGOC@GMAIL.COM
Ila Fiete ILAFIETE@MAIL.CLM.UTEXAS.EDU

University of Texas at Austin

Grid cells fire as a function of animal location. The spatial tuning of each cell forms a triangular lattice. These

lattices are translated, dilated and rotated across N different modules. Altogether, grid cells encode spatial location as a set of N two-dimensional phases modulo the lattice periods. The grid cell encoding capacity grows exponentially with N [Sreenivasan & Fiete 2011] and is the first known example of a near-Shannon-capacity population code in the brain. In theory, errors can be corrected by mapping to the nearest noise-free vector phase, which corresponds to maximum likelihood estimation under iid (truncated) Gaussian noise in each module's phase. Such a decoder can be naively implemented by a winner-take-all network, but the cost in neurons grows exponentially with N. The question is whether it is possible to construct a neurally plausible decoder that corrects errors at a nearly linear cost in N. Here, under the assumption that animal displacements in a short time interval are relatively small compared to the network capacity, we construct two small neural networks that can correct the grid code. Our 'decoder' does not build a separate representation of animal location, which in existing frameworks requires exponentially many neurons, exponentially large weights, or exponentially high firing rates. Rather, it computes the nearest coding state for noisy states, and maps the system back to it. With noiseless decoding neurons, this furnishes the first near-Shannon-capacity, neurally plausible encoder-decoder pair – that is, the information rate is asymptotically finite. If the decoding neurons are themselves noisy, the encoder-decoder pair loses a factor of $1/\log(N)$ in information rate, a considerable improvement over the $1/N$ scaling of most neural population coding models. These results are a key step in helping us understand whether the brain can in principle achieve super-polynomial representational performance as a function of neuron number.

T-18. Whole-brain imaging yields an embedding of the behavioral graph in neural activity space

Saul Kato KATO@IMP.AC.AT
Tina Schroedel TINA.SCHROEDEL@IMP.AC.AT
Harris Kaplan HARRIS.KAP@GMAIL.COM
Manuel Zimmer MANUEL.ZIMMER@IMP.AC.AT

Institute of Molecular Pathology, Vienna

Intrinsic brain dynamics are implicated in the generation of sophisticated behavior, although an explicit mapping between network dynamics and behavioral sequences has been elusive. Trial-to-trial variability is a seemingly inevitable consequence of intrinsic activity, hampering the interpretation of single-trial brain dynamics particularly when sparsely sampled. By combining whole-brain volumetric calcium imaging with neural recordings in freely moving C. elegans and a high-throughput analysis system, both co-developed by our group, we show, on a single trial basis, that the animal's high-level locomotory command state is robustly encoded in a large, fixed interneuron population. This suggests that a holistic behavioral state is widely communicated via shared activity; we additionally observe that many neurons possess activity subtly differentiated from the global signal by stereotyped phase relationships. We find that locomotory states correspond to sub-volumes of neural state space, segmenting the global attractor. Since locomotory state repeats but follows different sequences, behavior may be described as a cyclic state transition graph. By clustering population trajectory segments on the basis of individual neural transients, we can observe branching and merging of trajectory bundles corresponding to behavioral decisions, yielding, for the first time, an embedding of the cyclic behavioral graph in neural activity space. Next, by acutely inactivating specific input and output-oriented interneurons strongly participating in the global signal using chemical genetics, we show that overall dynamics is robust to single neuron perturbations even though strong behavioral effects can be elicited. The distributed nature of motor command dynamics we observe resolves overlapping results on behavioral roles of single neurons previously reported. Finally, we find that sensory circuit influences behavior by perturbing population trajectories rather than grossly changing attractor geometry. This study establishes a direct correspondence between intrinsic neural dynamics and behavioral state dynamics and defines the function of the global brain signal in this system.

T-19. A neural basis for the patterning of spontaneous behavior in larval zebrafish

Timothy Dunn[1] TWDUNN@FAS.HARVARD.EDU
Yu Mu MUY@JANELIA.HHMI.ORG
Sujatha Narayan NARAYANS@JANELIA.HHMI.ORG
Eva Naumann[2,1] EVA.AIMABLE@GMAIL.COM
Chao-Tsung Yang YANGC@JANELIA.HHMI.ORG
Owen Randlett OWEN.RANDLETT@GMAIL.COM
Alexander Schier[1] SCHIER@FAS.HARVARD.EDU
Jeremy Freeman[3,4] FREEMANJ11@JANELIA.HHMI.ORG
Florian Engert[1] FLORIAN@MCB.HARVARD.EDU
Misha Ahrens[3,4] AHRENSM@JANELIA.HHMI.ORG

[1]Harvard University
[2]University College London
[3]Janelia Farm Research Campus
[4]Howard Hughes Medical Institute

Even in the absence of sensory input or task demands, animals exhibit rich self-generated behavior that is essential for survival. We studied such self-generated behavior, and its neural correlates, in larval zebrafish. While swimming in featureless environments, zebrafish showed strong spatiotemporal structure in their locomotion, on time scales of up to tens of seconds. The animals, which swim in discrete swim bouts, strung together repeated turns to one direction, before switching to another. To identify the neural basis of this phenomenon, we used light-sheet imaging to record whole-brain activity while fish performed a fictive version of the same behavior. Whole-brain responses were characterized by fitting every imaging voxel with a unique behavioral tuning field capturing locomotion strength and directionality, analogous to visual receptive fields but for behavior rather than sensory input. Collapsing this tuning field onto one parameter described fictive turn angle, and whole-brain maps of this parameter revealed neuronal populations in the hindbrain that were either highly correlated with swim vigor or with turning. The populations correlated with turning exhibited similarly slow time courses as the behavior and included the previously identified “hindbrain oscillator” (HBO), a functionally defined neural structure conserved across fish. We interrogated this structure for causal relationships to spontaneous behavior, using cell ablations and stimulation to show that the HBO biases the direction of swimming. In addition, combinations of functional imaging and anatomical labeling revealed that the HBO comprises anatomically separated glutamatergic and GABAergic clusters, suggesting a mutual-inhibitory circuit motif. Finally, we modeled spontaneous swim behavior as a two-state Markov model and found that the observed swim statistics increase the efficiency of exploration on local scales when compared to a randomly walking fish. These findings establish a circuit underlying spatiotemporally structured spontaneous behavior that may support efficient exploration of environments in the absence of sensory cues.

T-20. The limits of Bayesian causal inference in multisensory perception

Luigi Acerbi[1] LUIGI.ACERBI@GMAIL.COM
Trevor Holland[2] FOXFACTION@GMAIL.COM
Wei Ji Ma[1] WEIJIMA@NYU.EDU

[1]New York University
[2]Baylor College of Medicine

In a large variety of perceptual tasks, humans integrate congruent multisensory cues near-optimally, that is according to their relative reliability. However, how exactly the nervous system determines that two cues belong together is still unclear. Bayesian causal inference has been proposed as a normative framework to explain how the percept of unity arises as the result of probabilistically inferring a common cause. Surprisingly, though,

this framework has never been put to the fundamental test of intermixing several levels of sensory noise in the same session. Here we perform a strong assessment of causal inference in multisensory integration by asking observers to estimate the location of audio-visual stimuli and perform yes/no judgements of unity. The key experimental manipulation is that in addition to audio-visual disparity, we also unpredictably changed visual cue reliability from trial to trial. To understand the data, we first fit the unimodal trials with a detailed Bayesian model that captures relevant features of observers' behavior such as local biases and eccentricity-dependent variability. We then use the fitted parameters to predict performance in the bimodal trials, with only a small number of additional free parameters (two to four). In all models we consider for the bimodal data, observers estimate the cue position for independent or fully fused cues in a Bayesian manner. The models differ in whether the observer uses Bayesian or non-Bayesian criteria to decide the amount of fusion between cues, i.e. to perform causal inference. Bayesian causal inference mostly fails to account for basic features of the data. Instead, the data are best fit by a model in which observers judge unity by applying a reliability-dependent distance criterion to the disparity of the auditory and visual measurements. This suggests that multisensory perception beyond the simple case of pure integration is probabilistic but far from Bayes-optimal.

T-21. Cognitive cost as optimal control of metabolic resources

S. Thomas Christie[1] TOM.CHRISTIE@GMAIL.COM
Paul Schrater[2] SCHRATER@UMN.EDU

[1]University of Minnesota
[2]University of Minnesota, Twin Cities

Cognitive processes that require vigilance, model-based lookahead, or extensive utilization of attention or working memory are said to incur a cost (Kool et al, 2010), to be aversive to use (Mcguire et al, 2010), and to be computationally expensive. General avoidance of such processes has led to characterizations of humans as lazy organisms, cognitive misers, who prefer fast and frugal heuristics or habits than deliberative thought. These statements raise the question of why certain types of cognition are costly. What, exactly, is being spent? We present a control-theoretic model of metabolic resource allocation that accounts for findings in the decision making literature as well as the effects of hypoglycemia on cognitive function. A critical component is the inclusion of astrocytic glycogen into the system energy dynamics. Glycogen acts as an energy buffer that can temporarily support high neural activity beyond the rate supported by blood glucose supply. Our model supersedes both the "cost/benefit" and "limited resource" models of cognitive cost while retaining valuable contributions of each. We show optimal energy allocation produces an effective cost on performance, for budgeting performance, a preference for model-free policies over look-ahead. We show our approach gives a unified explanation for the effects of drops in blood glucose on cognitive performance and provides a new mechanism for learning as resource investment.

T-22. Using speech-optimized convolutional neural networks to understand auditory cortex

Daniel Yamins YAMINS@MIT.EDU
Alexander Kell ALEXKELL@MIT.EDU
Sam Norman-Haignere SVNH@MIT.EDU
Josh McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

Despite many proposals, there is little consensus on the functional organization of auditory cortex. Here, we apply a novel computational approach to this problem, training a hierarchical convolutional neural network (CNN) to solve a high-level auditory task and comparing the representations that emerge with empirically-observed auditory cortex responses to natural sounds. We first optimized a variety of CNN architectures to perform a 600-way

word recognition task in the presence of high levels of complex background noise. Separately, using fMRI we measured responses throughout human auditory cortex to a broad set of natural sounds, including environmental sounds, music, mechanical sounds, speech, and animal vocalizations. We then computed the responses of each CNN model to these same natural sounds and used cross-validated linear regression to determine how well each model layer predicted each voxel's response profile. We have five main findings. First, across CNN architectures, there is a strong correlation between a model's word recognition performance and how well it predicts auditory cortical responses [r = 0.97]. Second, the CNN that performs best on the word recognition task achieves human-level performance. Third, this CNN explains nearly twice as much neural variance as a standard model of auditory cortex [Chi, Ru, Shamma, 2005]. Fourth, the CNN model recapitulates key known features of cortical organization, including a functional distinction between primary and non-primary cortical areas. Specifically, all model layers predict responses in primary auditory areas approximately equally well, but higher CNN layers predict non-primary cortical regions substantially better than lower layers. Fifth, we find that higher CNN layers predict posterior speech regions better than anterior speech regions, potentially suggesting a previously unknown cortical speech-processing hierarchy along the A-P axis. The results demonstrate the power of deep neural networks in understanding the organization of auditory cortex and the behavioral constraints that shape it.

T-23. Striatal dynamics explain duration judgments

Thiago Gouvea[1] THIAGO.GOUVEA@NEURO.FCHAMPALIMAUD.ORG
Tiago Monteiro[2] TIAGO.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG
Asma Motiwala[2] ASMA.MOTIWALA@NEURO.FCHAMPALIMAUD.ORG
Sofia Soares[2] SOFIA.SOARES@NEURO.FCHAMPALIMAUD.ORG
Christian Machens[2,3] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Joseph Paton[2] JOE.PATON@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]Champalimaud Neuroscience Programme
[3]Ecole Normale Superieure

Processing of time information is fundamental for sensation, cognition and action. However, the neural mechanisms that support time estimation are not well understood. We recorded from populations of neurons in the striatum, a brain area implicated in learning, motor function, and timing, in rats as they categorized time intervals as longer or shorter than a learned category boundary. We show that diverse firing dynamics exhibited by individual neurons allow for time to be robustly read out from the population. In addition, by comparing the categorization ability of neurons with categorization behavior of rats, we show that neural activity at interval offset from as few as 15 simultaneously recorded neurons was sufficient to explain behavioral performance. Continuous estimates of time decoded from striatal populations during interval stimuli ran faster or slower when rats judged an interval as longer or shorter, respectively. Lastly, by applying a classification analysis to high speed video frames taken during task performance we found that the categorization performance of neurons could not be explained by behavior that may unfold during the interval stimuli. These results demonstrate that striatal dynamics form a representation of time that is suitable to underlie perceptual report, and lend further support to "neural population clock" models of interval timing.

T-24. Human visual representations are predicted in space and time by convolutional neural networks

Radoslaw Cichy RMCICHY@MIT.EDU
Aditya Khosla KHOSLA@CSAIL.MIT.EDU
Dimitrios Pantazis PANTAZIS@MIT.EDU
Antonio Torralba TORRALBA@MIT.EDU
Aude Oliva OLIVA@MIT.EDU

Massachusetts Institute of Technology

The neural machinery underlying visual object recognition comprises a hierarchy of cortical regions in the ventral visual stream. The spatiotemporal dynamics of information flow in this hierarchy of regions is largely unknown. Here we tested the hypothesis that there is a correspondence between the spatiotemporal neural processes in the human brain and the layer hierarchy of a deep convolutional neural network. We presented 118 images of real-world objects to human participants (N=15) while we measured their brain activity with functional magnetic resonance imaging (fMRI) and magnetoencephalography (MEG). We trained an 8 layer (5 convolutional layers, 3 fully connected layers) convolutional neural network (CNN) to predict 683 object categories with 900K training images from the ImageNet dataset. We obtained layer-specific CNN responses to the same 118 images. To compare brain-imaging data with the CNN in a common framework, we used representational similarity analysis. The key idea is that if two conditions evoke similar patterns in brain imaging data, they should also evoke similar patterns in the computer model. We thus determined 'where' (fMRI) and 'when' (MEG) the CNNs predicted brain activity. We found a correspondence in hierarchy between cortical regions, processing time, and CNN layers. Low CNN layers predicted MEG activity early and high layers relatively later; low CNN layers predicted fMRI activity in early visual regions, and high layers in late visual regions. Surprisingly, the correspondence between CNN layer hierarchy and cortical regions held for the ventral and dorsal visual stream. Results were dependent on amount of training and training material. Our results show that CNNs are a promising formal model of human visual object recognition. Combined with fMRI and MEG, they provide an integrated spatiotemporal and algorithmically explicit view of the first few hundred milliseconds of object recognition.

T-25. Evidence that the ventral stream uses gradient coding to perform hierarchical inference

Elias Issa ISSA@MIT.EDU
Charles Cadieu C.CADIEU@GMAIL.COM
James DiCarlo DICARLO@MIT.EDU

Massachusetts Institute of Technology

Feedforward processing increases neural selectivity for objects across stages of the ventral visual hierarchy. It remains unclear, however, if any additional information is encoded in neural responses beyond what is available from bottom-up processing. Here, we recorded from the intermediate and output layers of the ventral visual stream corresponding to posterior and anterior inferior temporal cortex (IT) and show that, in posterior IT, face selectivity built up in the first feedforward step of processing is rapidly and completely reversed in the next 30 milliseconds of processing. In contrast, face selectivity in anterior IT did not undergo strong reversals. When we built models to explore existing hypotheses of the dynamics of ventral visual processing, we found that the observed neural dynamics were incompatible with signal propagation in a pure feedforward model or extensions implementing adaptation, lateral inhibition, or normalization. However, a simple model class that uses generative feedback connections displayed reversals of selectivity in its hidden layers. Importantly, these reversals corresponded to the overall magnitude of gradients, or update signals, for performing rapid inference across a hierarchy. We then show that a model using gradient coding provides a unified account of seemingly disparate neural phenomena in IT that had previously lacked a principled explanation. These phenomena include sublinear input integration, temporal sharpening of responses to more familiar images, and the mixed facilitatory and suppressive effects

of feedback. Without additional parameter modifications, our model accounted for these previous findings and our current results, unifying them under one coherent computational framework. Under this view, neurons more strongly encode the gradients and not the variables themselves in their firing rates which suggests that a cortical area combines bottom-up and top-down information to produce the signals required for efficient visual inference.

T-26. Efficient receptive field tiling in primate V1

Ian Nauhaus[1] NAUHAUS@AUSTIN.UTEXAS.EDU
Kristina Nielsen[2] NIELSEN@MAIL.MB.JHU.EDU
Edward Callaway[2] CALLAWAY@SALK.EDU

[1]University of Texas at Austin
[2]Salk Institute for Biological Studies

Each point of the retinal image is encoded by a local population of neurons in the primary visual cortex (V1), classically referred to as the “point-image”. Each point-image is a window onto V1’s feature maps that must encode all feature combinations to accurately represent the given point in visual space. In turn, two variables heavily impacting V1’s coding capacity are (1) the size of the point-image relative to the spatial periodicity of the feature maps, and (2) the alignment between the maps themselves. These properties have been investigated using electrodes and intrinsic signal imaging, but as we show here, two-photon imaging provides a different and more complete picture of the primate V1 micro-architecture. First, our data shows that the tiling of receptive fields is far more precise than previous electrode measurements, and that the V1 point-image is therefore much smaller than previous estimates (4x areal coverage). Precise retinotopy and a smaller point-image may confer benefits in read-out from higher areas, but requires greater efficiency in the joint-organization between maps to represent all features. Next, we used two-photon imaging to measure the joint-organization of three feature maps: orientation, ocularity, and spatial frequency. Previous studies showed that orientation map contours have a strict orthogonal alignment to those of ocularity and spatial frequency maps, thus improving coverage efficiency for these feature pairings. Here, we “close the loop” on how these three maps are related by showing that ocularity and spatial frequency maps have near-parallel contours, but unique spatial periods to maintain coverage. In addition, we found that binocular zones consistently overlap with high spatial frequency zones. In summary, primate V1 receptive fields have negligible scatter when the position of recorded neurons are accurately identified, yet feature maps overcome the constraints required of such a small point-image with a strict alignment strategy.

T-27. Addition by division: a recurrent circuit explains cortical odor response regulation by SOM cells

Fitz Sturgill FITZSTURGILL@GMAIL.COM
Paxon Frady PAXON.FRADY@GMAIL.COM
Jeffry Isaacson JISAACSON@UCSD.EDU

University of California, San Diego

Diverse types of local GABAergic interneurons shape the cortical representation of sensory information, but their individual contributions are not fully characterized. Here we show how somatostatin-expressing interneurons (SOM cells) contribute to odor coding in mouse olfactory or piriform cortex (PCx). By optogenetically identifying and suppressing SOM cells, we find that they regulate principal neuron output through a purely subtractive operation that is independent of odor identity or intensity. As SOM cells normally limit but do not abolish the spontaneous firing of PCx principal neurons, this operation enhances the salience of odor-evoked activity without changing cortical odor tuning. How do SOM cells uniformly suppress principal neuron firing rates in PCx even as they are recruited by odor stimuli? We find that SOM cells regulate principal cells (L2/3 pyramidal neurons) directly by providing sensory-evoked inhibition and indirectly by limiting the sensory-evoked recruitment of fast-spiking in-

terneurons (FS cells). Under our experimental conditions, the recruitment of FS cells via recurrent cortical circuits may compensate for the progressive loss of SOM cell-mediated inhibition across odor stimuli of increasing intensity. To test this idea, we built a simple, rectilinear rate model constrained by these observations. We find that in order to recapitulate the results of SOM cell inactivation on principal cell firing, we must include recurrent circuitry and a divisive mechanism that is disinhibited by SOM cell suppression. Taken together, these data provide the first description of how a single inhibitory cell type affects odor coding in PCx and argue that synaptic operations, patterns of neuronal recruitment by sensory stimuli and interplay with other cell types must all be considered to fully understand the computations performed by SOM cells and other cortical interneurons.

T-28. Correlative and causal evidence that attention improves communication between cortical areas

Douglas Ruff RUFFD@PITT.EDU
Marlene Cohen COHENM@PITT.EDU

University of Pittsburgh

Several recent studies have shown that in addition to affecting the firing rates of sensory neurons, attention decreases the extent to which fluctuations in response to repeated presentations of the same stimulus are shared between pairs of neurons in the same cortical area that have similar tuning. This decrease in so-called spike count correlations combined with attention-related improvements in the sensitivity of single neurons provides support for the hypothesis that attention improves perception by affecting the fidelity with which visual stimuli are encoded within a cortical area. However, attention has also been hypothesized to improve the communication of visual information between cortical areas. We tested the hypothesis that attention increases communication between areas on the timescale of behavioral trials using two independent and complementary approaches. First, we recorded simultaneously from populations of neurons in primary visual cortex (V1) and the middle temporal area (MT) using similar tasks and data analysis methods as those used to measure the effects of attention within an area. We found that in contrast to its effects on correlations within an area, attention increases correlations between pairs of neurons in different areas. Second, we made a causal manipulation to test the hypothesis that attention improves communication between areas by electrically stimulating V1 neurons during the attention task. We found that attention increases the extent to which manipulating V1 activity affects the activity of downstream neurons in MT. Together, our results provide evidence that attention acts on visual cortex in at least two ways: by affecting both the way visual stimuli are encoded within a cortical area and the extent to which visual information is communicated to downstream areas.

T-29. Stimulus driven inhibition by layer 6 underlies the neocortical processing of sensory change

Jakob Voigts[1] JVOIGTS@MIT.EDU
Christopher Deister[2] CHRISTOPHER_DEISTER@BROWN.EDU
Christopher Moore[2] CHRISTOPHER_MOORE@BROWN.EDU

[1]Massachusetts Institute of Technology
[2]Brown University

Neocortex learns predictive models of sensory input and detects and represents novel stimuli differently from expected or repeated stimuli. How this detection of stimulus deviations is implemented in the layered neocortical circuitry however, and how the novelty of a stimulus can be represented independently of the stimulus content is currently not understood. Here, using single-neuron recordings across neocortical layers and calcium imaging in awake mice, we find that layer 2/3 neurons encode heterogeneous, history dependent change signals, in contrast to layer 4 and 6 neurons that represent stimuli faithfully. We find evidence that layer 6 neurons play a key role in

this change representation through stimulus tuned inhibition. Layer 6 has long been hypothesized to play a central role in integrating bottom-up and top-down information because it is unique in receiving both direct thalamic, and long-range cortico-cortical inputs. It is sparsely but selectively sensory driven, and modulates sensory responses in superficial layers through inhibitory interneurons. Using calcium imaging in layer 6, we find that instead of selectively reacting to stimulus changes, layer 6 cells represent stimuli faithfully, similar to neurons in the main cortical input layer 4. This finding can explain why layer 2/3 neurons, which are driven by layer 4 and suppressed by layer 6 can represent stimulus changes without being constrained to representing the content of the most recent stimulus. Weak optogenetic stimulation of layer 6 suppresses superficial fast-spiking neurons and disrupts the heterogeneous change encoding in layer 2/3 neurons, causing them to linearly represent current stimuli, without changing overall firing rates. Subtle stimulus novelty also improves performance in a tactile detection task, and the optogenetic manipulation of layer 6 selectively removes this benefit. Our findings outline a mechanism for feature-independent change detection by normalization of neural responses to previous or expected stimuli through layer 6 mediated inhibition.

T-30. Interdigitated functional subnetworks in somatosensory cortex during active tactile behavior

Simon Peron[1] PERONS@JANELIA.HHMI.ORG
Freyja Olafsdottir[2] HAUFREY@GMAIL.COM
Jeremy Freeman[1] FREEMANJ11@JANELIA.HHMI.ORG
Karel Svoboda[1,3] SVOBODAK@JANELIA.HHMI.ORG

[1]Janelia Farm Research Campus
[2]University College London
[3]Howard Hughes Medical Institute

Pyramidal neurons in cortex exhibit functional diversity, even within a cell class and cortical area. Neurons with different behavior-related dynamics are often spatially intermingled and thus difficult to manipulate selectively. Consequently, little is known about the interactions between and within functionally defined cortical subnetworks, especially in the context of behavior. We employ multi-photon single cell ablation guided by volumetric two-photon calcium imaging to lesion neurons participating in specific sensory representations. Imaging was performed in mice performing a whisker-dependent object localization task. In L2/3 of vibrissal somatosensory cortex (vS1), ablating a handful ($<$50) of neurons encoding whisker-object contact (touch) reduces touch-related activity in the remaining, unablated members of the touch subnetwork. Neurons encoding whisker movements (whisking) are not affected. Ablation of inactive vS1 neurons produces no effect on either the touch or whisking subnetwork. This suggests that recurrent connections among members of the touch subnetwork amplify touch-related activity, enhancing the salience of the neural response to object contact. The lack of effect in the whisker movement subnetwork implies that these neurons do not depend on activity in the touch subnetwork.

T-31. Optogenetic investigation of dopamine D2 receptor signaling and loss-sensitivity in a model of problem gambling

Kelly Zalocusky KELLYZ@STANFORD.EDU
Thomas Davidson TJD@STANFORD.EDU
Charu Ramakrishnan CHARU.RAMAKRISHNAN@GMAIL.COM
Talia N Lerner TALIA.LERNER@GMAIL.COM
Brian Knutson KNUTSON@PSYCH.STANFORD.EDU
Karl Deisseroth DEISSERO@STANFORD.EDU

Stanford University

Impulse-control disorders, such as pathological gambling, can be conceptualized as a series of economic decisions, wherein the subject repeatedly chooses a risky but potentially large reward over a more certain outcome. Recent evidence implicates dopamine D2-receptor (D2R)-expressing neurons in these maladaptive decision processes. D2R agonist drugs increase incidence of problem gambling in clinical populations, and in the laboratory, these drugs decrease learning rates specifically in response to negative outcomes–an effect that potentially biases decisions in favor of impulsive reward-seeking. In a rodent model of risk preference, we find that behavior is well described by a logistic regression that accounts for the last 3 trial outcomes, indicating that even well trained rats are responding to trial-by-trial outcomes. Further, rats' long-term risk preference is largely explained by their acute response to losses. Risk-averse rats are more sensitive to losing outcomes than are risk-seeking rats. We find that a D2R agonist dose-dependently increases risk-seeking choices, both when delivered systemically and when delivered specifically in the nucleus accumbens (NAc), implicating the D2R-expressing NAc neurons in risk-preference. To further examine the behavioral role of D2R-expressing NAc cells, we expressed GCaMP6m in this population and recorded their activity using fiber photometry. We find that these cells differentiate reward size. These cells also predict rats' upcoming choices: their activity during the decision period is elevated if the rat is about to make a safe choice, as compared to a risky choice. By optogenetically manipulating the activity of these cells in a pseudo-random subset of trials, we find that ChR2-driven increases in activity during the decision period cause a significant decrease in risky choices on stimulated trials, suggesting that the D2R-expressing cells' activity during the decision period drives risk-preference on a trial-by-trial basis.

T-32. Characteristic dynamics of value coding in normalizing decision circuits

Kenway Louie[1] KLOUIE@CNS.NYU.EDU
Thomas LoFaro[2] TLOFARO@GUSTAVUS.EDU
Ryan Webb[3] RYAN.WEBB@UTORONTO.CA
Paul Glimcher[1] GLIMCHER@CNS.NYU.EDU

[1]New York University
[2]Gustavus Adolphus College
[3]University of Toronto

Normalization is a widespread neural computation, mediating divisive gain control in sensory processing and implementing a context-dependent value code in decision-related frontoparietal cortices. Although decision- making is a dynamic process with complex temporal characteristics, most models of normalization are time- independent and little is known about the dynamic interaction of normalization and choice. Here, we show that a simple dynamical systems model of normalization explains the characteristic phasic-sustained pattern of cortical decision activity and predicts specific normalization dynamics: value coding during initial transients, time-varying value modulation, and delayed onset of contextual information. Empirically, we observe these predicted dynamics in saccade-related neurons in monkey lateral intraparietal cortex. Furthermore, such models naturally incorporate a time-weighted average of past activity, implementing an intrinsic reference-dependence in value coding. Notably, value coding is strongest during initial rather than late stages of choice, suggesting that - in contrast to accumulator-based models of choice - economic decision-making may be more efficient at short timescales (a speed-accuracy complementarity). These results show that a recurrent model of divisive normalization captures both the dynamic activity patterns and equilibrium value coding exhibited by cortical decision circuits. In addition to elucidating potential network architectures for decision coding, dynamical circuit models such as this one can predict novel patterns of choice outside the scope of traditional behavioral analysis.

T-33. A multi-layered olfactory population code enhances odor detection speed

James Jeanne JAMES_JEANNE@HMS.HARVARD.EDU
Rachel Wilson RACHEL_WILSON@HMS.HARVARD.EDU

Harvard University

An animal's survival hinges on fast and accurate detection of sensory stimuli. However, individual neurons are noisy, which can impair both speed and accuracy. One mechanism to compensate for noise is to pool activity from many neurons that encode independent estimates of the sensory world. There are few opportunities to study this computation because of the scarcity of identifiable neural circuits with known connectivity. We used the Drosophila olfactory system to understand how multiple layers of neural circuitry process a population of noisy inputs. Using optogenetics, we drove spikes in a complete population of olfactory receptor neurons (ORNs) that converge on a common glomerulus. Each ORN synapses on all 6 of the antennal lobe projection neurons (PNs) in this glomerulus. Surprisingly, despite independent noise in ORNs, noise reduction in PNs was minimal. This sub-optimal performance results from a low PN spike threshold, which elevates spontaneous activity. However, the large ORN population reduced response latency in PNs. Due to their low threshold, PNs need to integrate only a handful of ORN spikes in order to spike themselves. By pooling from many ORNs—whose spike timing is independent—PNs can respond to those that, by chance, spike with the shortest latency. A low threshold means that PNs have high spontaneous firing rates. However, we found that driven spikes were more synchronized than spontaneous spikes between PNs. We identified a population of third-order olfactory neurons (in the lateral horn) that pool input from all 6 PNs. The rate of spontaneous activity was lower in lateral horn neurons than in PNs, with little sacrifice in driven response speed. Our results reveal a circuit that uses the power of pooling primarily for speed enhancement, rather than accuracy improvement; this circuit then suppresses spontaneous noise in a downstream layer of neurons that are selective for synchrony.

T-34. A common mechanism underlies changes of mind about decisions and confidence

Ronald van den Berg[1] RV284@CAM.AC.UK
Kavi Anandalingam[1] KAVI.ANANDALINGAM@GMAIL.COM
Ariel Zylberberg[2] AZ2368@COLUMBIA.EDU
Luke Woloszyn[2] LW2522@COLUMBIA.EDU
Roozbeh Kiani[3] ROOZBEH@NYU.EDU
Michael Shadlen[2] SHADLEN@COLUMBIA.EDU
Daniel Wolpert[1] WOLPERT@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Columbia University
[3]New York University

A decision is a commitment to a proposition or plan of action based on evidence, prior knowledge and expected costs. Here we examine the mechanism that underlies reaction times, choice and confidence, and why the latter two might undergo revision after initial expression. Subjects performed a reaction-time direction discrimination task, reaching to one of four targets simultaneously indicating direction choice and confidence level. Although the stimulus terminated on movement initiation, subjects occasionally changed their choice or confidence. Our hypothesis is that a common mechanism explains the rate of these occurrences, as well as the speed, accuracy and confidence of initial decisions. We modeled the decision process as a competition between two accumulators, one accruing evidence for right and another for left. The decision is rendered when one accumulator reaches a bound. Confidence is formalized as the probability of a correct decision (log-odds correct, LOC), given the state of two accumulators. For the initial decision, LOC is captured by the state of the losing accumulator and decision time, because the winning accumulator is at the bound. The confidence categories (low and high) are separated by a threshold on the LOC. After movement initiation, evidence in the processing pipeline (due to sensory and

motor delays) accumulates and LOC evolves. Depending on whether LOC crosses each of three bounds, the participant changes their confidence, their choice, both confidence and choice, or neither. This model provides an accurate fit to reaction times, choice and confidence of initial decisions and frequency of subsequent changes of choice and confidence as a function of task difficulty. Moreover, even on trials with no change-of-mind, the movement kinematics are predictive of confidence. Our model explains changes of choice and confidence as a continuation of the same process that underlies the initial choice and confidence.

T-35. Prefrontal and midbrain contributions to fast executive control in the rat

Chunyu A Duan[1] CDUAN@PRINCETON.EDU
Jeffrey Erlich[2] JERLICH@NYU.EDU
Charles Kopec[1] CKOPEC@PRINCETON.EDU
Carlos Brody[1,3] BRODY@PRINCETON.EDU

[1]Princeton University
[2]NYU Shanghai
[3]Howard Hughes Medical Institute

Rapid sensorimotor remapping in response to changing environmental demands is a hallmark of executive control. To study the underlying neural mechanisms, we developed an automated method to train rats in a behavior in which subjects were cued, on each trial, to either apply a sensorimotor association to orient towards a visual target ('Pro', or its reverse ('Anti', orient away). A 500-750 ms memory delay separated the task cue and the visual target. We found multiple behavioral asymmetries that suggest the Anti task is cognitively demanding while the Pro task is easier to learn and perform, consistent with current hypotheses that Anti requires prefrontal cortex (PFC) whereas Pro could be mediated by the midbrain superior colliculus (SC). Our inactivation and physiology data from the rat medial PFC confirmed its expected role in the Anti task. However, we provide three lines of evidence that challenge the hypothesis that SC is only required for stimulus-driven Pro responses and needs to be inhibited for correct Anti execution. First, bilateral pharmacological inactivation of SC substantially impaired Anti while leaving Pro essentially intact. Second, we conducted transient optogenetic inactivation of SC during either the task cue, memory delay or visual target periods, and found an Anti deficit only for delay period inactivation. These data reveal a surprising role for the SC in maintaining a memory of the Anti task set, separately from, and before any motor planning. Finally, consistent with the inactivation data, single-unit recordings in the SC showed non-spatial task encoding during the delay period, further supporting SC's contribution in maintaining a cognitively demanding task set. Taken together, our results suggest a distributed network underlying behavioral inhibition and flexible task switching, and argue against a strict segregation of cognitive functions in PFC versus motor functions in subcortical regions.

T-36. Single-trial motor and cortical variability are tightly and bi-directionally coupled

Marius Pachitariu[1] MARIUS10P@GMAIL.COM
Cora Ames[2] KCAMES@STANFORD.EDU
Byron Yu[3] BYRONYU@CMU.EDU
Gopal Santhanam[2] GOPAL@NERUR.COM
Stephen Ryu[2] SEOULMAN@STANFORD.EDU
Krishna Shenoy[2] SHENOY@STANFORD.EDU
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]Stanford University
[3]Carnegie Mellon University

Trial-to-trial variability in motor-cortical activity and instructed movements is often thought to reflect separate sources of noise. Here we show that in fact, for trained monkeys making directed reaches, fine variations in movement and in coordinated neural activity are tightly related. We first asked whether trial-to-trial fluctuations in population activity could predict small variations in hand trajectories over hundreds of reaches to the same target. Linearising the relationship about the corresponding means, we found we could predict motor variation with greater precision than typically reported for off-line decoders. Population activity preceding movements by 150ms accounted for about 60

T-37. Neural network model of 3D head-direction tuning in bats

Alon Rubin ALON.RUBIN@GMAIL.COM
Nachum Ulanovsky NACHUM.ULANOVSKY@WEIZMANN.AC.IL
Misha Tsodyks MTSODYKS@GMAIL.COM

Weizmann Institute of Science

Electrophysiological studies in Egyptian fruit bats demonstrate 3D head-direction neurons of several major types: pure azimuth cells, pure pitch cells, and cells tuned to a specific combination of head azimuth and pitch (conjunctive cells). Interestingly, the tuning of these cell types is consistent with a toroidal representation of head-direction. Here, we set out to investigate the theoretical conditions for the co-existence of these 3 types of cells in an interconnected neural network. We found that such a network can be modeled using the classical ring attractor and its extension to a toroidal attractor. This model allows mixed populations of neurons tuned to two cyclical variables — in our case, head azimuth and pitch. We solved the system both numerically and analytically, and showed that the coexistence of the conjunctive and the two types of pure representations is achieved over a broad range of parameters. Moreover, coexistence can occur even in cases of unequal population sizes — as has been observed in the experimental data. However, for certain areas of phase-space this coexistence collapses, and one population takes over — which illuminates the theoretical constraints on the co-existence of pure-variable and mixed-variable coding populations. Finally, our results suggest that manipulating the global balance of excitation/inhibition within a cortical brain circuit could fundamentally transform the qualitative tuning properties of neurons within that circuit. Our model can be generalized to additional brain regions that encode multiple cyclical variables, such as medial entorhinal cortex and primary visual cortex.

T-38. Zebra finch ventral tegmental area neurons encode song prediction error

Vikram Gadagkar VG55@CORNELL.EDU
Eliza Baird-Daniel ERB228@CORNELL.EDU
Alexander Farhang ARF66@CORNELL.EDU
Jesse Goldberg JESSEHGOLDBERG@GMAIL.COM

Cornell University

Many of our motor skills are not innately programmed, but are acquired by trial-and-error. It is well known that dopaminergic neurons in the ventral tegmental area (VTA) of mammals are activated by better-than-expected outcomes (positive prediction errors) and are suppressed by worse-then-expected outcomes (negative prediction errors) in a variety of reward-seeking tasks. Does the concept of dopaminergic prediction error apply to natural behaviors that lack external reinforcers such as food or juice? Songbirds provide a powerful model system because they have a discrete neural circuit, ‘the song system’, dedicated to a natural, learned motor sequence (song). Song learning requires motor exploration and auditory feedback, but it remains unknown how, and even if, vocal performance is evaluated during singing. To test if songbird VTA neurons carry conceptually similar error signals, we recorded (for the first time) from antidromically identified, basal ganglia (BG) - projecting VTA neurons (VTA-BG) in singing zebra finches during a protocol in which half the renditions of a targeted syllable were randomly distorted with broadband auditory feedback. VTA-BG neurons exhibited pauses after distorted syllables, consistent with a negative prediction error. These neurons also exhibited phasic bursts precisely time-locked to undistorted syllables, consistent with a positive prediction error. This suggests that an undistorted syllable was perceived as 'better than expected' because its expected value was diminished by the history of targeted distortions. Together, these findings (1) identify how the long sought-after auditory error signal reaches the song system; (2) reveal that songbirds can track the expected value of their ongoing vocalizations; and (3) demonstrate that principles of dopaminergic prediction error generalize to natural behaviors based on comparison of performance to an internal model.

T-39. Synapses represent and exploit estimates of uncertainty in their synaptic weight

Peter Latham PEL@GATSBY.UCL.AC.UK
Laurence Aitchison LAURENCE.AITCHISON@GMAIL.COM
Alexandre Pouget ALEXANDRE.POUGET@UNIGE.CH

Gatsby Computational Neuroscience Unit, UCL

Organisms face a hard problem: based on noisy sensory input, they must set a large number of synaptic weights. However, they do not receive enough information in their lifetime to learn the correct, or optimal weights (i.e. , the weights that ensure the circuit, system, and ultimately organism functions as effectively as possible). Instead, the best they could possibly do is compute a probability distribution over the optimal weights. Based on this observation, we hypothesize that synapses represent probability distribution over weights — in contrast to the widely held belief that they represent point estimates. From this hypothesis, we derive learning rules for supervised, reinforcement and unsupervised learning. This introduces a new feature: the more uncertain the brain is about the optimal weight of a synapse, the more plastic it is. This makes intuitive sense: if the uncertainty about a weight is large, new data should strongly influence its value, while if the uncertainty is small, little learning is needed. We also introduce a second hypothesis, which is that the more uncertainty there is about a synapse, the more variable it is. More concretely, the value of a synaptic weight at a given time is a sample from its probability distribution. The combination of these two hypotheses makes four predictions: 1) the more variable a synapse is, the more it should learn; 2) variability should increase with distance from the cell soma; 3) variability should increase as the presynaptic firing rate falls; and 4) PSP variance should be proportional to PSP mean. The first three predictions are a direct consequence of using a learning rule in which the learning rate is proportional to uncertainty. Predictions 2 and 4 are consistent with published data; we show data supporting prediction 1; and the third has yet to be tested.

T-40. Complex synapses as efficient memory systems

Marcus K Benna MKB2162@COLUMBIA.EDU
Stefano Fusi SF2237@COLUMBIA.EDU

Columbia University

The molecular machinery underlying memory consolidation at the level of synaptic connections is believed to employ a complex network of highly diverse biochemical processes that operate on a wide range of different timescales. An appropriate theoretical framework could help us identify their computational roles and understand how these intricate networks of interactions support synaptic memory formation and maintenance. Here we construct a broad class of synaptic models that can efficiently harness biological complexity to store and preserve

a huge number of memories, vastly outperforming other synaptic models of memory. The number of storable memories grows almost linearly with the number of synapses, which constitutes a substantial improvement over the square root scaling of previous models, especially when large neural systems are considered. This improvement is obtained without significantly reducing the initial memory strength, which still scales approximately like the square root of the number of synapses. This is achieved by combining together multiple dynamical processes that operate on different timescales, to ensure the memory strength decays as slowly as the inverse square root of the age of the corresponding synaptic modification. Memories are initially stored in fast varying variables and then progressively transferred to slower variables. Importantly, in our case the interactions between fast and slow variables are bidirectional, in contrast to the unidirectional cascades of previous models. The proposed models are robust to perturbations of parameters and can capture several properties of biological memories, which include delayed expression of synaptic potentiation and depression, synaptic metaplasticity, and spacing effects. We discuss predictions for the autocorrelation function of the synaptic efficacy that can be tested in plasticity experiments involving long sequences of synaptic modifications.

T-41. Selective rebalancing of cortical associations in humans via inhibitory plasticity

Helen Barron[1] HELEN.C.BARRON@GMAIL.COM
Tim Vogels[2] TIM.VOGELS@CNCB.OX.AC.UK
Uzay Emir[2] UZAY.EMIR@NDCN.OX.AC.UK
Tamar Makin[2] TAMARMAKIN@GMAIL.COM
Jacinta O’Shea[2] JACINTA.OSHEA@NDCN.OX.AC.UK
Ray Dolan[1] R.DOLAN@UCL.AC.UK
Timothy EJ Behrens[2] BEHRENS@FMRIB.OX.AC.UK

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

Synaptic input to cortical neurons is thought to be balanced such that excitatory and inhibitory (EI) currents are precisely matched and stable firing preserved. Changes in excitatory synaptic strength during learning and the acquisition of memory disturb this stable EI balance. It is not known how EI balance is restored after memory formation in humans, but a candidate mechanism implicates inhibitory synaptic plasticity in balancing excitatory connections with precisely mirrored inhibitory connections. Here we show evidence for such selective inhibitory plasticity in human cortex. We first developed an fMRI assay for indexing newly-formed Hebbian associations in human cortex, utilizing repetition suppression to measure the relative co-activation of neural representations supporting associated stimuli. Over time, we found a decrease in the measured strength of associations in cortex, but not in behaviour, a result consistent with the cortical circuit restoring EI balance via inhibitory plasticity whilst retaining the original memory. We tested our hypothesis by reducing cortical GABA levels via transcranial direct current stimulation to transiently modulate EI balance. By quantifying the change in GABA concentration using 7T MR spectroscopy, we found the cortical association reappeared in our fMRI assay in proportion to the GABA reduction. Our results support the idea of ‘antimemories’ in human cortex, i.e. proportional changes in inhibitory synaptic efficacies that balance the effects of potentiated excitatory synapses after memory formation.

T-42. Sparse distributed population codes of context support sequence disambiguation

Yan Wu YW331@CAM.AC.UK
Máté Lengyel M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

The storage and retrieval of past experience as sequences is important for humans and other animals. A main challenge in sequence learning is to discriminate specific sequences in the face of the ambiguity introduced by overlapping fragments shared by several sequences. Encoding contextual information together with the sequences of stimuli provides a computationally and biologically plausible way for disambiguation. Indeed, a number of cortical areas, most prominently the prefrontal cortex, the hippocampus, and the parahippocampal region, have been implicated in the encoding of contextual information. However, we lack a theoretical understanding of how contextual representations are extracted from experience in these cortical areas. Existing models of sequence disambiguation either take the representations of contextual information as given, or posit neural representations that are at odds with the distributed nature of cortical population codes. Here we present a recurrent neural network (RNN) model where contextual information is implicitly generated from the history of stimuli, without any explicit instruction, through optimising it for making sequential predictions. Although the RNN is only trained for one-step predictions, it supports cued recall of whole sequences without accumulating errors over subsequent items. The performances of both prediction and recall are robust under noise levels far beyond that applied during training. Our analysis shows this is because each memorised sequence is represented as a line of slow points (a quasi-line attractor) in the state space of the network. The resulting representation in the hidden layer of the RNN is sparse, and sequence disambiguation is achieved automatically in this sparse representational space. These results demonstrate that the sparse distributed representations of context found across the cortex provide optimal solutions for sequence disambiguation in neural networks.

T-43. Early cortical spontaneous activity provides a scaffold for constructing sensory representations

Matthias Kaschube[1] KASCHUBE@FIAS.UNI-FRANKFURT.DE
Bettina Hein[1] HEIN@FIAS.UNI-FRANKFURT.DE
Klaus Neuschwander[1] NEUSCHWANDER@FIAS.UNI-FRANKFURT.DE
David Whitney[2] DAVID.WHITNEY@MPFI.ORG
Gordon Smith[2] GORDON.SMITH@MPFI.ORG
David Fitzpatrick[2] DAVID.FITZPATRICK@MPFI.ORG

[1]Frankfurt Institute for Advanced Studies
[2]Max Planck Florida Institute for Neuroscience

The cortex is spontaneously active from the first moments that circuits form and there is ample evidence indicating that early cortical maturation relies on spontaneous activity. Yet, we know very little about how the pattern of cortical spontaneous activity prior to the onset of visual experience impacts circuit formation. Spontaneous patterns of activity that precede stimulus-evoked activity could provide a scaffold for the construction of sensory representations that are subsequently refined through sensory evoked activity. But, there is as yet no experimental evidence that spontaneous patterns of cortical activity prior to visual experience exhibit the spatial and temporal structure that is consistent with mature sensory evoked patterns. Here we took advantage of the robust columnar representation of orientation preference in visual cortex of the ferret to visualize patterns of spontaneous activity prior to the onset of visual experience and to determine how these patterns are related to stimulus evoked patterns in the same animal later in development. We found that there are robust columnar patterns of spontaneous activity that resemble the mature organization of the orientation preference map, several days prior to the time when an orientation preference map is evoked by visual stimulation. These observations were made by combining novel experimental techniques that employ the highly sensitive calcium indicator GCaMP6 to reveal population activity on a single trial basis in chronic recordings of the developing ferret visual cortex with novel data analysis approaches that uncover interpretable statistical relations from these data. We conclude that early spontaneous patterns of cortical activity exhibit an orderly columnar structure that forms the basis for building sensory evoked representations during cortical development.

I-1. Dorsal striatum encodes inferred position in a chain of behaviorally-utilized probabilistic events

Tanya Marton TANYA.MARTON@GMAIL.COM
Paul Worley PWORLEY@JHMI.EDU
Marshall Hussain Shuler SHULER@JHMI.EDU

Johns Hopkins University

The dorsal striatum (DS) is critical to learning reward contingencies and forming reward-accumulating policies. Primary sensory cortex also undergoes reward-dependent plasticity which may inform these policies. In particular, neurons in primary visual cortex (V1) come to represent the interval at which a visual stimulus predicts reward, exhibiting modulated activity from stimulus onset until the time of expected reward. To better understand how corticostriatal interactions contribute to reward-accumulating policies, we simultaneously recorded from V1 and DS in mice which enter a nosepoke, wait a random delay to probabilistically receive a stimulus and then lick to probabilistically receive a delayed reward. We find that while neurons' within V1 encode the time between visual stimulus and expected reward, a complete set of inter-event states (defined as positions between both expected and realized events) are encoded within neurons of the DS, plausibly through integration of cortical representations. Individual neurons' activity indicate a particular position within the chain of expected and realized events. DS neurons do more than subtend inter-event intervals. They also integrate temporal expectations with sensory and reward experience. For example, we find neurons that begin firing upon nosepoke entry and terminate firing upon stimulus delivery or after the mean time of stimulus delivery has elapsed, whichever occurs first. These neurons indicate the stimulus is neither expected nor received. Other neurons begin firing upon the mean time of stimulus delivery and terminate when the stimulus is delivered, indicating an expected, but unreceived stimulus. Together with many other state-encoding neurons, these neurons indicate the animals' position within a chain of probabilistic events. Upon the addition of a novel stimulus-reward contingency, some animals do not learn the optimal new policy, but retain learned relationships, despite many sessions. In these animals, a correlation of behavior and striatal response suggests that the DS only encodes policy-utilized positions.

I-2. Causal contribution and neural dynamics of the rat anterior striatum in an accumulation of evidence decision-making task

Michael Yartsev[1] MYARTSEV@GMAIL.COM
Timothy D Hanks[1] THANKS@PRINCETON.EDU
Carlos Brody[1,2] BRODY@PRINCETON.EDU

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

Accumulation of evidence is considered a fundamental computation accounting for many decision-making processes. Much attention has been devoted to cortical mechanisms underlying this process, with neural correlates of evidence accumulation found in both parietal and frontal cortices. The striatum serves as a gateway for cortical input into the basal ganglia and receives extensive projections from nearly all parts of cortex. Here, we used a combination of anatomical, pharmacological, optogenetic, electrophysiological and computational approaches to examine the role of the striatum in an auditory accumulation of evidence decision-making task. Anatomical tracing revealed a subregion of the rat anterior-striatum that receives convergent projections from the posterior parietal cortex (PPC) and frontal cortex (FOF) – both previously shown to encode variables related to accumulating evidence during this task. Using pharmacological and optogenetic inactivation, we showed that this striatal region is necessary for decisions driven by accumulation of evidence. Electrophysiological recordings revealed that many neurons exhibited strong pre-movement side-selective activity that gradually developed as the trial unfolded. Further, these firing rate modulations depended on the strength of accumulated evidence, and individual quanta of sensory evidence resulted in fixed firing rate increments that was sustained over time, as expected if striatal firing rates stably encode gradually accumulating evidence. Finally, we applied a recently developed computational method to examine the relationship between firing rates and the value of accumulated evidence and found that striatal neurons contained a graded representation of the accumulating evidence. Combined, using a wide range of experimental methods and analytical approaches, these data identify the anterior striatum as the first brain region shown to be both necessary for decisions driven by accumulating evidence, and to have the graded representation required to encode the value of the analog evidence accumulator. The anterior striatum may be the first known node of the evidence accumulation circuit.

I-3. Deciphering the neural representation of perceptual decisions with latent variable models

Kenneth Latimer[1] LATIMERK@UTEXAS.EDU
Jacob Yates[1] JLYATES@UTEXAS.EDU
Alex Huk[1] HUK@UTEXAS.EDU
Jonathan W Pillow[2] PILLOW@PRINCETON.EDU

[1]University of Texas at Austin
[2]Princeton University

Neural activity in the lateral intraparietal area (LIP) of macaque cortex has been hypothesized to represent an accumulation of evidence for decision-making. In particular, the ramp-like firing rates observed in LIP neurons have been interpreted as a neural correlate of the diffusion-to-bound model, which captures many of the aspects of decision-making behavior. An alternative hypothesis is that the spike rates of LIP neurons follow a discrete stepping process, in which the spike rate jumps stochastically between decision states on each trial. Attempts to differentiate between these hypotheses have relied on analyses of trial-averaged responses, which cannot reveal the dynamics of spike rates on individual trials. Here we address this problem within a Bayesian model comparison framework using latent dynamical models of LIP spike trains. We define explicit statistical models of spike responses in which the spike rate is governed by either a continuous diffusion-to-bound process or a discrete stepping process. We use Markov Chain Monte Carlo (MCMC) methods to fit these models to spike trains recorded from LIP neurons. These methods provide access to the full posterior distribution over parameters under each model, and allow us to infer bound-hitting or step times on each trial. In contrast to previous results, we find that the stepping model provides a better description than the diffusion-to-bound model for 31 out of 40 cells in a population of choice-selective LIP neurons.This indicates that a majority of neurons are better explained as stepping than as ramping during decision formation. Additionally, we extend our approach to model decision-related dynamics in multi-cell recordings, where traditional analyses have limited power to reveal shared representations of decision.

I-4. Linear dynamics of evidence integration in a contextual decision making task

Joana Soldado-Magraner[1] JMAGRANER@GATSBY.UCL.AC.UK
Valerio Mante[2] VALERIO@INI.UZH.CH
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University of Zurich

Circuits in PFC are believed to integrate sensory evidence and carry out context-dependent computations. In a recent study by Mante et al.[1], monkeys were trained to perform a dual-context decision-making task, learning to select a contextually-relevant stimulus while ignoring an irrelevant one. A non-linear RNN model was proposed, to explain how the circuit could be capable of integrating the relevant input while dynamically washing away the irrelevant one. The mechanism could be understood in terms of selection plus integration of input vectors

mediated by the left and right eigenvectors of the underlying linearized system. In the present study, we continue this work complementing it with a new analysis approach by inferring a linear dynamical system (LDS) model [2] directly from the data. The model-generated population trajectories captured the main features of the original trajectories. Furthermore, we were able to find a correspondence between the task-relevant dimensions found by regression in the real data[1], and the learned model parameters. Regression stimulus vectors were mostly aligned with the corresponding model input vectors. The regression decision axis, in turn, was closest to one of the right eigenvectors of the dynamical system with slowest dynamics. The complementary left eigenvector had the strongest projection coming from the contextually-relevant input. The particular dynamical solution that this model finds provides crucial support for the notion that changes in integration, rather than rebalanced input strengths, underlies the flexible behaviour. Specifically, it appears consistent with contextually-dependent alignment of inputs onto directions of dynamical persistence. Our analysis further suggests, that this persistent direction of integration is in fact one that provides an optimal readout of the system, as it is consistent with a dimension in state-space that better separates relevant inputs from irrelevant ones.

I-5. Bistable attractor dynamics explain the effects of rat PFC inactivation during decision making

Alex Piet[1] ALEXPIET@GMAIL.COM
Jeff Erlich[2] JERLICH@PRINCETON.EDU
Charles Kopec[1] CKOPEC@PRINCETON.EDU
Timothy D Hanks[1] THANKS@PRINCETON.EDU
Carlos Brody[1,3] BRODY@PRINCETON.EDU

[1]Princeton University
[2]NYU Shanghai
[3]Howard Hughes Medical Institute

A central goal of systems neuroscience is to understand how cognitive processes are implemented in the brain. Brain circuits can be directly manipulated through inactivation or stimulation. Explaining the effects of these perturbations on behavior is a key step for understanding the implementation of cognitive processes. Decision making is a well studied cognitive paradigm. A core process in decision making tasks is the gradual accumulation of noisy evidence. Unilateral inactivations of the rat Frontal Orienting Fields (FOF) during an auditory evidence accumulation task result in an ipsilateral bias. This bias is manifested as an unusual vertical scaling of the psychometric curve, which cannot be replicated by perturbations to drift-diffusion models that characterize evidence accumulation. However, a bistable attractor model easily recreates vertical scaling when used as a memory ‘read out’ of an accumulation process. This memory model encodes a binary decision variable, instead of an accumulating evidence variable. Vertical scaling is created by inducing biased flips in the binary encoding. This model makes a robust prediction that bias increases with the amount of time the FOF is an active memory process. However, a read out unit is only necessary at the end of the accumulation process, and should be active for the same amount of time regardless of stimulus duration. To clearly demonstrate time dependent bias, we turned to inactivation data from a memory guided orienting (MGO) task with an explicit memory period of variable duration. MGO data display a time dependent vertical scaling of the psychometric curve, consistent with the FOF maintaining a memory for an upcoming choice. A bistable attractor model explains normal and perturbed behavior from two different cognitive tasks, suggesting the FOF holds decisions about upcoming choices. By mapping cognitive processes to specific brain regions, more detailed questions about neural coding and dynamics can be investigated.

I-6. Reinforcement learning limits performance in categorical decision-making

Andre Mendonca ANDRE.MENDONCA@NEURO.FCHAMPALIMAUD.ORG
Maria Vicente MARIA.VICENTE@NEURO.FCHAMPALIMAUD.ORG
Alexandre Pouget ALEXANDRE.POUGET@UNIGE.CH
Zachary Mainen ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Neuroscience Programme

The source of uncertainty in standard integration-to-bound models of perceptual decision-making is rapid temporal fluctuations in sensory evidence. The possible contribution of variability in other processes, such as fluctuations in the weights use to map sensory evidence onto choice category, is less well understood. Here, we studied this issue in animals performing two decision tasks: categorization of binary odor mixtures and identification of odors at low concentrations. We found that, as problem difficulty was increased in the two tasks, for the same change in accuracy, the change in reaction time was substantially larger in detection than in categorization (41% vs. 13% increase). We first examined whether sensory uncertainty was sufficient to explain the differences between the two tasks. To do so we tested integration-to-bound models that have previously depicted the interplay between reaction time and performance. These standard models failed to capture performance in both tasks simultaneously. We hypothesized that reinforcement learning introduces trial-to-trial fluctuations in category boundary that limit performance in the categorization task. To test this hypothesis we expanded the standard sensory integration model to include a learning rule that changes the mapping from sensory to decision after a decision has been made. After fitting the performance data, this model predicted a specific magnitude and pattern of history-dependent choice biases. These predictions were in quantitative agreement with the data. We observed that learning stimulus-to-choice mapping was still occurring despite stable performance over sessions, implying that rats assume a non-stationary world while solving these tasks. This assumption induces more errors by amplifying stimulus noise, particularly for stimuli that are in the proximity of the 50/50 categorical boundary We therefore suggest that on-going reinforcement learning is a critical non- stochastic source of decision variability that can limit the benefit of evidence accumulation, favoring faster reaction times in some decisions.

I-7. Affective decision making: effects of arousal on a random dot motion task

Windy Torgerud[1] LYNCH174@UMN.EDU
Dominic Mussack[2] MUSS0080@UMN.EDU
Taraz Lee[3] TARAZ.LEE@PSYCH.UCSB.EDU
Giovanni Maffei[4] GIOVANNI.MAFFEI@UPF.EDU
Giuseppe Cotugno[5] GIUSEPPE.COTUGNO@KCL.AC.UK
Paul Schrater[2] SCHRATER@UMN.EDU

[1]University of Minnesota
[2]University of Minnesota, Twin Cities
[3]University of California, Santa Barbara
[4]Universitat Pompeu Fabra
[5]King's College London

The effects of emotion on decision making have not been fully characterized. In this pilot study we sought to computationally model the effects of arousal on a two alternative forced choice random dot motion coherence task using a Bayesian hierarchical drift diffusion model (HDDM). Subjects were exposed to trials consisting of one of three different levels of arousal condition immediately preceding performance on a random dot motion decision making task. Arousal's effect on decision making was modeled as a constant term modulating evidence accrual rate, decision threshold boundary, both, or neither. These four models were compared using hierarchical Bayesian estimation to determine which hypothesis best fit the data. Additionally, biometric data including galvanic skin response (GSR), heart rate (HR), eye tracking data, and facial emotional expression data was collected for

each participant. We found that the threshold only model most accurately characterizes our data. This implies that arousal influences decision making by decreasing participants' conservativeness in their responses, rather than by increasing the accrual of evidence via attentional allocation.

I-8. Categorical representations of decision variables within OFC

Alexander Vaughan AVAUGHAN@CSHL.EDU
Junya Hirokawa JHIROKAW@CSHL.EDU
Adam Kepecs KEPECS@CSHL.EDU

Cold Spring Harbor Laboratory

Orbitofrontal cortex (OFC) is strongly implicated in decision-making under uncertainty, especially when decisions require the evaluation of predicted outcomes. Neurophysiological studies have identified a number of outcome-related variables within OFC such as reward value, risk and confidence — as well as unified value signals. However, which of these representations dominates OFC, and how they relate to the behavioral role of OFC, remains controversial. Previous approaches to identifying cortical representations have relied on model-based regression and decoding analyses. These approaches are extremely powerful, but can suffer from confirmation bias. Alternative methods rely on dimensionality reduction, which can provide unbiased estimates of the major uncorrelated sources of population variance. However, these cannot distinguish between representations that are merely independent (ie., randomly mixed) from those that are categorically separable (ie., represented in distinct neuronal subpopulations). Critically, both approaches reinforce the assumption that cortical representations arise as random mixtures of variables that can only be understood by population-level decoders. We trained rats in a two-alternative forced-choice perceptual decision-making task with block-wise changes in reward size. Analysis of 495 vlOFC neurons recorded in this task reveals exemplar OFC neurons representing various decision variables, including reward size, confidence, and integrated value. PCA dimensionality reduction confirms the strong contribution of integrated value and other decision variables to OFC representations, although most neurons reflect a mixture of PCA components. However, cluster analysis shows that individual neurons do not simply carry a random mixture of these decision variables. Instead, non-euclidean clustering techniques reveal a set of distinct clusters, such that each neuron falls into one of 9 stereotyped representations. These clusters provide a model-free partition of OFC representations that nevertheless maps remarkably well onto the decision-variable models that arise from psychometric/neuroeconomic models of choice behavior. This modular architecture has significant implications for the representation and decoding of decision-variables within OFC.

I-9. Distinct neurobiological mechanisms of top-down attention

Thomas Luo TZLUO@UCHICAGO.EDU
John Maunsell MAUNSELL@UCHICAGO.EDU

University of Chicago

Attention is crucial for behavior and perception, but its neuronal basis remains poorly understood. Neuronal signals related to top-down visual attention are found in many different brain regions, and these widespread signals are generally assumed to participate in a unitary neural mechanism of attention. However, psychophysical studies show that the behavioral effects of attention can be separated into two distinct components: shifts in either the criterion or sensitivity of the subject. Depending on task demands, the behavioral difference between the attended and unattended visual locations could be a shift in criterion, sensitivity, or both. Single-neuron studies of attention have not examined this distinction. Here, we first show that a paradigm used by many neurophysiological experiments of attention conflates changes in criterion and sensitivity. Then, using a novel task to dissociate these two components, we asked whether attention-related neuronal modulations in monkey visual cortex are associated with selective changes in sensitivity, criterion, or both. We found that multiple aspects of attention-related neuronal

changes in area V4 all corresponded to behavioral changes in sensitivity but not criterion. This result suggests that neuronal signals in different brain regions are associated with separate components of attention. Attention is not a unitary process in the brain but instead consists of distinct neurobiological mechanisms.

I-10. Probing the causal role of area V4 neurons in attentional selection

Anirvan Nandy NANDY@SALK.EDU
Jonathan Nassi NASSI@SALK.EDU
John Reynolds REYNOLDS@SALK.EDU

Salk Institute for Biological Studies

Recent studies have shown that deployment of covert attention to a spatial location causes a significant decrease in correlated variability (common fluctuations) among neurons in area V4 whose receptive fields lie at the attended location (Cohen & Maunsell, 2009; Mitchell et al., 2009). This reduction is especially prominent in the low-frequency range ($<$ 10Hz). These studies have estimated, on theoretical grounds, that the reduction in common fluctuations accounts for a substantial fraction (80%) of the improvement in sensory processing that occurs when attention is directed toward a stimulus. However, the proposal that low frequency correlated variability does, in fact, impair perception, is purely hypothetical, and under some conditions decorrelation would be expected to impair encoding of sensory information (Averbeck et al., 2006). Here, we test this proposal directly by inducing low-frequency fluctuations in V4 via optogenetic stimulation, to see if this interferes with the animal's performance in an attention-demanding orientation-change detection task (Fig 1). On a subset of trials we stimulated neurons in area V4 using a recently developed approach to primate optogenetics (Fig 2) (Ruiz, Lustig, Nassi et al., 2013). We injected lentivirus carrying the CaMKII promoter to preferentially drive expression of the depolarizing opsin C1V1 (lenti-CaMKII-C1V1-ts- EYFP) in excitatory neurons in area V4. We find that low-frequency laser stimulation (5Hz sinusoidally modulated ramp) impairs the animal's ability to detect fine orientation changes (Fig 3A). Physiologically, this stimulation elevates correlations among pairs of neurons (Fig 3B) by biasing neuronal responses toward specific stimulation phases (Fig 3C), but without changes in mean firing rate (Fig 3D). These results demonstrate that correlated variability can impair perception, supporting the hypothesis that attention-dependent reductions in correlated variability contribute to improved perception of attended stimuli.

I-11. Hippocampal contributions to prefrontal decision making

Thomas Jahans-Price[1] TJ9733@BRISTOL.AC.UK
Rafal Bogacz[2] RAFAL.BOGACZ@NDCN.OX.AC.UK
Matt Jones[1] MATT.JONES@BRISTOL.AC.UK

[1]University of Bristol
[2]University of Oxford

We introduce a computational model describing rat behavior and the interactions of neural populations processing spatial and mnemonic information during a maze-based, decision-making task (Jones and Wilson, 2005). The model integrates sensory input and implements working memory to inform decisions at a choice point, reproducing rat behavioral data and predicting the occurrence of turn- and memory-dependent activity in neuronal networks subserving task performance. Analysis of medial prefrontal cortical (mPFC) neuron activity recorded from six adult rats during task performance showed a subset of mPFC neurons selective for current turn direction. Prefrontal turn-selective neurons displayed a ramping of activity on approach to the decision turn and turn-selectivity in mPFC was significantly reduced during error trials. These analyses complement data from neurophysiological recordings in non-human primates indicating that firing rates of cortical neurons correlate with integration of sensory evidence used to inform decision-making. In order link this cortical processing to input from subcortical structures such as the hippocampus (HPC), we apply Bayesian decoding methods (Zhang et al.,1998) to rat hippocampal recordings and estimate the position in space represented by population spiking activity. During moments of quiescence in decision making tasks, the decoded spatial representation has previous appeared predictive of future behaviour (Foster and Pfeiffer, 2013; Johnson and Redish, 2007), implicating the hippocampus in planning future actions. During this phenomenon, the animal's head often orients towards and switches between potential options, appearing to consider the alternatives in a process describes as vicarious trial and error (VTE). We use a spatial task based on Powell and Redish (2014) which reduces automated behaviour and increases instances of VTE. During these moments where the animal appears undecided and considering its options, we analyse information processing in prefrontal cortex, prospective hippocampal coding and the interactions between both structures.

I-12. Encoding of reward prediction errors by serotonin neurons revealed by bulk fluorescence recordings

Sara Matias[1] SARA.MATIAS@NEURO.FCHAMPALIMAUD.ORG
Eran Lottem[1] ERAN.LOTTEM@NEURO.FCHAMPALIMAUD.ORG
Guillaume P Dugue[2] GDUGUE@BIOLOGIE.ENS.FR
Zachary Mainen[3] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]Ecole Normale Superieur
[3]Champalimaud Neuroscience Programme

Serotonin (5-HT) has been proposed to carry a scalar signals related to predictions of reward and punishment, but this theory has not been rigorously tested due to the difficulty of recording 5-HT activity with conventional methods. To overcome this limitation we developed a method to stably and specifically monitor 5-HT neuron activity in behaving mice over several weeks. We expressed the genetically-encoded calcium indicator GCaMP6s and measured the total fluorescence population ('bulk') signal using an implanted optical fiber. To correct for movement artifacts, we implemented a robust method based on co-expression of the calcium-independent red fluorophore tdTomato. We applied this method in behaving mice using a classical conditioning paradigm in which four different odors were associated with outcomes of different valence. The same outcomes were also tested without predictive cues, allowing us to compare neuronal responses to expected vs. unexpected outcomes. To validate the recording method, we first expressed GCaMP6s/tdTomato in midbrain dopamine neurons (N = 4). As expected, we observed signals consistent with 'classical' reward prediction error coding. Next, we targeted 5-HT neurons in the dorsal raphe nucleus using the SERT-Cre mouse line (N = 10). Surprisingly, we found a very similar pattern of activity to dopamine neurons. 5-HT and dopamine neurons both responded to reward-predicting cues and to surprising rewards, but not to predicted rewards. These results indicate that, as theorized previously, DRN 5-HT and DA neurons jointly encode outcome predictions. But contrary to theory, they encode the coherent rather than opponent signals. These results suggest the need for new theoretical views on how neuromodulatory systems encode value predictions and how they interact to produce their behavioral effects. The recording methods deployed also present a simple and useful tool for further elucidation of the scalar encoding properties of neuromodulators and other genetically-definable neuronal populations.

I-13. Distinct neural processes for appetitive and informational rewards

Ethan Bromberg-Martin[1] NEUROETHAN@GMAIL.COM
Josh Merel[1] JSM2183@COLUMBIA.EDU
Tommy Blanchard[2] BLANCHARD.TOMMY@GMAIL.COM
Benjamin Hayden[2] BENHAYDEN@GMAIL.COM

[1]Columbia University
[2]University of Rochester

When we are uncertain about future rewards, we often have a strong preference to view informative sensory cues to learn what our future holds. This information seeking behavior is common in our everyday lives but is difficult to explain for current theories of reinforcement learning and decision making[1,2]. There are two major proposals to explain information seeking. According to two-process models the brain has two distinct motivational processes: one for seeking primary rewards (such as food and water) and one for seeking information about those rewards [3]. According to single-process models, however, these behaviors stem from a single process tied to primary rewards; information seeking is explained as a side-effect of suboptimal, nonlinear distortions in estimating the value of sensory cues [4,5]. Here we present neural and behavioral evidence for the two-process theory. We gave monkeys choices between gambles that varied in two features: water reward (e.g. the amount of water that could be won) and cue informativeness (whether the gamble provided visual cues that gave advance information about the outcome) [6,7]. Behaviorally, we show that monkeys assign high value to information, paying $>$20% of their water reward in exchange for informative cues. We show using theory and model-fitting that this behavior is well explained by two-process models. Neurally, we recorded from cells in the orbitofrontal cortex (OFC) which have been hypothesized to represent the subjective value of choice alternatives. Indeed, OFC neurons signaled the water reward and cue informativeness of each gamble. However, strikingly, OFC neurons had no statistical tendency to integrate these features to code subjective value. Instead they signaled water and informativeness in uncorrelated manners, resembling the distinct processes in two-process models. Thus, our data suggest that information-seeking is not merely a side-effect of seeking primary rewards; the brain appears to value information in its own right.

I-14. Towards a quantitative model of confidence: Testing the Bayesian confidence hypothesis

William Adler WILL.ADLER@NYU.EDU
Wei Ji Ma WEIJIMA@NYU.EDU
New York University

People intuitively know what it means to report their confidence in a decision. Despite the importance and long history of studying confidence, quantitative understanding of confidence is still poor. The Bayesian Confidence Hypothesis (BCH) states that confidence is a function solely of the posterior probability of the chosen option. This makes the specific prediction that combinations of stimulus features that produce the same posterior probability will also produce the same confidence report. We tested this hypothesis using a perceptual categorization task that incorporates aspects of naturalistic object recognition. Human subjects categorized stimuli as coming from two overlapping categories, reporting category choice and confidence simultaneously. The BCH makes a precise quantitative prediction for how observers take uncertainty into account, both in their choice and in their confidence report. We find that observers indeed incorporate trial-to-trial uncertainty. However, while the BCH provides a good fit, we found that a heuristic model that is qualitatively similar—but not Bayesian—outperforms the Bayesian model in formal model comparison for a majority of subjects. This suggests that observers perform an approximation to Bayesian inference when computing confidence ratings.

I-15. Trinary choices in a sequential integration paradigm

Santiago Herce Castanon SANTIAGO.HERCECASTANON@PSY.OX.AC.UK
Howard Chiu HOWARD.CHIU@WADH.OX.AC.UK
Konstantinos Tsetsos KONSTANTINOS.TSETSOS@PSY.OX.AC.UK
Christopher Summerfield CHRISTOPHER.SUMMERFIELD@PSY.OX.AC.UK
University of Oxford

According to rational choice theory, the relative value of two options should be unaffected by the value of a third option. However, it is known that humans violate this axiom of choice when making three-alternative economical and perceptual decisions. Recently, two computational models have been proposed to account for the influence of a third (distracter) option. These models focus respectively on how decision values are normalized during value encoding and comparison. Here we assessed how a low-value distracter influences choices between two higher-value options in a sequential decision-making paradigm. Participants viewed three streams of information (8 visual samples, comprising bars of variable length) and were asked to report which contained bars that were longer on average. Correct responses elicited feedback and a financial reward. Consistent with previous findings, as the value of the third option D increased, participants were less likely to choose the highest valued option HV over the second best, LV. However, we also observed that the momentary rank of each option was a key determinant of choices: samples where D was ranked the best, contributed most strongly to erroneous choices. This finding cannot be explained by previously proposed normalization models, or by a simple linear race model but is consistent with a model in which local winners capture attention and enjoy preferential integration. In other words, irrational choices may result from the limited capacity of information processing systems, and the tendency to attend selectively to the most valuable items. We recorded the pupil size of our participants during the execution of the task and found that it correlated positively with the value of the local maximum and negatively with the value of the local minimum. Together, these findings suggest that attention and capacity limits play an important role in provoking systematic violations of rational trinary choice.

I-16. Bridging single-trial dynamics of LFPs and unit activity via attractor models of decision-making

Laurence Hunt[1] LAURENCE.HUNT@UCL.AC.UK
Timothy EJ Behrens[2] BEHRENS@FMRIB.OX.AC.UK
Jonathan Wallis[3] WALLIS@BERKELEY.EDU
Steven Kennerley[1] S.KENNERLEY@UCL.AC.UK

[1]University College London
[2]University of Oxford
[3]University of California, Berkeley

The significance of prefrontal cortex for reward-guided choice is well known from both human imaging and animal neurophysiology studies. However, dialogue between human and animal research remains limited by difficulties in relating observations made across different techniques. A unified modelling framework may help reconcile these data. We have recently used attractor network models to demonstrate that varying decision values can elicit changes in local field potentials (LFP) dynamics, causing value correlates observable with human magnetoencephalography (Hunt et al., Nat Neurosci, 2012). Here, in light of this framework, we sought to relate simultaneously recorded LFP and single-unit data from prefrontal cortex of four macaque monkeys performing a cost-benefit decision task (Hosokawa et al., J Neurosci 2013). By performing principal component analysis of unfiltered LFPs timelocked to choice, components emerged that resembled the main choice-related LFP signature (PC1) and its temporal derivative (PC2). PC1 thus indexes LFP event-related amplitude, but crucially PC2 indexes its latency, reflecting the speed at which choice dynamics occurred on each individual trial. We found PC1 scores were correlated with overall value sum: the inputs of a decision process. PC2 scores, however, were correlated with chosen (but not unchosen) value: the outputs of a decision process. To relate LFP dynamics to single unit activity, we regressed single-trial PC1 and PC2 scores onto simultaneously recorded single-unit firing rates (including decision variables as nuisance coregressors). PC2, indexing the internal latency of each individual choice, predicted single-unit activity in nearly half of all recorded units — in particular those cells that showed a value-to-choice transformation. Similar results could be found by principal component decomposition of the attractor network model. This provides a novel bridge between LFP single-trial dynamics and simultaneously recorded single-unit activity. Perhaps more significantly, it allows us to relate value correlates in human neuroimaging studies to their cellular origins.

I-17. Detecting representation of the view expectation in mPFC and posterior parietal cortex

Yumi Shikauchi YUMI-S@SYS.I.KYOTO-U.AC.JP
Shin Ishii ISHII@I.KYOTO-U.AC.JP

Kyoto University

Humans use both external cues (e.g., scene views) and prior knowledge about the environment (e.g., cognitive map) to keep track of their position during spatial navigation. View expectation is an essential step to collate scene views with cognitive map. To determine how the brain performs view expectation during spatial navigation, we used an fMRI-based brain decoding methodology in which machine-learning models read out the contents of egocentric view expectation. Using a learned maze navigation task, we were able to decode subjects' expectations about their forthcoming scene view from fMRI signals in the posterior parietal and medial prefrontal cortices (mPFC), whereas activity patterns in the occipital cortex were found to represent various types of external cues. Moreover, the decoder's output reflected the subjects' expectations even when they were wrong, corresponding to subjective beliefs as opposed to objective reality. Furthermore, we were able to construct allocentric cognitive maps (i.e., the maze that participants "believed") using the egocentric view expectation decoded from the brain activities. View expectation is thus subjectively represented in the human brain in the identifiable form from fMRI voxels and fronto-parietal network is involved in integration of external cues and prior knowledge during spatial navigation.

I-18. Modulators of V4 population activity under attention

Neil Rabinowitz[1] NC.RABINOWITZ@GMAIL.COM
Robbe Goris[1,2] ROBBE.GORIS@NYU.EDU
Marlene Cohen[3] COHENM@PITT.EDU
Eero P. Simoncelli[1,2] EERO.SIMONCELLI@NYU.EDU

[1]New York University
[2]Howard Hughes Medical Institute
[3]University of Pittsburgh

Neural populations are affected by modulatory influences that fluctuate over time. Direct measurement of these influences is typically difficult or impossible, and one must resort to studying them indirectly, through their effects on individual neuronal responses. Visual attention is one such modulator. In macaque area V4, directed attention increases the firing rates of neurons, but also reduces their variability, as well as the noise correlations between them [Cohen & Maunsell, 2009]. We propose that this assortment of effects arises from a fluctuating, population-level modulation of response gain. We instantiate this hypothesis in a population response model, in which spike counts are drawn from Poisson processes, with instantaneous rates that are the product of a stimulus drive with slowly-varying private, and more rapidly-varying shared modulatory signals. We fit this model to spike trains from populations of ~100 V4 neurons, simultaneously recorded from both hemispheres while the animal performed a stimulus-detection task under directed attention [Cohen & Maunsell, 2009]. The model reveals two separate shared modulatory signals, each operating primarily within one hemisphere. For conditions in which the monkey was cued to attend to one hemifield, the corresponding modulator's mean is larger and its variance smaller. Attention thus increases and stabilizes population gain. These changes in shared modulation, in turn, explain the observed increases in individual neurons' firing rates, and the decreases in their variability and noise correlations. The magnitude of these changes for individual neurons is also predicted by the strength of their coupling to the shared modulators. Finally, the modulatory signals are correlated with the monkey's behavioral performance on each trial, and reflect the reward received on the previous trial. By exposing the detailed, time-varying behavior of internal signals that are coupled with behavioral observables, the model provides a parsimonious explanation for the effects of attention on neural responses.

I-19. Attention to items in working memory improves fidelity of population codes in human cortex

Thomas Sprague TSPRAGUE@UCSD.EDU
Edward Ester EESTER@UCSD.EDU
John Serences JSERENCES@UCSD.EDU

University of California, San Diego

Working memory (WM) is a core cognitive function that enables the maintenance of information no longer present in the environment for further computation and guidance of behavior. When macaques maintain a spatial position in WM, firing rates of prefrontal cortex neurons are sustained in a location-dependent manner (Funahashi et al, 1989). Though such single-neuron results are intriguing, to fully understand the population code for spatial information it is necessary to take into account all response properties within a brain region. We have recently developed a method that allows us to reconstruct images of stimuli viewed (Sprague & Serences, 2013) or held in visual WM (Sprague, Ester & Serences, 2014) by human observers using fMRI activation patterns in retinotopic occipital, parietal, and frontal visual regions of interest. We previously quantified changes in stimulus or WM representations within reconstructions and consistently observed changes in representation amplitude above baseline with changes in behavioral state (attention or WM load). Here, we sought to evaluate the effect of attention on reconstructed representations of stimuli held in spatial WM rather than representations of stimuli present on the screen. Participants precisely remembered the position of one or two dots on the screen. After 8 s, participants were cued to either continue maintaining one or both positions, or to focus attention on one of two remembered positions. After a second 8 s delay, participants recalled the position of one item held in WM. fMRI-based region-level reconstructions of the contents of WM revealed that directing attention to one of two items held in WM restored its representation to the fidelity of a single item held in WM, and behavioral performance also recovered. These results suggest that attention improves the information content of population codes for both physically-present and remembered visual stimuli in a similar manner.

I-20. Mechanisms underlying near-optimal evidence integration in parietal cortex

Hannah Tickle[1,2] H.TICKLE@UCL.AC.UK
Maarten Speekenbrink[1] M.SPEEKENBRINK@UCL.AC.UK
Konstantinos Tsetsos[2] KONSTANTINOS.TESTSOS@PSY.OX.AC.UK
Elizabeth Michael[2] ELIZABETH.MICHAEL@PSY.OX.AC.UK
Christopher Summerfield[2] CHRISTOPHER.SUMMERFIELD@PSY.OX.AC.UK

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

During sensorimotor behaviour, humans integrate evidence from multiple sensory modalities in an optimal manner, weighting each source according to its reliability (i.e. level of sensory noise). Paradoxically however, when evaluating hypothetical scenarios, humans are unable to use ‘higher order’ measures of reliability to produce optimal behaviour. For example, when estimating the gender balance in a group of individuals, people fail to given more credence to larger samples than smaller samples, instead relying simply on the ratio of evidence favouring either choice (Kahneman et al 1982; Griffin and Tversky 1992). Here, we asked whether humans weighted information by sample size during integration of approximate number, using a trial and error learning paradigm. In each trial, participants viewed a series of 8 samples of blue and pink dots representing handfuls of balls drawn from an urn. Their task was to judge whether the balls were drawn from an urn containing mainly blue or mainly pink balls (60:40 ratio). Using logistic regression to calculate the impact that each sample carried in the decision revealed that humans weighted samples by their reliability, giving more credence to larger than smaller sample sizes. Furthermore, EEG recording and subsequent time frequency analysis suggested that neural signals over the parietal cortex correlated with a computational model that approximated optimal behaviour by integrating the difference in the number of balls in each sample, providing a reliability-weighted signal. This neurobiologically plausible model provided a parsimonious account of behaviour and neural activity, and suggests that when making decisions from experience, humans integrate information optimally by adding up uncertainty-weighted inputs, as predicted by probabilistic neural encoding models.

I-21. Shifting stimulus for faster receptive fields estimation of ensembles of neurons

Daniela Pamplona DANIELA.PAMPLONA@INRIA.FR
Bruno Cessac BRUNO.CESSAC@INRIA.FR
Pierre Kornprobst PIERRE.KORNPROBST@INRIA.FR

INRIA, Sophia Antipolis

The spike triggered averaged (STA) technique has been widely used to estimate the receptive fields (RF) of sensory neurons [chichilnisky 2001]. Theoretically, it has been shown that when the neurons are stimulated with a white noise stimulus the STA is an unbiased estimator of the neuron RF (up to a multiplicative constant). The error decreases with the number of spikes at a rate proportional to the stimulus variance [Paninski 2003]. Experimentally, for visual neurons, the standard stimuli are checkerboards where block size is heuristically tuned. This raises difficulties when dealing with large neurons assemblies: When the block size is too small, neuron's response might be too weak, and when it is too large, one may miss RFs. Previously online updating the stimulus in the direction of larger stimulus-neural response correlation [Foldiak 2001] or mutual information [Machens 2002, Mackay 1992] has been proposed. However, these approaches can not be applied for an ensemble of cells recorded simultaneously since each neuron would update the stimulus in a different direction. We propose an improved checkerboard stimulus where blocks are shifted randomly in space at fixed time steps. Theoretically, we show that the STA remains an unbiased estimator of the RF. Additionally, we show two major properties of this new stimulus: (i) For a fixed block sized, RF spatial resolution is improved as a function of the number of possible shifts; (ii) Targeting a given RF spatial resolution, our method converges faster than the standard one. Numerically, we perform an exhaustive analysis of the performance of the approach based on simulated spiked trains from LNP cascades neurons varying RF sizes and positions. Results show global improvements in the RF representation even after short stimulation times. This makes this approach a promising solution to improve RF estimation of large ensemble of neurons.

I-22. Training brain machine interfaces without supervised data

Mohammad Gheshlaghi Azar MOHAMMAD.AZAR@NORTHWESTERN.EDU
Lee E Miller LM@NORTHWESTERN.EDU
Konrad P Kording KOERDING@GMAIL.COM

Northwestern University

Brain machine interfaces (BMIs) use neural recordings to decode user intent, allowing individuals with disabilities to interact with the outside world. Training a BMI decoder generally requires supervised data, where neural activities and user intents are simultaneously known. However such supervised datasets are often difficult or expensive to obtain. These problems are highlighted by the difficulties encountered when using observation based decoders. Ways of decoding without supervised data could facilitate BMI approaches. Here we introduce a procedure for training BMI which requires no supervised data. Everyday life behavior such as walking, jogging or feeding, produces movement distributions that are structured and non-isotropic. A mis-calibrated decoder will thus produce statistically detectable deviations in their predictions and minimizing such differences may be used to train the decoder. Our algorithm, Procrustes BMI, uses prior knowledge of movement distributions to adjust for deviations from the expected movement distribution. In a low-dimensional space found through PCA we search

for a linear transformation that makes the distribution of decoder outputs most similar to the prior distribution of movements (Procrustes problem). To solve the Procrustes problem, we look for a rotation/reflection transformation which minimizes the information divergence between the predicted and the prior distribution of movements. Our method replaces supervised data with geometry. We test the Procrustes BMI method offline using neural data from monkeys solving a center-out reaching task. Procrustes BMI makes accurate predictions of movements kinematics, performing almost as well as a supervised linear decoder. Our results suggest that Procrustes BMI can be effective without access to the supervised data when we have access to prior knowledge of movement distributions.

I-23. Deep learning in a time-encoded spiking network

Jonathan Tapson[1] J.TAPSON@UWS.EDU.AU
Saeed Afshar[1] S.AFSHAR@UWS.EDU.AU
Tara Julia Hamilton[1] T.HAMILTON@UWS.EDU.AU
Andre van Schaik[1,2] A.VANSCHAIK@UWS.EDU.AU

[1]University of Western Sydney
[2]The MARCS Institute

The utility of deep learning in conventional neural networks has been proven by recent successes in processing massive, complex image sets. Biological networks display many of the same features as deep networks, such as multiple layers of interconnected neurons, with unsupervised learning in the early stages of processing, and supervised learning at higher levels; but biological networks encode signals with spikes. O'Connor et al. (2013) have produced a spiking deep belief network using rate-encoded spikes, which is trained offline using continuous-valued signals; it is a translation of a trained continuous network into a rate-coded spiking network. Neftci et al. (2014) have translated the classic contrastive divergence algorithm into the spiking domain, using STDP in stochastic rate-encoded neurons, to produce a spiking deep belief network. We have developed a biologically plausible deep learning network schema that is intrinsically sensitive to individual spike times, and is able to recognize complex spatio-temporal spike patterns. It combines localized synaptic adaptation with dendritic integration to implement both unsupervised and supervised learning. The unsupervised learning uses adaption of post-synaptic potentials and spiking thresholds to learn features of input patterns. The supervised learning algorithm synthesizes the correct dendritic weights in extensive arbors to achieve classification. Both algorithms can use simple weight update rules based on local neural information - no information needs to be back-propagated from a neuron to a pre-synaptic neuron. We have tested this network on a spike-time-encoding of the MNIST handwritten digit dataset. Unsupervised learning was used to build local receptive field networks in which laterally-inhibiting neurons learnt local features of the image. The outputs of this unsupervised first layer were fed to a supervised layer of spiking neurons, which classified the digits. This network learns in an online mode and converges on accuracies comparable to conventional machine learning techniques.

I-24. Online sparse dictionary learning via matrix factorization in a Hebbian/anti-Hebbian network

Tao Hu[1] TAOHU@TEES.TAMUS.EDU
Cengiz Pehlevan[2,3] PEHLEVANC@JANELIA.HHMI.ORG
Dmitri Chklovskii[4] MITYA@SIMONSFOUNDATION.ORG

[1]Texas A&M University
[2]Janelia Farm Research Campus
[3]Howard Hughes Medical Institute
[4]Simons Center for Data Analysis

Olshausen and Field (OF) proposed that neural computations in the primary visual cortex (V1) can be partially modelled by sparse dictionary learning[1]. By minimizing the regularized representation error they derived an online algorithm, which learns Gabor-filter features from a natural image ensemble in agreement with physiological experiments. The OF algorithm can be mapped onto the dynamics and synaptic plasticity in a neural network[1]. However, the originally proposed reciprocally connected two-layer network is not supported by cortical anatomy. At the same time, in the single-layer network[2] the derived learning rule is non-local - the synaptic weight update depends on the activity of neurons other than just pre- and postsynaptic ones - and hence biologically implausible. Previously, local learning rules have been suggested [3,4] but have not been derived from a principled cost function. Here, to overcome this problem, we derive sparse dictionary learning from a novel cost function - a regularized error of the symmetric factorization of the input's sample covariance matrix. Our algorithm maps onto a neural network of the same architecture as OF but using only biologically plausible local learning rules. When trained on natural images our network learns Gabor filters and reproduces the correlation among synaptic weights hard-wired in the OF network. The steps of the derived online algorithm, as well as the output statistics, capture several salient physiological properties of biological neurons. Specifically, the local Hebbian/anti-Hebbian synaptic learning rules we derived are consistent with those previously abstracted from biological observations of synaptic plasticity. The learning rate in the synaptic weight update is inversely proportional to the cumulative activity of the postsynaptic neuron in agreement with observations of LTP magnitude decaying with age in an activity dependent manner[5]. The distribution of neuronal firing has a peak at zero and a heavy tail in agreement with physiological measurements[6].

I-25. Bayesian filtering, parallel hypotheses and uncertainty: a new, combined model for human learning

Marco Lehmann MARCO.LEHMANN@EPFL.CH
Alexander Aivazidis ALEXANDER.AIVAZIDIS@EPFL.CH
Mohammadjavad Faraji MOHAMMADJAVAD.FARAJI@EPFL.CH
Kerstin Preuschoff KERSTIN.PREUSCHOFF@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Human learning and decision making in unstable environments has been studied intensively in recent years. A number of algorithms has been proposed to model how human subjects react to change-points. In this study we compare six models (including Nassar2010, Wilson2013, Payzan-LeNestour2011, Adams/MacKay2007) quantitatively and qualitatively and propose a new, combined model. We are interested in two questions: a) how well do these algorithms adapt to abrupt changes in the environment and b) how well do the algorithms explain human behaviour? We conduct this discussion in a simple framework which highlights the conceptual strengths of each algorithm and allows a comparison of their characteristics. We then combine fundamental concepts into a new behavioural model. This simple algorithm solves the change-point problem with remarkable performance, fits well to experimental data, and has an intuitive interpretation. The task solved by algorithms and humans is taken from Wilson et al, 2013: numbers are drawn from a normal distribution. With hazard rate h, the mean can abruptly change. The goal is to sequentially estimate the underlying mean from noisy observations. While there is a Bayes optimal solution, humans are likely to use a less complex algorithm. Indeed, fitting and model selection shows that half of the subjects (14/28) are best fit by simple TD learning. However, the other subjects are best fit by Payzan-LeNestour (n=7) and our combined model (n=6). Both models implement a reduced Bayesian approach and keep track of estimation uncertainty. In addition, our model has two parallel run-lengths. By allowing the run-lengths to grow dynamically, we provide a simple mechanism for evidence accumulation. Our results corroborate previous findings and combine them into a novel algorithm. It integrates three cognitive processes: Bayesian filtering, parallel memory traces at long and short timescales, and attentional selection by taking into account unexpected uncertainty.

I-26. Modeling activity-dependent constraint-induced therapy by spike timing dependent plasticity (STDP)

Won Joon Sohn WONJSOHN@GMAIL.COM
Terence D Sanger TSANGER@USC.EDU
University of Southern California

Unilateral brain injury in the developing brain leads to the development of an abnormally strengthened ipsilateral corticospinal tract (CST) and a weakened contralateral CST, as can be seen in subjects with hemiplegic cerebral palsy. In feline model, Martin etc. demonstrated that ipsilateral control persists after the injury resolves unless constraint therapy is applied. Such impairment in motor control by early brain injury bears resemblance to amblyopia in that it involves inter-hemispheric synaptic-competition as a principal mechanism for the impairment. Previously, amblyopia has been explained within the framework of BCM theory, a rate-based synaptic modification theory. Here, we hypothesized that a realistic spiking neuron model implements the same clinical neurological phenomenon if spike-timing-dependent-plasticity (STDP) is simulated. In this study, we modeled the developmental process of contralateral and ipsilateral CST under unilateral brain inactivation and reverse-inactivation of the unaffected side as a constraint-induced therapy. To this purpose, we developed a simplified corticospinal system with four Izhikevich-type spiking neurons representing CST connections and four synapses in the neuronal connections. We used programmable Very-Large-Scale-Integrated (VLSI) hardware to test the long-term developmental outcome of the injury and intervention in extremely high speed. Our model demonstrates (1) the establishment of dominant ipsilateral connection under unilateral-inactivation, (2) the inability of weakened contralateral connection to re-establish spontaneously, and (3) the re-establishment of contralateral connection by reverse-inactivation, a constraint-induced therapy, all rooted on the mechanism of synaptic-competition. Although our model is a highly simplified and limited representation of CST pathways, it showed that STDP with realistic spiking neurons is sufficient to capture behaviors of inter-hemispheric synaptic-competition, a key mechanism in the development and treatment of motor disorders due to unilateral early brain injury. We believe the model could further be honed as a tool to develop customized intervention strategies for individual patients with early unilateral brain injury.

I-27. A two-layer neural architecture underlies descending control of limbs in Drosophila

Cynthia Hsu CTH12@DUKE.EDU
Katherine A Tschida KAL14@DUKE.EDU
Vikas Bhandawat VB37@DUKE.EDU
Duke University

Flexible, adaptive behaviors require integrating sensory inputs and internal state to produce well-timed, coordinated motor output. Descending neurons (DNs), which are located in the brain and project to the body, are the sole pathway through which the information the brain integrates reaches the body's motor circuits. In mammals, the interaction of DNs across tracts and vast number of DNs per tract impedes a comprehensive understanding of descending motor control. In contrast, we find Drosophila only have ~900 DNs. Despite this numerical simplicity, the distribution of DNs is similar to other animals; thus, flies are a uniquely suitable model for the study of descending motor control. We employ a multi-pronged approach to gain insight into descending motor control in Drosophila. To characterize how individual DNs vary in their anatomy, sensory responses, and relationship to movement, we established methods for in-vivo whole cell patch clamp recordings, while presenting the flies with stimuli from multiple modalities and tracking individual leg movements. To determine the role of large populations of DNs, we activate and inactivate DN clusters while tracking leg movements. To characterize the circuit downstream of the DNs, we use an ex-vivo prep to activate DNs and record the electrophysiological response of motor neurons in the thorax. We have characterized DNs from two clusters. DNs from one cluster are strongly tuned to sensory modality, have anatomy suggestive of direct connection to motor neurons, and are likely to play a role

in movement initiation. DNs in a second cluster are not tuned to sensory information, do not project to the motor neurons, and their spiking activity correlates with fast leg movements but does not precede movement initiation. These results suggest a model of motor control in which one DN cluster drives movement initiation while another cluster modulates the parameters of ongoing movement.

I-28. Neural generative models with stochastic synapses capture richer representations

Maruan Al-Shedivat[1] ALSHEDIVAT.MARUAN@GMAIL.COM
Emre Neftci[2] NEMRE@UCSD.EDU
Gert Cauwenberghs[2] GERT@UCSD.EDU

[1]KAUST
[2]University of California, San Diego

Stochasticity in synaptic vesicle release is one of the major sources of noise in the brain. While the concept of cellular neural noise gave rise to computational models of biological learning such as deep belief networks and algorithms such as spike-sampling, the functional implications of synaptic stochasticity on learning remain unascertained and are often limited to filtering, decorrelation, or regularization. In this work, we approach synaptic stochasticity from the perspective of representations learning showing that it can improve fast concept learning in real situations where labeled data is scarce. We study a two-layer neural network that implements a Boltzmann machine with probabilistic connections. Noisy synaptic strengths lead to a notion of stochastic ensembles of generative models. We demonstrate how such ensembles can be tuned using variational inference methods. Analytically marginalizing synaptic noise for Bernoulli and Gaussian cases, we further use stochastic optimization techniques based on Gibbs sampling for learning the synaptic distributions. The results have three interesting implications. First, our network does not use noise for mere regularization, as it is used by dropout techniques in artificial networks. Instead, the ensemble that results from synaptic stochasticity is fitted to the data and hence stores the information about its variability. Second, during the learning process only the strongest inhibitory synapses become very reliable while the rest remain unreliable. This relates our model to experimental cortical data. Finally, we demonstrate that knowledge represented in the stochastic ensembles learned in an unsupervised way can leverage the performance of the subsequent classification and enable one-shot learning—the ability to learn categories from a handful of examples. We hypothesize that synaptic stochasticity in the brain may encode large amounts of observed data in such stochastic ensembles and further use them for fast learning from limited discriminative information.

I-29. Inferring structured connectivity from spike trains under Negative-Binomial Generalized Linear Model

Scott Linderman[1] SLINDERMAN@SEAS.HARVARD.EDU
Ryan Adams[1] RPA@SEAS.HARVARD.EDU
Jonathan W Pillow[2] PILLOW@PRINCETON.EDU

[1]Harvard University
[2]Princeton University

The steady expansion of neural recording capability provides exciting opportunities for discovering unexpected patterns and gaining new insights into neural computation. Realizing these gains requires flexible and accurate yet tractable statistical methods for extracting structure from large-scale neural recordings. Here we present a model for simultaneously recorded multi-neuron spike trains with negative binomial spiking and structured patterns of functional coupling between neurons. We use a generalized linear model (GLM) with negative-binomial observations to describe spike trains, which provides a flexible model for over-dispersed spike counts (i.e., responses

with greater-than-Poisson variability), and introduce flexible priors over functional coupling kernels derived from sparse random network models. The coupling kernels capture dependencies between neurons by allowing spiking activity in each neuron to influence future spiking activity in its neighbors. However, these dependencies tend to be sparse, and to have additional structure that is not exploited by standard (e.g., group lasso) regularization methods. For example, neurons may belong to different classes, as is often found in the retina, or they may be characterized by a small number of features, such as a preferred stimulus selectivity. These latent variables lend interpretability to otherwise incomprehensible data. To incorporate these concepts, we decompose the coupling kernels with a weighted network, and leverage latent variable models like the Erdős-Renyi model, stochastic block model, and the latent feature model as priors over the interactions. To perform inference, we exploit a recent innovations in negative binomial regression to perform efficient, fully-Bayesian sampling of the posterior distribution over parameters given the data. This provides access to the full posterior distribution over connectivity, and allows underlying network variables to be inferred alongside the low-dimensional latent variables of each neuron. We apply the model to neural data from primate retina and show that it recovers interpretable patterns of interaction between different cell types.

I-30. Sampling in a hierarchical model of images reproduces top-down effects in visual perception

Mihály Bányai BANYAI.MIHALY@WIGNER.MTA.HU
Gergö Orbán ORBAN.GERGO@WIGNER.MTA.HU

CSNL, Wigner RCP

Sensory perception can be understood as inference of latent causes underlying stimuli. This requires animals to learn the statistics of the environment in terms of a generative model with hierarchically organized latent variables corresponding to features of increasing complexity, which is characteristic of efficient probabilistic models (Hinton, 2006). While layers of probabilistic image models map intuitively to the hierarchical structure of the ventral stream, details of the correspondence have remained elusive. Mean activity at the lowest level of visual hierarchy, V1 simple cells, are well covered by independent linear filters adapted to the statistics of natural scenes (Olshausen, 1996) further elaborated by introducing specific dependencies between filters (Schwartz, 2001). Effects of top-down interactions on mean responses have been formalised in terms of covariance components of latent variables (Karklin, 2008). However, understanding the full response statistics of V1 neurons requires establishing computational principles that link efficient computations with higher-order statistics. A computationally appealing and neurally feasible method to perform Bayesian inference is the construction of stochastic samples (Lee, 2003). Sampling establishes a link between the neural response distribution and probability distributions in Bayesian models. We explore sampling in a hierarchical model of vision and its consequences on neural response statistics. A hierarchical model of the visual system was built that bears similarities with Karklin's covariance model, and used sampling to perform Bayesian inference. Activity of model neurons were regarded as samples from the posterior distribution resulting from inference. Top-down effects in the visual hierarchy are particularly noticeable in the phenomenon of illusory contours, thus we used synthetic images with the model to predict the neural response to such stimuli. The sampling scenario predicts variance and covariance changes in V1 and reproduces magnitude modulation and temporal evolution of neural responses to real and illusory contours in V1 and V2.

I-31. PrAGMATiC: a probabilistic and generative model of areas tiling the cortex

Alexander Huth ALEX.HUTH@BERKELEY.EDU
Wendy de Heer DEHEER@BERKELEY.EDU
Thomas Griffiths TOM_GRIFFITHS@BERKELEY.EDU
Frederic Theunissen THEUNISSEN@BERKELEY.EDU
Jack Gallant GALLANT@BERKELEY.EDU

University of California, Berkeley

In the best-studied regions of the human brain, functional representations appear to be organized into topographic cortical maps: retinotopic and semantic maps in visual cortex, tonotopic maps in auditory cortex, somatotopic maps in somato-motor cortex, and numerosity maps in parietal cortex. It seems likely that other regions of the brain—such as prefrontal cortex—are also organized into cortical maps, but those maps have yet to be discovered. One major obstacle for defining cortical maps is that both anatomy and functional representation can differ substantially between individuals, thwarting current methods for combining data across subjects. We have overcome this obstacle by developing PrAGMATiC: a probabilistic and generative model of areas tiling the cortex. This hierarchical Bayesian modeling framework accounts for individual differences in anatomy and function by treating the cortical map observed in each subject as a sample from a single underlying probability distribution. This distribution consists of two parts: an arrangement model and an emission model. The arrangement model uses a novel spring-based approach to describe cortical topography. This flexible approach can account for substantial individual differences in the shape, size, and anatomical location of functional brain areas. The emission model can use a number of different methods to generate functional cortical maps based on the map arrangement. Comparing these emission methods enables us to quantitatively test hypotheses about the principles of map organization, such as whether a cortical map is more likely composed of distinct functional areas or continuous gradients. Here we show that the PrAGMATiC modeling framework provides a compact and faithful characterization of cortical maps in prefrontal and parietal cortex based on voxel-wise models of semantic selectivity from an fMRI language experiment.

I-32. 3D motion sensitivity in a binocular model for motion integration in MT neurons

Pamela Baker PMBAKER@UW.EDU
Wyeth Bair WYETH0@UW.EDU

University of Washington

A recurring theme in neurophysiology is that while neurons may be characterized by, and organized according to, selectivity for multiple submodalities concurrently, attributes are often studied in isolation. For example, motion and binocular disparity sensitivities coexist in the early motion pathway, but a unified circuit model integrating these signals remains elusive. Such a model is important for studying the visual perception of motion in depth (MID), which involves both fronto-parallel (FP) visual motion and binocular signal integration. Two very recent studies (Czuba et al. 2014, Sanada and DeAngelis 2014) now show that many MT neurons are MID sensitive, contrary to the prevailing view (Maunsell and van Essen, 1983). These novel data are ideal for constraining models of binocular motion integration in MT. We have built binocular models of MT neurons to show how MID sensitivity can arise via inter-ocular velocity differences (IOVDs). Our modeling framework encompasses features common to established monocular MT models and extends the model of Rust et al. (2006) to be image-computable. We found that the characteristic differences between pattern and component computations make different predictions for MID tuning: FP motion-tuned neurons were better represented by the component cell model, whereas MID tuned cells could only be reconciled with an IOVD pattern cell with binocularly imbalanced input. By implementing binocular mixing in V1, we were able to generate robust IOVD-based MID tuning even without strictly monocular signals. Interestingly, our IOVD models predict a characteristic change in direction tuning with dichoptic plaids,

with the preferred direction shifted by 90o with dichoptic presentation. Overall, our models integrate motion and binocular processing to explain recent novel findings, make testable predictions relating disparate forms of motion sensitivity, and provide a foundation for building a unified binocular disparity-motion model of the dorsal stream.

I-33. Phoneme perception as bayesian inference with a narrow-tuned multi-modal prior

Daeseob Lim DAESEOB@GMAIL.COM
Seng Bum Yoo SBYOO@SNU.AC.KR
Sang-Hun Lee VISIONSL@SNU.AC.KR
Seoul National University

Sensory perception can be viewed as statistical inference. According to this view, observers estimate true events in the world based on newly acquired bits of sensory information and other previously formed beliefs about the world. Intriguingly, recent studies demonstrated that a distribution of prior could be revealed by fitting a Bayesian model to perceptual bias and variability exhibited by observers (Girshick et al., 2011). Here we attempted to infer a prior belief utilized by human observers engaged in phoneme perception, where the presence of a strong prior is expected. Observers performed classification and discrimination tasks on acoustic stimuli that varied gradually along the spectrum encompassing three stop consonant phonemes — /ba/, /da/, and /ga/. We modeled a prior as a mixture of three Gaussian distributions, the means and variances of which reflect the centers and spreads of phoneme stimuli encountered by listeners in the past. The likelihood was modeled as a Gaussian distribution, with a mean matched to a probed stimulus and a variance equal for all stimuli. With only a few parameters set to be free, our model successfully accounted for the two hallmarks of phoneme perception simultaneously, 'categorical perception' in the classification task and 'enhanced discriminability around phoneme boundaries' in the discrimination task. In goodness of fit, our model excelled a model with a uniform prior and paralleled a model with a free-form prior. Furthermore, our model, with all the model parameters inherited, predicted dramatic changes in phoneme classification after stimulus adaptation, large repulsive shifts away from adapting stimuli. Our Bayesian model offers a parsimonious yet rich explanation for phoneme perception by revealing a multimodal prior narrow-tuned to prototype phoneme stimuli. A simulation exercise demonstrates that our brain may implement this multimodal prior using probabilistic population codes based on basis functions varying in tuning width across phoneme stimuli.

I-34. Discovering switching autoregressive dynamics in neural spike train recordings

Matthew Johnson MATTJJ@CSAIL.MIT.EDU
Scott Linderman SLINDERMAN@SEAS.HARVARD.EDU
Sandeep Datta SRDATTA@HMS.HARVARD.EDU
Ryan Adams RPA@SEAS.HARVARD.EDU
Harvard University

Generalized linear models (GLM) are powerful tools for identifying dependence in spiking populations of neurons, both over time and within the population. The GLM identifies these dependencies by modeling spiking patterns through a linear regression and an appropriately-selected link function and likelihood. This regression setup is appealing for its simplicity, the wide variety of available priors, the potential for interpretability, and its computational efficiency. However, the GLM suffers from at least three notable deficiencies. First, the model is linear up to the link function, which only allows a limited range of response maps from neural spiking histories. Second, the model's parameters are fixed over time, while neural responses may vary due to processes that are exogenous to the population. Third, the generalized linear model presupposes a characteristic time scale for all dynamics,

when there may be multiple, varying time scales of neural activity in a given population. Here we seek to address these deficiencies via a switching variant of the generalized linear model. A switching system is one that evolves through a set of discrete states over time, with each state exhibiting its own low-level dynamics. For example, the latent state of a hidden Markov model (HMM) can be used to determine the parameters of an autoregressive (AR) process. These HMM-AR models can be used to identify common patterns of linear dependence that vary over time. Bayesian nonparametric versions of HMM-AR models extend these ideas to allow for an infinite number of such patterns to exist a priori, and semi-Markov variants allow the different states to have idiosyncratic duration distributions. Here we develop GLM variants of these switching AR processes and specialize them for neural spiking data. In particular, we exploit recent data augmentation schemes for negative binomial likelihood functions to make inference tractable in HDP-HSMM-AR models with count-based observations.

I-35. Advancing models of shape representation for mid-level vision

Dina Popovkina DPOPOVKINA@GMAIL.COM
Anitha Pasupathy PASUPAT@UW.EDU
Wyeth Bair WYETH.BAIR@GMAIL.COM

University of Washington

Theories about object recognition in the ventral visual pathway differ regarding the relative importance of boundary vs. surface properties in building an effective visual representation. To address this, we manipulated boundary and surface properties of simple shapes to examine their influence on the representation in cortical area V4 in a combined electrophysiology and modeling study. The current best model for V4 shape selectivity is a hierarchical-Max() model that achieves translation invariance and boundary contour tuning similar to that reported in V4. We compared the ability of this model and V4 neurons to maintain their tuning across a large battery of simple shapes that were presented either in filled or outline form. We found that the model, trained only on filled shapes, maintained robust tuning for outlines. Surprisingly, we also found that most V4 neurons did not maintain their shape tuning (as measured with filled shapes) for outlines, with many revealing a complete collapse of tuning (low firing rates to all outline shapes). We retrained the model on shapes and outlines together, and found the model was unable to simultaneously be selective for filled shapes and unselective for their outlines. A spectral analysis of the stimuli revealed that power is increased at higher frequencies for outlines while the underlying orientation structure is maintained. This indicated that the model is emphasizing orientation patterns that characterize the boundary while discarding phase information relevant to the surface fill. We used these insights to restructure the model, allowing shape-tuned V4 subunits to access phase-dependent inputs from early in the hierarchy, and found that this new architecture could explain the data. Our results suggest that boundary orientation and surface information are both maintained until at least the mid-level visual representation. We hypothesize that this arrangement is crucial for image segmentation.

I-36. Data-driven mean-field networks of aEIF neurons and their application to computational psychiatry

Loreen Hertog[1] LOREEN.HERTAEG@ZI-MANNHEIM.DE
Nicolas Brunel[2] NBRUNEL@GALTON.UCHICAGO.EDU
Daniel Durstewitz[3] DANIEL.DURSTEWITZ@ZI-MANNHEIM.DE

[1]BCCN Heidelberg-Mannheim
[2]University of Chicago
[3]Central Institute of Mental Health

Within recent years, there has been an increasing demand for biologically highly realistic and large-scale neural network studies, especially within the newly developing area of computational psychiatry. However, inclusion of a lot of physiological and anatomical detail into network models makes parameter estimation and simulation slow and tedious, and limits a deep understanding of the system's dynamical mechanisms. Thus, reduced but still physiologically valid mathematical model frameworks that ease the problem of fast parameter estimation, and allow for analytical access to the system dynamic would be highly desirable. Based on our previous work (Hertog et al. 2012, 2014), we present a data-driven, mathematical modelling approach which can be used to efficiently study the effect of changes in single cell and network parameters on the steady-state dynamics. For this purpose, we had derived two approximations for the steady-state firing rate of the adaptive exponential integrate-and-fire neuron (aEIF): In the first approach, we combine the 1-dimensional Fokker-Planck (FP) solution of the EIF with distributional assumptions for the adaptation current. The second method is based on solving perturbatively the 2-dimensional FP equation describing the aEIF with noisy inputs in the long adaptation time constant limit. Theoretical f-I curves are compared to single neuron simulations for a large number of parameter settings derived from in-vitro electrophysiological recordings of rodent prefrontal cortex (PFC) neurons probed with a wide range of inputs. A mean-field network model of the PFC is then developed which captures a number of physiological properties like different pyramidal and interneuronal cell types, short-term synaptic plasticity and conductance-based synapses with realistic time constants. This mean-field-based model is then used to compare dynamics in "healthy wild-type" networks to networks in which neuronal and synaptic parameters have been estimated from slice recordings of psychiatrically relevant genetic animal models or pharmacological manipulations.

I-37. Spatial sound-source segregation using a physiologically inspired network model

Junzi Dong JUNZID@BU.EDU
H. Steven Colburn COLBURN@BU.EDU
Kamal Sen KAMALSEN@BU.EDU

Boston University

The neural mechanisms underlying the "cocktail party problem" remain poorly understood. A key step in solving this problem is the ability to segregate sound sources from a mixture of sounds. Here, we present a computational network model of spatial sound-source segregation based on physiological responses from cortical-level neurons in songbirds. The goals of this study are two-fold: 1) to provide a physiological explanation for how the cortical level integrates and remodels spatial information from peripheral inputs to segregate a single sound-source from multiple sources, and; 2) to use the model to construct an engineering solution for segregating single sounds-sources from multi-source environments. For the first goal, we show that the responses of the network model match a population of recorded neurons, and propose physiological experiments that can be performed to test the model configuration. For the second goal, we propose networks for an engineering solution e.g., a beamformer. We are in the process of constructing an integrated engineering solution by adding two additional processing modules. The first is a peripheral input stage, which extracts spatial cues from binaural acoustic inputs to provide the neural inputs to the network model. The last module is stimulus reconstruction, whereby the network model output representing a segregated single-source in neural code is mapped back to acoustic waveforms that human subjects can listen to and comprehend. The integrated engineering solution may be useful for hearing assistive devices, which experience difficulty with the cocktail party problem.

I-38. A spatially extended rate model of the primary visual cortex

Dina Obeid DNO2103@COLUMBIA.EDU
Kenneth Miller KEN@NEUROTHEORY.COLUMBIA.EDU

Columbia University

We present a model that shows how the primary visual cortex (V1) of animals with orientation maps integrates

multiple inputs. Our work builds on the idea of a Stabilized Supralinear Network (SSN) which was previously studied by D. Rubin, S. Van Hooser, K.D. Miller (Neuron 2014, in press) and Y. Ahmadian, D. Rubin, K.D. Miller (Neural Computation, 2013) and proposed as a unified circuit motif for multi-input integration in the sensory cortex. This model produces a transition from supralinear to sublinear summation with increasing stimulus strength that, in the context of simple connectivity assumptions, reproduces surround suppression, "normalization" and their dependencies on stimulus strengths as well as multiple other phenomena. Our previous studies focused on the steady state behavior of 0d and 1d SSN models. Rubin et al.~(2014) also studied a 2d model, however, it was in a different parameter regime than that which theoretical study predicts should show the strongest sublinear behavior for stronger inputs, and we have found that it failed to show the decrease in inhibition observed during surround suppression (Ozeki et al, Neuron 2009). Here we study a 2d SSN rate-neuron network in the strongly sublinear regime of the SSN, with V1-like functional and spatial connection probabilities. This network produces the observed decrease in inhibition during surround suppression, while continuing to reproduce multiple steady-state V1 behaviors as before. Having found a more appropriate parameter regime for the 2d model, we are currently moving beyond steady state behavior to study response dynamics in this network. Surround suppression, normalization, and many related phenomena of multi-input integration are seen across multiple cortical areas. Understanding the underlying circuit is crucial to understanding fundamental brain computations. To do so, it is key to characterize model behavior in a regime well matched to cortical behavior, as we do here.

I-39. Internal models for interpreting neural population activity during sensorimotor control

Matthew Golub MGOLUB@CMU.EDU
Byron Yu BYRONYU@CMU.EDU
Steven Chase SCHASE@CMU.EDU

Carnegie Mellon University

To successfully guide limb movements, the brain takes in sensory information about the limb, internally tracks the state of the limb, and produces appropriate motor commands. It is widely believed that this process uses an internal model, which describes our inner conception about how the limb responds to motor commands. Although much is known about how sensory information is encoded and how motor commands drive movements, relatively little is known about these internal models and how they impact the selection of motor commands. Using multi-electrode recordings in rhesus monkeys, we leveraged a brain-machine interface (BMI) paradigm and novel statistical tools to study internal models in unprecedented detail. We extracted subjects' internal models of the BMI from neural population activity at the temporal resolution of tens of milliseconds and at the spatial resolution of individual neurons. These internal models describe subjects' prior beliefs about how their neural activity drives movement of the device. We discovered that mismatch between subjects' internal models and the actual BMI explains roughly 65% of their movement errors and substantially reduces subjects' dynamic range of movement speeds. Taken together, these findings show that outwardly aberrant behavior can be explained by taking into account a subject's prior beliefs, as represented by an internal model.

I-40. Convolutional spike-triggered covariance analysis for estimating subunit models

Anqi Wu[1] ANQIW@PRINCETON.EDU
Il Memming Park[2,3] MEMMING@AUSTIN.UTEXAS.EDU
Jonathan W Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Stony Brook University
[3]University of Texas at Austin

Subunit models provide a powerful yet parsimonious description of neural spike responses to complex stimuli. They can be expressed by an “LN-LNP” cascade, in which the first stage involves linear projection onto a bank of identical, spatially shifted subunit filters, and a fixed nonlinearity transforms each subunit filter output. The outputs of these convolutional subunits then drive spiking via a linear-nonlinear-Poisson (LNP) cascade. Recent interest in such models has surged due to their biological plausibility and accuracy for characterizing early sensory responses. However, fitting subunit models poses a difficult computational challenge due to the expense of evaluating the log-likelihood and the ubiquity of local optima. Here we address this problem by forging a theoretical connection between spike-triggered covariance analysis and nonlinear subunit models. Specifically, we show that a “convolutional” decomposition of the spike-triggered average (STA) and covariance (STC) provides an asymptotically efficient estimator for the subunit model under certain technical conditions: the stimulus is Gaussian, the subunit nonlinearity is well approximated by a second-order polynomial, and the final nonlinearity is exponential. In this case, the subunit model represents a special case of a canonical generalized quadratic model (GQM), which allows us to apply the expected log-likelihood trick (Park & Pillow 2011) to reduce the log-likelihood to a form involving only the moments of the spike-triggered distribution. Convolutional STC thus has fixed computational cost that does not scale with dataset size. We show that it outperforms highly regularized versions of the GQM, and achieves nearly the same prediction performance as the full maximum-likelihood estimator, yet with substantially lower cost. We apply these methods to neural data from macaque primary visual cortex and examine the number of subunits needed to model responses of different cell types.

I-41. Hierarchies for representing temporal sound cues in primary and secondary sensory cortical fields

Heather Read HEATHER.READ@UCONN.EDU
Christopher Lee CML.EMAIL@GMAIL.COM
Ahmad Osman AHMAD.OSMAN@UCONN.EDU
Monty Escabi ESCABI@ENGR.UCONN.EDU

University of Connecticut

Animals perceive loudness and timbre shape cues caused by changes in the sound envelope amplitude over time (Irino and Patterson 1996). In addition, they perceive periodicity temporal cues when sounds are repeated (Joris et al., 2004). Though primary (A1) and secondary cortical fields are implicated in perception of temporal sound cues, the underlying mechanisms remain unknown. In A1, sound shape and periodicity cues are represented with spike rate changes (Lu et al., 2001; Joris et al., 2004). Here we examine temporal sound cue sensitivities in A1 and two candidate acoustic object areas, ventral (VAF) and caudal suprarhinal (SRAF) auditory fields. Fifty-five unique shaped periodic noise sequences are generated to probe sensitivities to sound shape and modulation frequency in anesthetized rats. Shuffled auto-correlograms and modified Gaussian fits are generated to quantify spike-timing errors, jitter and reliability (Zheng and Escabi, 2008). We find jitter decreases logarithmically with increasing shape cutoff filter (Fc) in all three cortical fields. In contrast, reliability decreases proportionally with increasing modulation frequency (Fm). Thus, jitter and reliability provide plausible neural codes for sound shape and periodicity cues, respectively. Significantly, this suggests jitter and reliability errors have distinct biological sources and are not simply ‘noise’. A1 responses are sound onset driven with low spike-timing jitter indicating more precise temporal cue encoding than ventral fields; whereas, sustained spiking with larger jitter follows the shape of noise in ventral fields. Across regional neuron populations, there is a rank order increase in average jitter with: A1 $<$ VAF $<$ cSRAF. Likewise, there is a rank order decrease in upper cutoff frequencies for reliable spiking with: A1 $>$ VAF $>$ cSRAF. This suggests a general hierarchy for representing temporal cues on multiple time scales across distinct auditory cortical fields in a manner akin to that observed in visual sensory cortices (Andermann et al., 2011).

I-42. A neuronal network model for context-dependent perceptual decision on ambiguous sound comparison

Chengcheng Huang[1] CHENGCHENGHUANG11@GMAIL.COM
Bernhard Englitz[2,3] BENGLITZ@GMAIL.COM
Shihab Shamma[4] SAS@ISR.UMD.EDU
John Rinzel[1] RINZELJM@GMAIL.COM

[1]New York University
[2]Radboud University
[3]Donders Institute
[4]University of Maryland

Many natural stimuli contain perceptual ambiguities that can be cognitively resolved by the surrounding context, e.g. preceding stimuli. In audition, preceding context can bias the perception of speech and non-speech stimuli. Here, we develop a neuronal network model that accounts for context-dependent perception of pitch change between two successive stimuli - Shepard tones (each comprised of many octave- spaced pure tones). The pair's pitch change is ambiguous for Shepard tones separated by one-half octave, the tritone comparison. Listeners experience opposite percepts (either ascending or descending) if an ambiguous pair is preceded by a sequence of Shepard tones that lie in the half-octave above or below the first tone. We developed a recurrent, firing-rate network model, based on frequency-change selective units, that successfully accounts for the context-dependent perception demonstrated in behavioral experiments. The model consists of two tonotopically organized, excitatory populations, Eup and Edown, that respond preferentially to ascending or descending stimuli in pitch, respectively. These preferences are generated by an inhibitory population that provides inhibition asymmetric in frequency to the two populations; context dependence arises from slow facilitation of inhibition. We show that contextual influence depends on the spectral distribution and number of preceding tones. Further, using phase-space analysis, we demonstrate how the mechanistic basis, facilitated inhibition from previous stimuli and the waning inhibition from the just-preceding tone, shapes the competition between the Eup and Edown populations. In conclusion, our model accounts for contextual influences on the pitch change perception of an ambiguous tone pair by introducing a novel decoding strategy based on direction-selective units. The model's network architecture and slow facilitating inhibition emerge as predictions of neuronal mechanisms for these perceptual dynamics. Since the model structure does not depend on the specific stimuli, it may well generalize to other contextual effects and stimulus types.

I-43. Reading a cognitive map: a hippocampal model of path planning

Christopher Nolan[1,2] C.NOLAN@UQ.EDU.AU
Janet Wiles[1] JANETW@ITEE.UQ.EDU.AU
Allen Cheung[2] A.CHEUNG@UQ.EDU.AU

[1]The University of Queensland
[2]The Queensland Brain Institute

The mammalian hippocampal formation is important for storing episodic memories and performing spatial navigation tasks. The 'place cells' of the hippocampus strongly correlate with an animal's position in an environment, and as a population appear to provide a unique position code. Existing models explain how position correlates may arise from self-motion and/or external cues, but provide little insight into how a position code per se can be used for navigation. Importantly for path planning, the physical connectedness between two positions in space is not guaranteed a priori. A barrier or obstacle may prevent a nearby location from being reached, except via a detour. We have developed a model of hippocampal subregions CA3 and CA1. Our CA3 model implements the connectivity of the recurrent network of the rodent CA3 region, and learns the connectedness of space in an environment. Our CA1 model receives feed-forward projections from CA3, and can use the learned spatial connectivity to successfully plan paths to goals. This model simultaneously accounts for place cell data at the

anatomical, neurophysiological and behavioural levels, notably: 1) theta phase precession occurs in both CA3 and CA1 cells; 2) ventral place fields have larger spatial extents than dorsal place fields; 3) place fields enlarge with experience; 4) there is no topographic relationship between adjacent place cells. The model's results suggest: 1) multiple detours can be generated to bypass a given set of obstacles; 2) planning-dependent spatial behaviour will be affected by CA3 inactivation; 3) place fields are affected by barriers because of interference with spatial connectivity; 4) CA3 recurrent connections may be used for encoding spatial connectivity (not just pattern completion); 5) large place fields represent the region from which an animal can reach a location, not merely the current location.

I-44. Functional network & homologous brain region discovery using hierarchical latent state space models

Kyle Ulrich KYLE.ULRICH@DUKE.EDU
Jeff Beck JEFF.BECK@DUKE.EDU
Katherine Heller KHELLER@STAT.DUKE.EDU

Duke University

Dynamic latent state space models are often used to infer functional connectivity between brain regions from which we have simultaneously recorded data. However in many experiments high fidelity simultaneous recordings of all brain regions of interest are often unavailable due to limitations on the number of implantable electrodes or the accessibility of the brain regions of interest. Here we propose a nonparametric hierarchical Bayesian latent state space model designed to leverage data recorded from different brain regions, during different recording sessions, and even different sensor modalities (FMRI, LFP, multiunit electrophysiology, behavioral output etc.) allowing us to discover functional connectivity between brain regions and, possibly, homologous regions across species. The model assumes that a low dimensional, real valued, state vector summarizes the task relevant information contained in each task relevant brain region and that these region specific state vectors are sparsely connected to one another to form a linear (or weakly non-linear) dynamical system. Similarly, task relevant stimuli that drive the brain play the role of a sparsely connected driving force, while behavioral output is modeled as a noisy observation of one of the 'brain regions'. Critically, we assume that the neural network that transforms inputs into behavior does not change throughout the entire course of the experiment. Across recording sessions, however, changes do occur in sensor modality and sensor location. These changes are modeled by resampling sensor dynamics and the region specific state vectors that drive them at the start of each recording session. In this poster, we will present the generative model, a sketch of the associated variational Bayesian inference algorithm, and results on simulated and real data from a motivation for reward task.

I-45. Simulating cochlear implant modulation detection thresholds with a biophysical auditory nerve model

Elle O'Brien ANDRONOVHOPF@GMAIL.COM
Nikita Imennov IMENNOV@GMAIL.COM
Jay Rubinstein RUBINJ@U.WASHINGTON.EDU

University of Washington

Modern cochlear implants represent natural sounds with time-varying amplitude modulation of a carrier signal. To improve the efficacy of cochlear implant stimulation strategies, we seek to understand the neural code for the detection of envelope modulation. A psychophysical measure of sensitivity to these modulations is the modulation detection threshold (MDT). In this study, we present a computational model of the auditory nerve fiber population and discrimination process that predicts three behavioral trends: (1) improved sensitivity at higher stimulus intensity, (2) decreased sensitivity at higher modulation frequencies, and (3) decreased sensitivity at higher carrier pulse rates. Unlike previous efforts to characterize MDTs with computational modeling, we implement a stochastic, biophysical neural population model with physiologically distributed diameters. This model accurately describes temporal response characteristics of single fibers and predicts sensitivity to temporal fine structure changes. MDTs were computed with an ideal observer discriminator; threshold is defined as the minimum modulation depth where a modulated and unmodulated stimulus are discriminable 79.4% of the time. In addition to predicting major behavioral trends, our model suggests that a small subset of fibers have a strong effect on the shape and limens of the MDT curves. For all stimuli tested, certain fibers are capable of excellent discrimination, but the behavioral trend is best described by discrimination with all fiber inputs pooled. The temporal precision of the likelihood estimator influences the overall limens but not the shape of the threshold curve. These results demonstrate the importance of using a multi-diameter fiber population in modeling the modulation transfer function and indicate a wider applicability of our model to estimating behavioral performance in cochlear implant listeners.

I-46. Evaluating ambiguous associations in the amygdala by learning the statistical structure of the environment

Tamas Madarasz[1,2] TAMASMADARASZ1@GMAIL.COM
Lorenzo Diaz-Mataix[1] LDM5@NYU.EDU
Joseph LeDoux[1] JEL1@NYU.EDU
Joshua Johansen[2] JJOHANS@BRAIN.RIKEN.JP

[1]New York University
[2]RIKEN Brain Science Institute

Recognizing predictive relationships in the environment is critical for survival, but a detailed understanding of the underlying neural mechanisms remains elusive. In particular it is not clear how the brain distinguishes predictive relationships from spurious ones, and how it computes associations when these relationships are ambiguous. A simple example of an ambiguous cue-outcome relationship arises when an outcome occurs both in the presence and absence of a sensory cue (so-called contingency degradation). Established accounts of classical conditioning characterize contingency learning as fitting parameters in a fixed generative or discriminative model of the environment, with different sensory cues competing with each other to predict important outcomes. In a series of auditory fear conditioning experiments we show instead that interference from competing associations is not required to learn a reduced cue-outcome contingency, and that changes in the strengths of different associations are dissociable. Instead, building on previous work (Griffiths and Tenenbaum, 2005), we propose a computational model for conditioning that directly assesses different models of the environment's statistical structure (the set of predictive or causal relationships) when encountering novel stimuli. This model employs structure learning and Bayesian model averaging over Bayesian Network representations of the environment, and gives a good quantitative account for our behavioral data incorporating different known conditioning phenomena, including contingency degradation, overshadowing, partial reinforcement and second-order conditioning. Using optogenetic and electrophysiological techniques, we also identify a well-defined cell population, pyramidal cells in the lateral amygdala (LA), that regulate contingency evaluations in auditory aversive conditioning. Previous work has shown that cells in the LA and its input structures show a diversity of response properties, with many cells responding to different combinations of sensory stimuli, paralleling the diversity of graph structures seen in our statistical model. Together with those results, our findings suggest a novel view on how plasticity in the amygdala might encode environmental associations and uncertainty.

I-47. Dynamic Bayesian integration of vision and proprioception during feedback control

Frederic Crevecoeur[1] FCREVECOEUR@GMAIL.COM
Douglas Munoz[2] DOUG.MUNOZ@QUEENSU.CA
Stephen Scott[2] STEVE.SCOTT@QUEENSU.CA

[1]Catholic University of Louvain
[2]Queen's University

Static Bayesian models of multi-sensory integration suggest that the central nervous system (CNS) combines distinct sensory signals by weighting them according to their reliability. Such models explain many aspects of decision-making or movement planning. However, it remains unclear how the CNS performs multi-sensory integration during movement execution, which in theory requires considering not only sensory variability but also differences in sensory delays. We addressed this issue by instructing subjects (N = 10) to visually track their fingertip or a white cursor perturbed either by applying mechanical perturbation on their limb or by displaying a trajectory matching participants fingertip motion (visual perturbations). We found that saccadic reaction times following mechanical perturbations were similar with or without vision (P=0.1), and both were significantly shorter than following visual-only perturbations ($P<0.01$). The end-point error at the end of the first saccade was also similar with or without vision (P=0.21), whereas the error following visual perturbations was significantly smaller ($P<0.05$). These results fit with a dynamic Bayesian model, in which proprioception plays a dominant role resulting from smaller delays. Finally, the saccade end-point variance was only slightly reduced with vision (one-sided comparison: P=0.029), but remained greater than following visual perturbations (P=0.016). This is inconsistent with a static model predicting that the posterior variance should be smaller than the variance of each sensory modality taken independently. In contrast our results are well captured in a dynamic model that considers differences in sensory delays. Altogether our data emphasize a rapid correction of state estimation mediated by limb afferent feedback, and suggest that dynamic Bayesian estimation underlies real-time multi-sensory integration and feedback control.

I-48. Varied network connections yield a gamma rhythm robust to variation in mean synaptic strength

Mark Reimers[1] REIMERSM@MSU.EDU
Steven Hauser[2] SCH5ZC@VIRGINIA.EDU

[1]Michigan State University
[2]University of Virginia

The gamma rhythm has been implicated in many cognitive functions and disruptions in gamma have been linked with mental illnesses such as schizophrenia. Numerous mechanisms have been proposed for gamma rhythm dynamics and many have been implemented as computational models. Although GABA-ergic inhibitory signalling is widely held to be important, nevertheless a consensus has not yet been reached on the mechanisms of the cortical gamma rhythm. In particular little is known about how necessary a balance between excitatory and inhibitory inputs is to a well-formed gamma rhythm. Furthermore recent data suggest that the overall strength of inhibitory signaling varies more than twenty-fold between individual human brains; little is known about how dynamic compensation might stabilize gamma rhythm phenotypes across such large variation. Few computational models of any biological phenomenon are robust to such dramatic variation. We modified the Börgers-Kopell model for gamma rhythms in order to conform more closely to neuroanatomical and physiological data. In particular we implemented a sparse topology of connections, a long-tailed distribution of synaptic weights, and a much shorter inhibitory decay time constant, following measures reported in the literature; we also removed strong drive to the inhibitory neurons. We show that the gamma rhythm produced by the modified model conforms much more closely to the statistics of gamma rhythms measured in vivo. Furthermore the matching of excitatory and inhibitory currents, which is increasingly seen as important, is roughly trebled in the modified model. Other aspects

of emergent dynamics, such as recovery times also conformed much better to experimental observations. Finally although the usual instantiation of this model is not robust to large variation in the synaptic strength parameter, the modified version maintains power and frequency peaks over the range of variation implied by genomic data.

I-49. Invariant representations for action recognition in the visual system

Leyla Isik LISIK@MIT.EDU
Andrea Tacchetti ATACCHET@MIT.EDU
Tomaso Poggio TP@AI.MIT.EDU

Massachusetts Institute of Technology

The human brain can rapidly parse a constant stream of visual input. The majority of visual neuroscience studies, however, focus on responses to static, still-frame images. Here we use magnetoencephalography (MEG) decoding and a computational model to study invariant action recognition in videos. We created a well-controlled, naturalistic dataset to study action recognition across different views and actors. We find that, like objects, actions can also be read out from MEG data in under 200 ms (after the subject has viewed only 5 frames of video). Action can also be decoded across actor and viewpoint, showing that this early representation is invariant. Finally, we developed an extension of the HMAX model, inspired by Hubel and Wiesel's findings of simple and complex cells in primary visual cortex as well as a recent computational theory of the feedforward ventral stream, which is traditionally used to perform size- and position-invariant object recognition in images, to recognize actions. We show that instantiations of this model class can also perform recognition in natural videos that are robust to non-affine transformations. Specifically, view-invariant action recognition and action invariant actor identification in the model can be achieved by pooling across views or actions, in the same manner and model layer as affine transformations (size and position) in traditional HMAX. Together these results provide a temporal map of the first few hundred milliseconds of human action recognition as well as a mechanistic explanation of the computations underlying invariant visual recognition.

I-50. Deep Gaze I: Boosting saliency prediction with feature maps trained on ImageNet

Matthias Kummerer MATTHIAS.KUEMMERER@BETHGELAB.ORG
Thomas Wallis TOM.WALLIS@BETHGELAB.ORG
Lucas Theis LUCAS@BETHGELAB.ORG
Matthias Bethge MATTHIAS@BETHGELAB.ORG

University of Tübingen

Among the wide range of complex factors driving where people look, the properties of an image that are predictive for fixations under “free viewing” conditions have been studied the most. Here we frame saliency models probabilistically as point processes, allowing the calculation of log-likelihoods and bringing saliency evaluation into the domain of information theory. We compare the information gain of all high-performing state-of-the-art models to a gold standard and find that only one third of the explainable spatial information is captured. Thus, contrary to previous assertions, purely spatial saliency remains a significant challenge. Our probabilistic approach also offers a principled way of understanding and reconciling much of the disagreement between existing saliency metrics. Finally, we present a novel way of reusing deep neural networks that have been pretrained on the task of object recognition in models of fixation prediction. Using the well-known network developed by Krizhevsky et al. (2012), we propose a new saliency model, “Deep Gaze I”, that accounts for high-level features like objects and popout. It significantly outperforms previous state-of-the-art models on the MIT Saliency Benchmark and now explains more than half of the explainable information.

I-51. Explaining monkey face patch system as deep inverse graphics

Ilker Yildirim[1,2] ILKERY@MIT.EDU
Tejas Kulkarni[1] TEJASK@MIT.EDU
Winrich Freiwald[2] WFREIWALD@MAIL.ROCKEFELLER.EDU
Josh Tenenbaum[1] JBT@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Rockefeller University

The spiking patterns of the neurons at different fMRI-identified face patches show a hierarchical organization of selectivity for faces: neurons at the most posterior patch (ML/MF) appear to be tuned to specific view points, AL (a more anterior patch) neurons exhibit specificity to mirror-symmetric view points, and the most anterior patch (AM) appear to be largely view-invariance, i.e., neurons there show specificity to individuals (Freiwald & Tsao, 2010). Here we propose and implement a computational characterization of the macaque face patch system. Our main hypothesis is that face processing is composed of a hierarchy of processing stages where the goal is to “inverse render” a given image of a face to its underlying 3d shape and texture. The model wraps and fine-tunes a convolutional neural network (CNN) within a generative vision model of face shape and texture. We find that different components of our model captures the qualitative properties of each of the three face patches. ML/MF and AL are captured by different layers of the fine-tuned CNN, whereas AM is captured by the latent variables of the generative model. This modeling exercise makes two important contributions: (1) mirror symmetry (as in the patch AL) requires dense connectivity from the layer below and requires the agent to observe mirror-symmetric images that belong to the same identity, (2) AM is best explained by a representation that consists not only of latent shape and texture variables but also of the latent variables for generic scene variables such as pose and light location, indicating that this most anterior patch should be equivariant.

I-52. Inference of functional sub-circuits in the dynamical connectome of the C. elegans worm

Eli Shlizerman[1] SHLIZEE@UW.EDU
S Grant Babb[2] SGRANTBABB@GMAIL.COM

[1]University of Washington
[2]Intel

To infer functional sub-circuits in the neuronal network of Caenorhabditis elegans (C. elegans) worm we construct a model for neuronal dynamics and develop an algorithm for constructing a probabilistic graphical model (PGM) for the network from simulated neural responses. To construct the dynamical model, we combine known connectome data with physiologically appropriate neuron model to simulate the dynamics of the C. elegans worm network in response to stimuli over time. We then verify the model by comparing simulated and experimentally measured I-V curves at equilibrium and by stimulation of commanding neurons, which invoke experimentally characterized dynamics. To infer dominant sub-circuits/pathways being activated in response to stimuli, we propose an algorithm that performs individual stimulations of neurons and analyzes the simulated network dynamics using singular value decomposition to derive a map of dependencies between neurons (i.e. functional connectome). We find that the functional connectome is significantly different than the static connectome as it reflects recurrent interactions and nonlinear responses within the network. We then construct an undirected PGM in which nodes correspond to neurons and edges correspond to potentials computed from the functional connectome. By performing posterior inference, i.e., conditioning the PGM with known commanding neurons as excited (setting high probability to these nodes), we are able to infer the motorneurons expected to respond to such stimulation. Such a framework can potentially allow for identifying unknown commanding neurons and subcircuits that they activate.

I-53. Energy-based latent variable models as efficient descriptors of multi-neuron activity patterns

Dominik Spicher[1] SPICHER@PYL.UNIBE.CH
Gerrit Hilgen[2] GERRIT.HILGEN@NCL.AC.UK
Hayder Amin[3] HAYDER.AMIN@IIT.IT
Alessandro Maccione[3] ALESSANDRO.MACCIONE@IIT.IT
Sahar Pirmoradian[4] SPIRMORA@INF.ED.AC.UK
Gasper Tkacik[5] GASPER.TKACIK@IST.AC.AT
Evelyne Sernagor[2] EVELYNE.SERNAGOR@NCL.AC.UK
Luca Berdondini[3] LUCA.BERDONDINI@IIT.IT
Matthias Hennig[4] M.HENNIG@ED.AC.UK

[1]University of Bern
[2]University of Newcastle
[3]Italian Institute of Technology
[4]University of Edinburgh
[5]IST Austria

Local populations of neurons are commonly held to be the substrate of neuronal computation. Insights into interdependencies among participating neurons are thus a prerequisite for understanding the emergence of meaningful computation from many discrete units in neural systems. Maximum Entropy (ME) models, which explicitly reproduce some features of an experimental distribution and are otherwise maximally unstructured, are a popular approach for constructing models of neuronal populations and understanding their emergent behaviour in terms of underlying factors. Here build upon work by Köster and colleagues and investigate another energy-based model, the Restricted Boltzmann Machine (RBM), based on introducing unobserved hidden units, to study the composedness of the correlational structure of neuronal populations. Using large scale multi-electrode array recordings from hippocampal cultures and mouse retina, we show that RBMs compare favourably with ME models in terms of constructing low-dimensional descriptions of the potentially exponential correlation structure within a neuronal population. We first compared the two approaches for small groups of neurons, systematically varying model complexity, and found that RBMs require fewer parameters than ME models to reproduce state probabilities with a given accuracy. For large groups of 60 neurons, we then compared RBMs with pairwise and k-pairwise ME models. The ease of increasing expressive power of RBMs by adding hidden units allows the construction of accurate models with few parameters, removing biases in estimating state likelihoods and triple correlations exhibited by the ME models we use. Notably, RBMs also allow for exact computation of many of those measures that require sampling-based approaches in large ME models. We conclude that in terms of ease of construction, goodness of fit and ease of evaluation, RBMs provide a viable approach to ME models in elucidating underlying factors for the emergent behaviour of local neuronal networks.

I-54. A neural network model for arm posture control in motor cortex

Hagai Lalazar[1] HL2757@COLUMBIA.EDU
Larry Abbott[1] LFA2103@COLUMBIA.EDU
Eilon Vaadia[2] EILON.VAADIA@ELSC.HUJI.AC.IL

[1]Columbia University
[2]The Hebrew University of Jerusalem

Neurons with nonlinear mixed-selectivity are being discovered in an increasing number of cortical areas. During an arm posture control task, M1 neurons exhibited diverse forms of complex nonlinearities including high spatial frequencies and non-separable mixing of concurrent task parameters. While several theoretical models used the equations of motion of a model arm to predict canonical single-neuron responses, they do not account for this diversity of complex responses, nor how such diversity contributes to M1 computation. We constructed a feedforward neural network model using biologically realistic inputs and random connectivity that, with only a small number of free parameters, accurately replicates the M1 data. This model offers a simple mechanism that accounts for the existence of both the linear and nonlinear components in the neuronal responses, without the need to adjust synaptic weights or explicit single-neuron encoding. To our knowledge, this is one of the first examples where the diversity of nonlinear responses of a population of real neurons is replicated by a neural network model. This model offers novel predictions about aspects of sensorimotor processing along the dorsal stream. It suggests a mechanistic explanation for how the increased visual acuity from foveal vision (primates automatically saccade to a reaching target) can also increase motor accuracy. Moreover, the model predicts that the multimodal visual-and-posture responses along the dorsal stream are not only evidence of sequential sensorimotor transformations, but also of the parallel mixing of task parameters required to ultimately activate the same muscles. By decoding the EMG of major muscles in this task, we show that the complex nonlinear components are as informative as the linear ones. This study supports the view that M1 neurons serve as basis-functions for generating muscle patterns during arm posture control and that this basis consists of diverse complex nonlinear responses, not canonical (parametric) single-neuron responses.

I-55. Modeling reactivation of neuronal network spiking patterns across human focal seizures

Felipe Gerhard[1] FELIPE_GERHARD@BROWN.EDU
Sydney Cash[2] SCASH@MGH.HARVARD.EDU
Wilson Truccolo[1] WTRUCCOLO@GMAIL.COM

[1]Brown University
[2]Harvard University

Hypersynchronous and excessive neuronal activity are believed to cause epileptic seizures. Recent studies of the spiking activity in ensembles of single neurons in people with intractable focal epilepsy have shown, however, that neuronal dynamics during seizures can be highly heterogeneous, yet reproducible across consecutive seizures. Here, we show for the first time that the reproducible structure underlying the spatiotemporal patterns can be captured by point process models of the ensemble dynamics. We studied two main types of human focal seizures: spike-wave complex (SWC) and gamma-band seizures. Neural dynamics during SWC seizures were best described by a two-dimensional latent state-space model (cross-validated AUC = 0.91, 10ms resolution). Gamma-band seizure dynamics were less predictable and higher dimensional (optimal dimensionality: 3-6, AUC = 0.70). Remarkably, in both cases, models trained on one seizure generalized to a second seizure recorded in the same patient with equal predictive power (AUC = 0.91 and 0.72 for SWC and gamma-band, respectively). Additionally, for both types of seizures, generalized linear models (GLMs) that included effective network interactions and global fluctuations in population activity captured the same information as the low-dimensional latent model (AUC = 0.92 and 0.68 for SWC and gamma-band, respectively). GLMs generalized across seizures equally well: cross-seizure prediction AUC = 0.93 (SWC) and 0.69 (gamma-band seizures). Predictions by both models were partially redundant: correlation coefficient r = 0.78 (SWC) and 0.76 (gamma-band). Overall, our findings suggest the existence of recurring network patterns at the level of ensembles of single neurons, rather than a simple random recruitment of neuronal subpopulations to engage in each seizure. Furthermore, low-dimensional dynamics emerge at seizure onset and tend to consistently reactivate in consecutive seizures. We expect the above models of spatiotemporal neuronal dynamics to contribute to the development of prediction and early-detection algorithms for seizure control systems.

I-56. Time versus gesture coding in songbird motor cortex

Galen Lynch[1] GLYNCH@MIT.EDU
Tatsuo Okubo[1] OKUBO@MIT.EDU
Alexander Hanuschkin[2,3] ALEXANDER.HANUSCHKIN@INI.PHYS.ETHZ.CH
Richard Hahnloser[3] RICH@INI.PHYS.ETHZ.CH
Michale Fee[1] FEE@MIT.EDU

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

Learned motor behaviors require animals to learn and deploy precisely timed motor sequences. The origin of timing precision in motor behaviors remains unclear: while some premotor brain regions have activity related to motor timing, most are thought to encode features of motor gestures. Songbirds learn and produce complex sequences of vocal gestures with millisecond precision, and have emerged as a model system for understanding the neural basis of complex learned motor behaviors. Bird song requires the premotor nucleus HVC; projection neurons in HVC burst sparsely at stereotyped times in the song, with different neurons active at different times in the motif. Two models of HVC coding have been proposed: one hypothesis states that projection neuron bursts form a continuous sequence during song that represents motor time to support both song production (Leonardo & Fee 2005) and learning (Fee & Goldberg 2011); another hypothesis proposed by Amador et al (2013) states that bursts only occur at specific times in the song corresponding to discontinuities and extrema in the trajectories of vocal/motor gestures (called GTE). In the GTE model, HVC would become inactive between GTE times, rendering these models incompatible with each other. Using a large dataset of projection neurons recorded in singing birds, we tested predictions of this GTE model, including the alignment of neural activity to GTE, and the temporal clustering of activity expected to arise from such alignment. We find that the data do not support the predictions of the GTE model. Furthermore, the distribution of projection neuron bursts throughout song is consistent with neurons being equally likely to burst at every time in song. Our findings suggest that while each HVC projection neuron is sparsely active, as a population they are continuously active during song and may form a time basis used to learn and produce the song.

I-57. Predicting the relationship between sensory neuron variability and perceptual decisions

Yuwei Cui[1,2] YWCUI@UMD.EDU
Liu Liu[3] LIU.LIU2@MAIL.MCGILL.CA
James McFarland[1] JAMES.MCFARLAND@GMAIL.COM
Christopher Pack[3] CHRISTOPHER.PACK@MCGILL.CA
Daniel Butts[4] DAB@UMD.EDU

[1]University of Maryland, College Park
[2]Numenta Inc.
[3]McGill University
[4]University of Maryland

Cortical neuron responses can be quite different across repeated presentations of the same stimulus. In perceptual discrimination tasks, differences in the response of individual cortical neurons can be correlated with perceptual decisions, but little is known about how such correlation arises. We bridged these two vastly different scales between single neuron activity and behavior by detecting the influence of ongoing cortical network activity, which we found has direct relationships to both single neuron activity and perceptual decisions. We simultaneously recorded spikes and local field potentials (LFPs) with a multi-electrode array from the middle temporal area (MT) of the alert macaque monkey during performance of a motion discrimination task, or during passive viewing of optic flow stimulus. Through the application of data-constrained models that used the LFP and multi-unit

activity to predict MT neuron spikes, we demonstrated that a significant fraction of the MT neuron variability was predictable from the LFPs during both tasks. Spiking responses of individual MT neurons were strongly linked to two frequency bands in the LFP: gamma oscillations (~35 Hz), and delta oscillations (~2 Hz), which provided complementary information to stimulus-driven and multi-unit activity in both passive and task conditions. Moreover, individual MT neurons exhibited choice-related fluctuations in firing rate that was predicted by MT neurons' phase preference in the delta-band LFP. These results therefore identify signatures of network activity related to the variability of cortical neuron responses, and suggest their central role in modulation of sensory neuron function.

I-58. Dynamics of multi-stable states during ongoing and evoked cortical activity

Luca Mazzucato LUCA.MAZZUCATO@STONYBROOK.EDU
Alfredo Fontanini ALFREDO.FONTANINI@STONYBROOK.EDU
Giancarlo La Camera GIANCARLO.LACAMERA@STONYBROOK.EDU

SUNY at Stony Brook

The concerted activity of ensemble neurons in sensory cortex of alert animals can be characterized in terms of sequences of metastable states, each state defined by a specific and quasi-stationary pattern of firing rates. While single-trial metastability has been studied for its role in sensory coding, memory and decision-making, very little is known about the network mechanisms responsible for its generation. A common assumption, at least for sensory cortices, is that the emergence of state sequences is triggered by an external stimulus. Here we show that metastability can be observed also in the absence of overt sensory stimulation. Multi-electrode recordings from the gustatory cortex of alert rats revealed that ongoing activity can be characterized in terms of metastable sequences wherein many single neurons spontaneously attain 3 or more different firing rates, a feature we termed 'multi-stability.' Single neuron multi-stability represents a challenge to existing spiking network models, where typically each neuron is at most bi-stable. We present a spiking network model that accounts for both network metastability in the absence of stimulation and multi-stability in single neuron firing. In the model, each state results from the activation of a fraction of neural clusters with potentiated intra-cluster connections, with the firing rate in each cluster depending on the number of active clusters. Simulations show that the model's ensemble activity hops between the different states, reproducing the metastable dynamics observed in the data. When probed with external stimuli, the model reproduced two key properties of evoked activity: the quenching of single neurons multi-stability into bi-stability and a reduction of trial-by-trial variability. The model further predicts that stimulus onset reduces the dimensionality of ensemble activity, a new phenomenon that was confirmed in the data. Altogether, our results provide a unifying and mechanistic framework that accounts for ongoing and evoked dynamics in cortical circuits.

I-59. Altered cortical neuronal avalanches in a rat model of schizophrenia.

Saurav Seshadri SAURAV.SESHADRI@NIH.GOV
Dietmar Plenz PLENZD@MAIL.NIH.GOV

National Institute of Mental Health

In various experimental systems, ranging from organotypic slices to the rodent, non-human primate, and human brain, resting state activity in the cortex organizes as neuronal avalanches. These are spontaneous spatiotemporal cascades of synchronized activity that exhibit a precise organization, described by power laws. Avalanche dynamics have been shown to optimize several aspects of cortical processing, including information capacity and dynamic range. Pharmacological studies show that avalanche dynamics are dependent on excitatory-inhibitory (E-I) balance, as well as dopaminergic signaling. NMDA receptor hypofunction, which affects E-I balance, is thought to be a core pathophysiological feature of schizophrenia. We therefore hypothesized that avalanche dy-

namics may be disrupted in a rodent model of schizophrenia. To address this, we injected rats with the NMDAR antagonist phencyclidine (PCP, s.c., 10 mg/kg) at postnatal day 7, 9, and 11. Neuronal activity was recorded in superficial layers of cortex by in vivo 2-photon imaging of pyramidal neurons expressing the genetically encoded calcium indicator YC 2.60. We found that cluster size distributions significantly deviated from a power law in PCP-treated rats, compared with vehicle-treated littermates (DKS - Vehicle: 0.06, PCP: 0.11, p = 0.02). Statistical measures of activity at the individual neuron level (such as the coefficient of variation in interspike intervals) were unchanged, indicating that altered avalanche dynamics represent an emergent, network-level phenotype. PCP-treated rats also exhibited deficits in visual working memory in the novel object recognition task. Rats were treated with D-serine, an NMDAR co-agonist which has shown clinical efficacy in schizophrenic patients, to rescue these phenotypes. These results have two important implications: first, that ongoing neuronal group activity at single-cell resolution organizes as neuronal avalanches and is sensitive to E-I disruption in vivo, and second, that avalanche dynamics could characterize cortical dysfunction in schizophrenia, and potentially serve as a biomarker for diagnosis or drug screening.

I-60. Effects of sprouting and excitation-inhibition imbalance on epileptiform activity

Christopher Kim[1] C.KIM@BCF.UNI-FREIBURG.DE
Ulrich Egert[1] ULRICH.EGERT@IMTEK.UNI-FREIBURG.DE
Arvind Kumar[2,1] ARVKUMAR@KTH.SE
Stefan Rotter[1] STEFAN.ROTTER@BIOLOGIE.UNI-FREIBURG.DE

[1]Bernstein Center Freiburg
[2]KTH Royal Institute of Technology

Human patients and animal models of mesial temporal lobe epilepsy exhibit pronounced mossy fiber sprouting in dentate gyrus and imbalance of excitation-inhibition (EI) in surrounding regions. In this study, we considered a computational model of recurrently connected excitatory and inhibitory neurons to investigate the effects of mossy fiber sprouting and EI imbalance on the network dynamics. Mossy fiber sprouting was modeled by allowing a subpopulation of excitatory neurons to produce additional recurrent connection to itself and feedforward projections to the rest of the neurons in the network. The effects of EI imbalance was examined by changing the inhibitory synaptic strength onto sprouting and nonsprouting excitatory neurons. The population activity was investigated with Fokker-Planck equations, and network simulations were performed to verify the predictions. Main results of the study are that (1) the presence of a population of sprouting neurons large enough alters the network dynamics and (2) promoting activity in nonsprouting neurons stabilizes the increased activity due to sprouting. For networks whose sprouting population size is too small, e.g. 5% of the excitatory population in our model, adjusting EI balance of nonsprouting neurons towards excitation dominance increases the mean firing rate as expected. However, beyond a cirtical size of sprouting neurons (approximately 10%), promoting activity in nonsprouting neurons suppresses the elevated network activity caused by mossy fiber sprouting. Network simulations confirm that pushing EI balance of nonsprouting neurons towards excitation, by reducing inhibitory synaptic strength onto them or applying stimulus to nonsprouting neurons, can decrease the mean firing rate and terminate synchronized activity. Analyzing the stationary state of the Fokker-Planck equations provide concise explanations of this phenomenon. Our study points to the importance of not only mossy fiber sprouting in inducing epileptic activity but the role of surrounding regions in regulating aberrant activity.

I-61. Crowding of orientation signals in macaque V1 neuronal populations

Christopher Henry CHRISTOPHER.HENRY@EINSTEIN.YU.EDU
Adam Kohn ADAM.KOHN@EINSTEIN.YU.EDU

Albert Einstein College of Medicine

Spatial context can exert strong influences on the perception of visual objects or features. One well-studied example is visual crowding, in which the discriminability of an isolated target stimulus is impaired by the simultaneous presentation of surrounding distractor stimuli. Crowding has been studied extensively in perceptual work and, more recently, using functional imaging. This work has led to the suggestion that crowding involves the accurate representation of sensory signals in early cortex, followed by the unavoidable pooling of target and distractor stimuli in higher cortex. However, there has been little neurophysiological work on crowding, and its effects on neuronal population representations in early cortex are unknown. We recorded the responses of neuronal populations in macaque V1 to target sinusoidal gratings, presented in isolation or when surrounded by a set of distractor gratings ("crowders"). Crowders reduced the average neuronal firing rate by ~20%, accompanied by a slight reduction in response variability (Fano factors) and shared variability between neurons (spike count correlations). To assess the effect of crowding on the population representation of orientation, we applied linear discriminant analysis to V1 responses to pairs of targets with offset orientations. Decoders trained on responses to isolated stimuli had significantly worse performance on responses to crowded stimuli, across a wide range of task difficulties (4-20 degrees orientation offset). Performance deficits were also apparent when decoders were trained on crowded responses. Shuffling the neuronal responses to remove correlations reduced decoding performance, but the detrimental effects of crowding were still evident, suggesting that crowding arises in part from effects on the responses of individual neurons. Our study demonstrates that distractors reduce target feature representation in neuronal populations in primary visual cortex. Thus, while maladaptive pooling by higher cortical neurons may also contribute, crowding involves a loss of sensory information at the first stages of cortical processing.

I-62. Rapid context-dependent switching of auditory ensembles during behavior

Kishore Kuchibhotla KUCHIBH@GMAIL.COM
Robert Froemke ROBERT.FROEMKE@MED.NYU.EDU

New York University

Behavioral context helps to determine the value of external stimuli. The same sensory cues can have multiple meanings when presented in different environments (Gilbert and Li 2013; Fritz et al. 2003). How do neural circuits in the sensory cortex flexibly represent the same stimuli under different conditions? Using two-photon calcium imaging and whole-cell voltage clamp recordings in head-fixed behaving mice, we monitored neural responses to the same auditory cues in primary auditory cortex (A1) during behavior (active listening) and outside task context (passive hearing). For active listening, mice were trained to lick for a water reward after hearing a tone (CS+) and withhold from licking after hearing a different tone (CS-). In passive hearing, mice were exposed to the same tones without opportunity for reward. Surprisingly, calcium imaging demonstrated that the two contexts activated distinct neuronal ensembles. Changing one context to the other rapidly switched the neuronal ensemble encoding the same tone. Moreover, the topographic organization of A1 transiently lost structure during active behavior and returned immediately after task completion, suggesting that sensory stimuli may have context-dependent coding mechanisms. We then asked how the ensemble patterns might rapidly switch depending on context. Whole-cell voltage-clamp recordings in behaving mice showed that cortical inhibitory currents were significantly modulated between the two contexts in all measured neurons (ΔInh=45$\pm$9%, n=9/9 neurons with $\triangle$Inh$>$15%) while excitatory currents were less modulated (ΔExc=21$\pm$9%, n=3/9 neurons with ΔExc$>$15%). Moreover, the inhibitory currents altered the excitatory-inhibitory balance in a bidirectional manner: some neurons increased their E:I ratio for the active-context while others showed a preference for the passive context. Thus, cortical inhibition enables multiple representations of a sensory stimulus to co-exist in the same microcircuit by

dynamically repurposing sensory maps to focus cortical processing on task-relevant stimuli.

I-63. Dynamic integration of visual features by parietal cortex

Guilhem Ibos GUILHEMIBOS@GMAIL.COM
David J Freedman DFREEDMAN@UCHICAGO.EDU

University of Chicago

Feature-based attention (FBA) is the process by which behaviorally relevant stimulus features are selected. Recently we have shown that FBA shifts the feature tuning of neurons of the lateral intraparietal area (LIP) toward task-relevant feature values (Ibos and Freedman, 2014). In an additional experiment, monkeys attended both the color and the direction of visual stimuli located outside the receptive field of the neurons being recorded and ignored some distractors located inside. We have shown that spatial attention partially confined the effect of FBA on the representation of visual feature in LIP to the relevant position. We developed a neural network model which accounts for the flexible feature integration effects observed in LIP, by showing that tuning shifts in LIP arise as a consequence of linear integration of the attention-related response-gain changes which are known to occur in upstream visual areas, such as MT and V4. The model consists of two interconnected neuronal layers, a first layer (L1) sending feed-forward connections to a second layer (L2). L2 neurons integrate multiple inputs from a population of direction-tuned L1 neurons. Previous work in MT and V4 found that feature selective neurons show changes in response gain due to feature-based attention. Thus, we applied such gain modulations to L1 neurons and considered the impact on tuning in L2. We then considered the impact of gain-modulated L1 activity on L2, which is assumed to simply linearly integrate L1 activity. This revealed shifts of direction tuning in L2, highly similar to that described in our recent work in LIP, suggesting that the flexible integration of visual inputs to LIP is controlled by attentional modulations of sensory encoding in upstream cortical areas.

I-64. What does LIP do when a decision is dissociated from planning the next eye movement?

NaYoung So[1,2] NS2911@COLUMBIA.EDU
Michael Shadlen[1] SHADLEN@COLUMBIA.EDU

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

Decisions are often based on the accumulation of evidence. Such deliberation operates over a time span that outlasts the momentary evidence, which may be fleeting. For perceptual decisions, sensory neurons convey momentary evidence to higher order structures, which are capable of holding and incrementing persistent activity—that is, integrating. These high order structures are often associated with planning the behavior used to communicate the decision. For example, neurons in the lateral intraparietal area (LIP) represent the accumulation of evidence when the decision is communicated by a saccadic eye movement. Often, however, a decision is not reported immediately, but retained and retrieved later. This poses a challenge for structures like LIP, because other behaviors might ensue during retention (e.g., if LIP represents only the next saccade; Barash et al., 1991). We constructed a simple paradigm to test this. A monkey discriminated the direction of random dot motion, but completed a predictable sequence of irrelevant eye movements before reporting its decision. The intervention did not impair performance. During decision formation, LIP activity reflected evidence accumulation, despite the fact that the next eye movement was always to an irrelevant target (T0) outside the neuron's response field (RF). Upon fixation of T0, decision-related activity vanished, but appeared in other LIP neurons as the gaze change brought one of the choice targets into their RF. The decision information then reappeared when the gaze returned to the fixation point, accompanying preparation of the final eye movement report. Thus, even when the decision is to be reported later, (1) LIP neurons represent accumulating evidence if the outcome of the decision is associated with a particular saccade, and (2) the LIP population retains information about the decision as the information is passed among neurons in accordance with their RF location relative to a planned response.

I-65. Multiple mechanisms for stimulus-specific adaptation in the primary auditory cortex

Ryan G Natan RNATAN@MAIL.MED.UPENN.EDU
John Briguglio JOHNBRI@SAS.UPENN.EDU
Laetitia Mwilambwe-Tshilobo LAETITIA@MAIL.MED.UPENN.EDU
Maria Geffen MGEFFEN@MAIL.MED.UPENN.EDU

University of Pennsylvania

Adaptation to stimulus context is a ubiquitous property of cortical neurons and is thought to enhance efficiency of sensory coding. Yet the specific neuronal circuits that facilitate cortical adaptation remain unknown. In the primary auditory cortex (A1), the vast majority of neurons exhibit stimulus-specific adaptation (SSA), responding weakly to frequently repeated tones and strongly to rare tones. This form of history-dependent adaptation may increase cortical sensitivity to rare sounds. Here, we identify three distinct components shaping cortical SSA. The current source density sink amplitude profile across cortical layers in response to common and rare tones revealed that thalamo-cortical depression contributes significantly to cortical SSA. Furthermore, we found that two types of inhibitory interneurons control SSA in a complementary fashion. Optogenetic suppression of parvalbumin-positive interneurons (PVs) led to an equal increase in the firing rate of principal neurons to common and rare tones, suggesting that PVs control SSA by non-specific inhibition. Suppression of somatostatin-positive interneurons (SOMs) led to an increase in neuronal responses to frequent, but not to rare tones, suggesting that SOMs contribute to SSA directly, by stimulus-selective facilitation. Indeed, inhibitory effect of SOMs on excitatory neurons increased with successive tone repeats. In contrast, PVs provided constant inhibition throughout the stimulus sequence. The effects of PVs and SOMs differed across cortical layers, suggesting that the two interneuron types differentially transform responses to common and rare tone between input and output stages of the cortex. Interestingly, while excitatory neurons in A1 have been shown to exhibit complex patterns of adaptation to stimuli that include more than two tones, inhibition from PVs and SOMs was not modulated by higher order stimulus deviance. Taken together, our results demonstrate that SSA in the auditory cortex is a product of multiple adaptation mechanisms.

I-66. Hypothalamic control of self-initiated aggression seeking behavior

Annegret Falkner ANNEGRET.FALKNER@NYUMC.ORG
Dayu Lin DAYU.LIN@NYUMC.ORG

New York University

Aggressive motivation can be defined as the internal state that drives animals to seek out opportunities to perform aggressive actions. As with other motivational state behaviors, internally generated motivations may promote seeking behaviors that are independent of incoming sensory information. However, the neural substrates controlling aggressive motivation are not known. While the ventromedial hypothalamus, ventrolateral area (VMHvl) of male mice has recently been identified to have a critical role in promoting aggressive actions1, its role in aggression-seeking has been unclear2. In order to assay aggressive motivation, we have trained mice on a novel operant paradigm where animals self-initiate aggression “trials” by nosepoking for access to a social “reward” of a submissive male mouse. We find clear task learning in ~50% of trained mice and demonstrate bi-directional control of task response rate using both a pharmacogenetic and optogenetic approach. Reversible inactivation using the DREADD Gi system reduces task response rate for the social “reward” but not for a corresponding nonsocial reward (water). Optogenetic activation of VMHvl neurons or disinhibition of surrounding GABAergic

neurons decreases response initiation latency for single trials. To further determine the role of the VMHvl we recorded from populations of VMHvl neurons during this task and used PCA to identify neurons with preferred response selectivity during distinct task phases, including the aggressive social interaction, the operant response, and the delay period. We find that activity recorded during the interpoke interval of a subclass of VMHvl neurons is correlated with poke initiation latency, consistent with a model by which spikes accumulated during the inter-poke interval drive self-initiated aggression seeking behavior. Furthermore, cross correlation of the spontaneous activity of co-recorded VMHvl neurons demonstrates preferential coupling between different behaviorally defined classes of neurons such that information may be directionally passed through this circuit as behavior progresses from motivation to social action.

I-67. Population encoding of sound in central auditory neurons of the fly brain

Anthony Azevedo ANTHONY_AZEVEDO@HMS.HARVARD.EDU
Rachel Wilson RACHEL_WILSON@HMS.HARVARD.EDU

Harvard University

A common way of processing temporal information in a stimulus is to compute its spectral content, as the vertebrate auditory system does for sound. The Drosophila auditory system plays a central role in encoding the fly's courtship song and thus in guiding mating, but it is not known if temporal frequency is represented in the fly brain, nor how such a representation might arise without the mechanical filtering of a cochlea. Sound causes the fly's antenna to rotate, activating mechanosensitive neurons which turn project to the fly's brain. A key gap in our knowledge is how central circuits in the fly brain process and relay sound information and which neurons participate in these circuits. Recent work has identified a class of central neurons called AMMC-B1 neurons that respond to sound. Silencing AMMC-B1 neurons abolishes female receptivity, suggesting these neurons are a bottleneck in the auditory pathway. This work has prompted us to develop stimulation and recording techniques to understand the electrophysiology and sound encoding properties of these neurons. We have found that AMMC-B1 neurons are a diverse population, within which individual neurons have different responses to brief rotations of the antenna (steps), are sensitive to antennal oscillations (sinusoids) in distinct frequency bands, and exhibit diverse and unexpected intrinsic properties that likely underlie frequency selectivity. We have also recorded responses from 3rd order auditory neurons that likely receive input from AMMC-B1 neurons. These 3rd order neurons appear to use frequency tuned inputs to selectively represent specific features of courtship song. Our results show that the fly auditory system contains parallel channels that encode distinct sound features. This is an important step in understanding how neural computations implemented at the cellular and population level can support a behavior that is critical for the survival of the species.

I-68. Rapid path planning and preplay in maze-like environments using attractor networks

Dane Corneil DANE.CORNEIL@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Animals navigating in a well-known environment can rapidly learn and revisit observed reward locations, often after a single trial (Bast et al., 2009). The mechanism for rapid path planning remains unknown, though evidence suggests that the CA3 region in the hippocampus is important (Nakazawa et al., 2003), with a potential role for 'preplay' of navigation-related activity (Pfeiffer and Foster, 2013). Here, we consider an attractor network model of the CA3 region, and show how this model can be used to represent spatial locations in realistic environments with walls and obstacles. The synaptic weights in the network model are optimized for stable bump formation, so that neurons tend to excite other neurons with nearby place field centers and inhibit neurons with distant place field centers. Using these simple assumptions, we initialize the activity in the network to represent an initial location in the environment, and weakly stimulate the network with a bump at an arbitrary goal location. We find that, in networks representing large place fields, the network properties cause the bump to move smoothly from its initial location to the goal location along the shortest path, around obstacles or walls. Reward-modulated Hebbian plasticity during the first visit to a goal location enables a later activation of the goal location with a broad, unspecific external stimulus, representing input to CA3. These results illustrate that an attractor network that produces stable spatial memories, when augmented to represent large-scale spatial relationships, can be parsimoniously extended to rapid path planning.

I-69. Changes in striatal network dynamics during discrimination learning

Konstantin Bakhurin BAKHURIN@UCLA.EDU
Victor Mac VMAC1105@UCLA.EDU
Peyman Golshani PGOLSHANI@MEDNET.UCLA.EDU
Sotiris Masmanidis SMASMANIDIS@UCLA.EDU

University of California, Los Angeles

The striatum is thought to play a critical role in linking specific environmental cues with appropriate actions, enabling us to learn to guide our actions under different circumstances. However, it is still unclear how networks of striatal neurons are allocated to a specific task and how their coordinated network activity influences behavior. Our approach studies the striatal network by sampling spiking activity from large populations of striatal neurons. We used custom built 256 channel silicon microprobes to simultaneously record from over 100 striatal units in head-fixed mice learning a cue discrimination task. Among recorded striatal units, we detected a subset of putative medium spiny neurons (MSNs) and fast spiking interneurons (FSIs) that discriminated between an olfactory cue that predicted a reward (CS+), and a second cue that had no predictive value (CS-). We compared the network dynamics of discriminating and non-discriminating cells in two cohorts of mice corresponding to early and late stages of training. Early training was characterized by low behavioral discriminability (d'), with d' correlating with the fraction of task-responsive neurons in the striatum, but not the fraction of discriminating neurons. As expected, animals in the late stage of training had significantly higher behavioral discriminability, and intriguingly performance was correlated to the fraction of discriminating FSIs, but not MSNs. This suggests a learning-dependent change in striatal network dynamics in relation to behavioral discrimination. Furthermore, we studied the temporal relationship among simultaneously recorded units using noise and resting state correlations. We found that performance in the late, but not early stage of training was correlated to the coupling strength between discriminatory and non-discriminatory cells. Finally, we discuss the implications of these novel findings for the dynamic computational role of striatal circuits during the course of learning.

I-70. Auditory network optimized for noise robust speech discrimination predicts auditory system hierarchy

Fatemeh Khatami KHATAMI@ENGR.UCONN.EDU
Monty Escabi ESCABI@ENGR.UCONN.EDU

University of Connecticut

Sound recognition, including speech recognition, is accomplished by multiple hierarchical neural processing levels. Starting with the cochlea, sounds are decomposed into frequency components and subsequent auditory levels (e.g, cortex and midbrain) further decompose sound into spectro-temporal components. This multi-level processing scheme produces a robust neural representation enabling sound recognition in acoustically challenging environments. Yet how this high performance is accomplished is poorly understood. We develop a biologically inspired computational auditory network consisting of multiple hierarchical levels of excitatory and inhibitory spiking neurons and test its performance in a speech recognition task.We demonstrate that functional and structural properties of the optimal network mirrorhierarchies observed in the auditory pathway and show that this organization is essential for robust speech recognition. Upon optimizing the network to maximize discriminability of speech words, high performance is achieved even under variable conditions (e.g., multiple speakers). Furthermore, the optimal network shares structural and functional properties seen in the mammalian auditory system. The integration time constants of the network increase across the network layers such that the early layers are fast and deep layers are slow, analogous to transformation for temporal processing in the auditory system. At the early layers, connections between layers are largely restricted to nearby frequencies. Connectivity becomes progressively more divergent from layer-to-layer, analogous anatomical transformations observed between the auditory periphery and cortex. Finally, selectivity increases and spike rates decrease across the network layers producing a sparse representation that is analogous to transformations observed in the auditory system. Compared to a constrained network that lacks these structural features, the optimal network outperforms it in acoustically variable conditions (e.g., multiple speakers). Thus, the hierarchical organization and functional transformations of the network, and possibly the auditory system as a whole, are critical for achieving robust performance in variable acoustic conditions.

I-71. Thalamic inhibitory barreloids created by feedback projections from the thalamic reticular nucleus

Kazuo Imaizumi[1] KAZUO_IMAIZUMI@MSN.COM
Yuchio Yanagawa[2] YUCHIO@GUNMA-U.AC.JP
Charles C Lee[1] CCLEE@LSU.EDU

[1]Louisiana State University
[2]Gunma University

The thalamic barreloids are the obligate neural structures in the somasosensory thalamus that transmit ascending whisker information from the brainstem to the somatosensory barrel cortex. However, thalamic barreloids are much less understood because of the difficulty with delineating their structure in live experimental preparations. To address this fundamental problem, we have developed a live in vitro slice preparation that enables the straightforward identification of each barreloid for targeted physiological recordings. Our preparation takes advantage of transgenetically labeled vesicular GABA transporter (VGAT)-Venus axons originating from the thalamic reticular nucleus (TRN). We have used this novel slice preparation to understand the functional synaptic topography of the reticulothalamic circuit using laser-scanning photo-stimulation via uncaging glutamate. We found that the barreloids receive topographic inhibitory input from the TRN. Because it is generally accepted that inhibitory circuits in the central sensory systems develop later than excitatory circuits, we investigated the development of these inhibitory barreloids and quantified their developmental delay by comparing to excitatory barreloids formed by VGLUT2 receptors using immunofluorescence. We found that the inhibitory barreloids are developmentally established by postnatal day 7-8 (P7-8), which lag 4-6 days behind the establishment of excitatory barreloids formed by the ascending axons from the brainstem (P2-3). This strongly suggests that the critical period of structural plasticity in inhibitory barreloids substantially differ from our current knowledge (the critical period of the barrel formation closes by P4), which we are investigating through peripheral manipulations of structural plasticity. Overall, we propose a novel model system for studying developmental plasticity in the somatosensory thalamus using our transgenic slice preparation.

I-72. Effects of serotonin and dopamine manipulation on speed-accuracy trade-offs

Vincent Costa VINCENT.COSTA@NIH.GOV
Laura Kakalios LAURA.KAKALIOS@NIH.GOV
Eunjeong Lee EUNJEONG.LEE@NIH.GOV
Bruno Averbeck AVERBECKBB@MAIL.NIH.GOV

National Institute of Mental Health

The neural systems that mediate speed-accuracy trade-offs are largely unknown. To explore the neural systems that contribute to the speed accuracy trade-off, we manipulated dopamine and serotonin signaling in monkeys and compared their performance to sessions in which the animals were administered saline while they performed a perceptual inference task. In the task, the color of each pixel was independently chosen to be either read or blue according to an underlying color bias which takes on values from 0.5 to 0.7. The colors of the pixels are updated every 100 ms, but always according to the color bias. In other words, the subjects may be viewing a stimulus which is 65% blue pixels, on average. Their task is to indicate whether the stimulus contains more blue or more red pixels. Because the colors of the pixels are updated independently every 100 ms, performance is more accurate when subjects spend more time observing the stimulus. Dopamine was manipulated by systemically injecting GBR-12909 (1.3 mg/kg). This drug blocks the dopamine transporter (DAT), and leads to larger phasic and tonic dopamine responses. Serotonin was correspondingly manipulated by injecting escitalopram (1.0 mg/kg). This drug blocks the serotonin transporter, and therefore also leads to larger levels of serotonin, although less is known about the exact dynamic effects on phasic vs. tonic serotonin. We found that increasing dopamine levels shifts animals to faster responses with a corresponding decrease in accuracy. Further modeling will assess whether the rate of information processing is approximately preserved in this condition. When serotonin levels are increased, however, animals appear to have a decrease in the rate of information processing, because their accuracy is lower at corresponding reaction times.

I-73. Scaling properties of dimensionality reduction for neural populations and network models

Ryan Williamson[1] RCW30@PITT.EDU
Benjamin Cowley[2] BCOWLEY@ANDREW.CMU.EDU
Ashok Litwin-Kumar[3] AK3625@COLUMBIA.EDU
Brent Doiron[1] BDOIRON@PITT.EDU
Adam Kohn[4] ADAM.KOHN@EINSTEIN.YU.EDU
Matt Smith[1] SMITHMA@PITT.EDU
Byron Yu[2] BYRONYU@CMU.EDU

[1]University of Pittsburgh
[2]Carnegie Mellon University
[3]Columbia University
[4]Albert Einstein College of Medicine

Recent studies have used dimensionality reduction techniques to study neural population activity during decision-making, motor control, olfaction, working memory, and other behavioral tasks. To date no study has systematically investigated how the outputs of dimensionality reduction scale with number of recorded neurons and trials. Such a study would indicate whether the results obtained from a limited sampling of a neural circuit reliably represents activity in the larger circuit. To this end, we applied factor analysis to spontaneous activity recorded in visual area V1 of monkeys. We found that dimensionality increases proportional to the number of neurons and trials. In contrast, we found that percent shared variance remains constant over a wide range of neurons and trials. To extrapolate these results to larger numbers of neurons and trials, we investigated whether network models of spiking activity with a balance between excitation and inhibition show similar scaling properties as biological networks.

We found that, given a large number of trials, both non-clustered and clustered balanced networks showed similar scaling properties up to the number of neurons recorded in the biological network. However, the two models make different predictions about scaling beyond this experimental regime, and will require recordings from larger numbers of neurons before further comparisons with biological networks can be made. These results suggest that sampling tens of neurons and hundreds of trials is sufficient for identifying how much of a neuron's activity co-fluctuates with other neurons in the larger network. However, depending on the type of network connectivity structure, sampling additional neurons and trials can provide a richer description of the co-fluctuations across the network.

I-74. Learning in STDP networks of neurons by a sequence of spatiotemporally patterned stimuli

Ho Jun Lee SNIPER0712@KOREA.AC.KR
June Hoan Kim LOVEPHY@GMAIL.COM
Kyoung Lee KYOUNG@KOREA.AC.KR

Korea University

'Spike-timing dependent plasticity' is perhaps the most important concept as the learning mechanism of biological neural systems. Although the original inception of STDP was based on a system of only two neurons, most learning process occurs in large populations of neurons. Synchronized bursts are a popular dynamic event occurring in coupled networks of neurons in general, and it's been suggested that they greatly improve the efficacy of synaptic modification [Froemke et al., 2006, Harnett et al., 2009]. So, we set out to explore how spatiotemporal patterns of extrinsic electrical stimulations can alter the bursting dynamics in cultured networks of neurons that are affecting each other presumably through some STDP mechanism [Choi et al., 2012]. The extrinsic stimulation is a repeated sequence of paired electrical pulses delivered to two different subgroups of neurons with a Δt time delay. Interestingly, the arrival times of 'recurrent burst', which is the second burst following the very first burst that was incurred immediately by a single probing pulse were significantly precise only for training with some particular values of Δt [Fig.1]. Subsequently, we recapitulated most of our experimental observations very well, employing the well-established Izhikevich neural network model, whose synaptic weight landscape evolves in time based on a STDP function. We also found that depending on the subgroups of cells chosen for stimulations and the specific value of Δt, a particular subset of neurons formed a synaptically facilitated strong sub-network [Fig.2]. The morphology of the enhanced subnetwork was also strongly related to the overall excitability of the system. We also found that the two most important factors governing the existence of enhanced subnetwork are the degree of dispersion in the conduction time delay from the cells that are receiving the stimulations directly and the number of strong synaptic weights attached to the stimulated cells [Fig.3].

I-75. Learning temporal integration in balanced networks

Erik Nygren NYGREN@PYL.UNIBE.CH
Robert Urbanczik URBANCZIK@PYL.UNIBE.CH
Walter Senn SENN@PYL.UNIBE.CH

University of Bern

Neural integrators have been recognized to play an important role in many different functional areas of the brain such as the oculomotor control, head direction and in short term memory. The experimental evidence for neural integrators has lead to intensive theoretical studies in the last decades, and many promising models have come forth. Whereas earlier models often suffered from stability issues, more recent works often overcome these through different mechanisms such as balanced inhibition or bistable neurons. How the specific structures for stabilization could arise in nature through learning is not fully understood. In this work we propose a biologically realistic model which learns stable integration through a combination of different coding paradigms observed in experimental data. On one hand information is coded in the mean firing rate of compartmental neurons and on the other hand inhibitory neurons are recruited to enhance capacity and stability when information is accumulated. We show that the recruiting behavior of the network is a direct consequence of the log-normal distribution of synaptic weights observed experimentally. Furthermore we use a simple gradient based learning for spiking compartmental neurons to achieve the necessary network structure for robust integration. Feedforward and recurrent connections projecting to the dendritic compartment are plastic and the learning paradigm for temporal integrations consist of providing the derivative of the teacher signal through the feedforward connections to the dendrite while the somatic compartment is presented the teaching signal. In this way a network structure is formed such that external input is integrated to persistent activity of the compartmental neurons, while inhibitory neurons are recruited as the stored signal grows. In absence of external stimuli, the balance between the inhibitory and excitatory neurons keeps the network stable. This model provides a realistic explanation for how integrator networks can emerge in biological networks through learning.

I-76. When feedforward nets pretend to oscillate

Alexandre Payeur APAYE026@UOTTAWA.CA
Leonard Maler LMALER@UOTTAWA.CA
Andre Longtin ALONGTIN@UOTTAWA.CA
University of Ottawa

In the context of neural coding, the dynamics of a neuronal network is specified by the sequence of times at which its neurons spike in response to a given stimulus. The structure of these spike sequences depends on single-neuron and synaptic properties, as well as on the network's connectivity. Here, we study how important is this connectivity in determining a network's dynamics. To this end, we compared the spiking activity of a delayed recurrent network of inhibitory integrate-and-fire neurons driven by a correlated input with that of a feedforward network having the same properties and input. Under these circumstances, the recurrent net is known to display oscillations. When the networks were local (i.e. when axonal delays were short) and the neurons intrinsically noisy, the postsynaptic neurons of the feedforward net possessed spike-train statistics very similar to those of the recurrent neurons. Using linear response theory, we showed that the feedforward net may be seen as a first order approximation of the recurrent network. We also found that heterogeneous axonal delays in the feedforward net may accentuate its resemblance to the recurrent net. We conclude that a strict connectomic analysis of a network is not a sufficient constraint for its dynamics. We also advocate detailed spike-train analyses to assess the properties of neural activity, especially in the context of brain rhythms.

I-77. SpikeFORCE: A FORCE algorithm for training spike timing based neuronal activity

Ran Rubin RR2980@COLUMBIA.EDU
Haim Sompolinsky HAIM@FIZ.HUJI.AC.IL
Larry Abbott LFA2103@COLUMBIA.EDU
Columbia University

Neurons encode and convey information through the timing of their action potentials, but most models of neural learning have focused on training firing rates. Recent theoretical work has characterized the types of computations neurons can perform using spike-timing information and how model neurons can be trained to produce spike-timing based codes. However, because of the nonlinearity associated with the post-spike potential reset in spiking neurons, existing training algorithms for implementing precise spike-time-based input-output transformations in recurrent networks suffer from either low convergence rates or stability issues. For rate-based models of recurrent neural networks, the FORCE learning algorithm allows networks to implement complex computations and harness the computational powers of randomly connected recurrent neural networks to produce desired firing-rate patterns in an efficient, robust and stable fashion (Sussillo and Abbott, 2009; Laje and Buonomano, 2013). The main idea behind FORCE learning is to keep the error of the network, the mismatch between the output of the network and the desired output, small at all times by rapidly changing weights in the network. SpikeFORCE is a FORCE-inspired algorithm for training a spike-timing-based, rather than a rate-based, neuronal response. SpikeFORCE can be applied to spiking neurons in feedforward and recurrent networks to entrain desired spike-timing relationships with a pre-assigned precision. It is online, relatively fast and converges efficiently. Training is done by using FORCE like learning to ensure the neurons' membrane potentials cross threshold only at the allowed desired times but without enforcing a specific membrane potential trace. Importantly, SpikeFORCE can automatically construct robust solutions that are resistant to noise and to the harsh nonlinear effects of threshold crossing dynamics. SpikeFORCE can by used to train various recurrent architectures, and it provides an approach for improving our understanding of the computational power and limitations of spiking-timing effects in neural networks.

I-78. Divisive inhibition of rate response via input timescale

Timothy D Oleskiw OLESKIW@UW.EDU
Wyeth Bair WYETH.BAIR@GMAIL.COM
Eric Shea-Brown ETSB@UW.EDU

University of Washington

The divisive inhibition of a neuron's firing rate is considered a fundamental operation in many neural computations, including gain control and normalization. However, the precise mechanisms underlying divisive inhibition are not completely understood, making it difficult to predict when this phenomenon will occur in spiking network models. Research has shown that balanced synaptic input to a neuron can divisively scale its spike rate response, both in vitro and in conductance-based leaky-integrate-and-fire (LIF) models. Interestingly, a review of analytic solutions to similar models driven by stochastic input show response inhibition to depend strongly on the input autocorrelation timescale. To better understand this, we present numeric and analytic results exploring the effects of input timescale on divisive inhibition of the rate response in a model LIF neuron. First, simulations demonstrate that both the variance and timescale of injected current contribute to the slope of the rate response curve. We then derive estimates for the response of simplified threshold neurons driven by coloured Gaussian noise, solved exactly in a stationary case. These findings suggest rate response division to arise from an interaction between the input's timescale and the neuron's filtering properties, i.e. leak current and post-spike refractory effects. Further, we show numerically that by increasing synaptic timescales across a biophysically plausible range, firing rates are divisively inhibited. Surprisingly, this effect occurs even while the level of balanced synaptic activity and net conductance is held constant. These findings are discussed in the context of spiking models for normalization, with preliminary data illustrating how an input's timescale can be shaped by correlating its excitatory and inhibitory synaptic events. Together, our results demonstrate input timescale as a novel mechanism for the divisive inhibition of firing rates, and provide insight into the implementation of effective spiking models for gain control and normalization.

I-79. Extracting latent structure from multiple interacting neural populations

Joao D Semedo[1,2] JSEMEDO@CMU.EDU
Amin Zandvakili[3] AMIN.ZANDVAKILI@EINSTEIN.YU.EDU
Adam Kohn[3] ADAM.KOHN@EINSTEIN.YU.EDU
Christian Machens[4,5] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Byron Yu[1] BYRONYU@CMU.EDU

[1]Carnegie Mellon University
[2]Instituto Superior Tecnico
[3]Albert Einstein College of Medicine
[4]Champalimaud Neuroscience Programme
[5]Ecole Normale Superieure

Developments in neural recording technology are rapidly enabling the recording of populations of neurons in multiple brain areas simultaneously, as well as the identification of the types of neurons being recorded (e.g., excitatory vs. inhibitory). As a result, there is a growing need for statistical methods to study the interaction among multiple, labeled populations of neurons. Rather than attempting to identify direct interactions between neurons (where the number of interactions grows with the number of neurons squared), we propose to extract a smaller number of latent variables from each population and study how these latent variables interact. Common methods for extracting latent variables from neural population activity, such as factor analysis (FA), are not tailored for the study of interactions between populations. We propose using probabilistic canonical correlation analysis (pCCA), which can be interpreted as the multivariate extension of pairwise correlation. We applied FA and pCCA to populations of neurons recorded simultaneously in visual areas V1 and V2 of an anesthetized macaque monkey and found that the between-population covariance has lower dimensionality than the within population covariance. The between-population interaction also shows a rich temporal structure, as the two areas are correlated at multiple time delays throughout the trial. In order to further explore this temporal structure, we propose an extension to pCCA, termed group latent auto-regressive analysis (gLARA). gLARA captures the temporal structure of the population activity by defining an auto-regressive model on the latent variables, while still distinguishing within-population dynamics from between- population interactions. We found that gLARA was better able to capture the structure in the V1-V2 population recordings than pCCA. Moreover, it best describes the population activity when the auto-regressive model uses the past 15 ms of activity, suggesting that the most significant interactions between these areas may occur within this time window.

I-80. Interaction of slow network integration and fast neural integration towards spike generation

Annabelle Singer[1] ASINGER@MIT.EDU
Giovanni Talei Franzesi[1] GEGGIO@MIT.EDU
Suhasa Kodandaramaiah[1] SUHASABK@MEDIA.MIT.EDU
Francisco Flores[1] FJFLORES76@GMAIL.COM
Craig Forest[2] CRAIG.FOREST@ME.GATECH.EDU
Nancy Kopell[3] NK@MATH.BU.EDU
Ed Boyden[1] ESB@MEDIA.MIT.EDU

[1]Massachusetts Institute of Technology
[2]Georgia Tech
[3]Boston University

Models of neural code generation have often focused on the brief membrane time constant as a fundamental duration over which subthreshold information is integrated by a neuron towards spike threshold. Here we report that neurons in hippocampal area CA1 and barrel cortex of awake mice exhibit gradual rises in voltage, lasting up to hundreds of milliseconds, before the brief rises that precede spikes. Fast rises rarely yielded spikes unless

they were preceded by a gradual rise. Spike occurrences could be in part predicted by the presence of a gradual rise alone. Robotic whole-cell recording of pairs of neurons show that gradual rises are often coordinated across nearby cells, whereas fast rises are cell-specific. This suggests that slower, network-level integration might complement the faster, classical integration within single neurons to mediate the generation of spike patterns. These coordinated gradual rises could serve a computational role by helping pre-select neurons within a network to fire in response to subsequent inputs.

I-81. Excitatory to inhibitory connectivity shaped by synaptic and homeostatic plasticity

Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK
Jacopo Bono J.BONO13@IMPERIAL.AC.UK
Ulysse Klatzmann ULYSSEKLATZMANN@GMAIL.COM

Imperial College London

Recent experimental techniques allowed to study the relationship between neurons' stimulus-preference and connectivity. In particular, in the layer II/III of primary visual cortex, it was shown that excitatory neurons with the same orientation preference have a high probability of being bidirectionally connected (Ko et al. Nature 2011). However, the intracortical connectivity is only getting refined after eye-opening (Ko et al. Nature 2013). We have recently hypothesized that this process is a result of experience-dependent plasticity, modelled by a Hebbian learning rule (Clopath et al. 2010, Ko et al. Nature 2013). In contrast to excitatory neurons, parvalbumin-expressing (PV) inhibitory cells are less input-specific (Hofer et al. Nat. Neur. 2011, Bock et al. Nature 2011): Hofer et al. showed that PV neurons receive excitatory inputs from neurons with different orientation preferences. In this work, we investigate the mechanism by which excitatory to inhibitory connections are formed (how) and their potential function (why) in a small recurrent network. We found that a model combining Hebbian learning with homeostatic plasticity, which allows PV neurons to spike at a high rate (i.e reproducing the fast-spiking intrinsic property of the cells), develops unspecific excitatory-to-inhibitory connections. We then tested the role of inhibition by simulating our model with and without inhibition after learning convergence. We found that inhibition ensures less fluctuation of the synaptic weights over time, hence stabilizes the network. We therefore propose that unspecific excitatory to PV connections can be a result of the intrinsic homeostatic property of PV neurons, and can allow the network to be more stable.

I-82. High degree neurons tend to contribute more and process less information in cortical networks

Nicholas Timme[1] NMTIMME@UMAIL.IU.EDU
Shinya Ito[2] SHIXNYA@GMAIL.COM
Maxym Myroshnychenko[1] MMYROS@GMAIL.COM
Fang-Chin Yeh[1] FANYEH@INDIANA.EDU
Emma Hiolski[2] EHIOLSKI@GMAIL.COM
Alan Litke[2] ALAN.LITKE@CERN.CH
John Beggs[3] JMBEGGS@INDIANA.EDU

[1]Indiana University
[2]University of California, Santa Cruz
[3]Indiana University Bloomington

Understanding the computations performed by small groups of neurons is of great importance in neuroscience. To investigate such computations, we used tools from information theory (transfer entropy and the partial information decomposition) to study information processing in timescale-dependent effective connectivity networks (i.e. mul-

tiplex neural networks). These networks were derived from the spiking activity of thousands of neurons recorded from 25 cortico-hippocampal slice cultures using a high-density 512-electrode array with 60 micrometer inter-electrode spacing and 50 microsecond temporal resolution. To the best of our knowledge, this preparation and recording method combination yields recorded neuron quantities and temporal and spatial recording resolutions that are not currently available in any in vivo recording system. By utilizing transfer entropy - a well established method for detecting linear and nonlinear interactions in time series - and the partial information decomposition - a powerful, recently developed tool for dissecting information processing into novel parts - we were able to move beyond important but well traveled questions about network topology to novel questions about how neurons process information in these networks. Given the field's great deal of interest in high degree neurons (neurons with many connections; so called “hubs"), we chose to examine the relationship between neuron degree and information processing. For interactions that ranged in time from 1.6 to 300 ms, we found that high degree neurons tended to contribute more information and process less information than low degree neurons. For slower interactions ranging from 350 to 3000 ms, we found that information contribution was generally uncorrelated with number of connections, but that high degree neurons tended to process more information than low degree neurons. These results provide new insights into the relationship between network topology and information processing in neural networks.

I-83. Enhancement and modelling of spike timing reliability in-vivo using noise evoked by juxtacellular stimulation

Guy Doron[1] GUY.DORON@CHARITE.DE
Jens Doose[2] JENS.DOOSE@BCCN-BERLIN.DE
Michael Brecht[2] MICHAEL.BRECHT@BCCN-BERLIN.DE
Benjamin Lindner[2] BENJAMIN.LINDNER@PHYSIK.HU-BERLIN.DE

[1]Humboldt University Berlin
[2]BCCN Berlin

We have recently shown that the juxtacellular nanostimulation technique can be used to parametrically control spike frequency and number in identified single cortical neurons in vivo (Houweling et al., 2010). Specifically, we found that spike number in barrel cortex neurons varies linearly both with stimulus intensity and stimulus duration, using step current injections. However, using this method we were not able so far to achieve spike timing reliability. Here we are extending these findings and show that driving pyramidal cells in anesthetized rat vibrissal motor cortex with fluctuating stimuli (frozen bandpass-limited white noise) results in increased spike timing reliability. Specifically, we report that parametrically increasing the nanostimulation noise level results in increased spike train synchronization. In addition, it results in increased cross spectra between the stimulus and the spike trains, which leads to increased coherence and change in the coherence shape. We further explore how well the spike train in response to this stimulus can be captured by an exponential integrate-and-fire neuron (Fourcaud-Trocme et al., 2003), a simple model that has been successfully applied for reproducing spike times of pyramidal cells under noisy current stimulation in vitro (Badel et al., 2008). In contrast to the latter situation, our model also includes an appreciable amount of intrinsic noise, accounting for fluctuating input from the surrounding network. Nanostimulation therefore permits enhanced control of spike timing in single cortical neurons and therefore holds great potential for elucidating how spike timing reliability in single neurons may contribute to behavior.

I-84. Enhancing spike inference precision by combining independent trials of calcium imaging recordings

Josh Merel[1] JSM2183@COLUMBIA.EDU
Eftychios Pnevmatikakis[2] EPNEVMATIKAKIS@SIMONSFOUNDATION.ORG
Kalman Katlowitz[3] AKATLOWITZ@GMAIL.COM
Michel Picardo[3] MICHEL.PICARDO@NYUMC.ORG
Michael Long[3] MLONG@MED.NYU.EDU
Liam Paninski[1] LIAM@STAT.COLUMBIA.EDU

[1]Columbia University
[2]Simons Foundation
[3]New York University

Calcium imaging methods enable monitoring of large population of neurons in neuroscience. However, typical raster scanning approaches suffer from a trade-off between the imaging frame rate, the measurement signal-to-noise ratio and the spatial area/resolution that is monitored. Moreover, the spiking signal is observed indirectly through the concentration of the calcium bound chemical indicators that have significantly slower dynamics. As a result, when spike inference is performed in discrete time, the resolution is dictated by the frame rate, whereas continuous time estimates tend to exhibit high uncertainty at a finer-precision, thus limiting our ability to perform spike time super-resolution. In this work we show how this problem can be alleviated in the case that we know a-priori that the imaged neurons exhibit high temporal precision. We extend our previous work on Bayesian spike inference and present a Bayesian approach that combines multiple independent trials of the same experiment. Each spike can be described by a global random variable that is common across trials, plus a trial dependent jitter that is modeled by a zero mean Gaussian distribution. We then use a Gibbs algorithm to sample over the posterior distribution of the shared spike times, trial dependent jitters, jitter variance, as well as other model parameters (baseline, initial concentration, spike amplitude, and observation noise) that can also be trial dependent. We apply this method to calcium imaging recordings from neurons in the HVC area of the zebra finch, where neurons exhibit bursts during singing that are remarkable precise across multiple song repetitions. Our method shows that combining N trials leads to an order square root of N increase in the precision of spike timing inference, comparable to the precision previously observed with standard electrophysiology recordings.

I-85. Fast multi-neuron spike identification via continuous orthogonal matching pursuit

Karin Knudson[1] KARIN.C.KNUDSON@GMAIL.COM
Jacob Yates[1] JLYATES@UTEXAS.EDU
Alex Huk[1] HUK@UTEXAS.EDU
Jonathan W Pillow[2] PILLOW@PRINCETON.EDU

[1]University of Texas at Austin
[2]Princeton University

Electophysiological recordings of spiking neurons can be represented as a sparse sum of continuously translated and scaled copies of stereotyped spike waveforms. Here we present a fast method for estimating the identities, amplitudes, and continuous times of spike waveforms, given the aggregate signal and a waveform shape for each neuron. The method, which we call Continuous Orthogonal Matching Pursuit (COMP), proceeds by iteratively selecting a set of basis vectors for representing one spike waveform over some continuous range of times, and then refining estimates of all spike amplitudes and times, alternating between steps in a process analogous to the well-known Orthogonal Matching Pursuit (OMP) algorithm. Our approach for modeling continuous translations builds on the Continuous Basis Pursuit (CBP) algorithm [Ekanadham et al., 2011], which we extend in several ways: by selecting a basis that optimally approximates translated copies of the waveforms, replacing convex optimization with a greedy optimization approach, and moving to the Fourier domain to more accurately estimate

continuous time shifts. Using both simulated and real neural data, we show that COMP gracefully resolves multiple superimposed spike waveforms and achieves gains over CBP in both speed and accuracy.

I-86. Statistical assessment of sequences of synchronous spiking in massively parallel spike trains

Emiliano Torre[1] E.TORRE@FZ-JUELICH.DE
Carlos Canova[1] C.CANOVA@FZ-JUELICH.DE
George Gerstein[2] GEORGE@MULAB.PHYSIOL.UPENN.EDU
Moritz Helias[1] M.HELIAS@FZ-JUELICH.DE
Michael Denker[1] M.DENKER@FZ-JUELICH.DE
Sonja Grün[1] S.GRUEN@FZ-JUELICH.DE

[1]INM-6 / IAS-6, Research Center Juelich
[2]University of Pennsylvania

Propagation of spike synchronization across groups of neurons has been hypothesized as a mechanism of information processing and transmission (Abeles M, 1991). From a theoretical standpoint, the possibility to build biologically realistic neural networks that enable stable propagation of volleys of synchronous spikes (synchronous events) from one group of neurons to the next has been confirmed by the synfire chain model (Abeles M, 1991; Diesmann M et al, 1999; Trengove C et al, 2013). However, to date there is no experimental evidence of sequences of synchronous events (SSEs) in the cortex. Recent technological advances to record from hundreds of cells simultaneously increase the chance to observe such concerted activity. Still, its quantitative assessment requires suitable statistical tools. We here propose a statistical method to detect repeated SSEs in massively parallel spike trains. The method builds on a previous technique to identify repeated SSEs as diagonal structures in an intersection matrix, whose entries quantify the normalized overlap in neuron identities between any two time bins (Schrader S et al, 2008; Gerstein GL et al, 2012). Our method replaces the normalized overlap with its cumulative probability under the hypothesis of independence. The method then tests for each entry that a) the entry is not statistically significant under the hypothesis of independence, and b) it is not part of a significant diagonal structure, given the observed zero-lag correlations. Upon rejection of both null hypotheses, the method identifies repeated SSEs and assesses their neuronal composition. We calibrate the analysis with stochastic simulations containing SSEs embedded in background spiking activity. In biologically realistic scenarios large portions of the sequences are identified, with low false positive and false negative levels. We demonstrate robustness against various types of spike train non-stationarity and conclude with preliminary results on electrophysiological recordings.

I-87. Non-linear input summation in balanced networks can be controlled by activity dependent synapses

Sara Konrad SARA.KONRAD@BRAIN.MPG.DE
Tatjana Tchumatchenko TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

Max-Planck Institute for Brain Research

The current inputs to cortical neurons are under many circumstances highly variable. This situation arises due to a balance of net excitatory and inhibitory currents which leads to strongly fluctuating sub-threshold neuronal inputs which in turn generate highly irregular firing. Spiking network models, consisting of sparsely and randomly connected excitatory and inhibitory neurons with external drive, have been shown to reproduce this irregular cortical activity [Van Vreeswijk and Sompolinsky, 1996]. In the modeled balanced state, the relation between external input strength and network activity is always linear. However, in order to perform more complex computations like normalization, classical receptive fields or surround suppression the network response needs to be non-linear.

So, how can non-linear network responses be implemented and controlled in spiking balanced networks? To address this question, we extended the mean-field theory for spiking balanced networks with synapses displaying short-term synaptic plasticity published by [Mongillo et al., 2012] and show that the response of balanced networks with activity dependent synapses to external inputs is in general non-linear. When depression and facilitation act on similar time-scales, two stable activity states occur for a wide range of external inputs. We show, that the summation properties in these two stable states differ significantly and exhibit supralinear summation in the down state and sublinear summation in the up state. We therefore conclude that input summation rules as well as system noise, can be controlled and changed drastically for the same input by choosing one of the two possible states. We define and analytically derive the degree of linear summation for general activity dependent synaptic transmission functions in steady state. Furthermore, we corroborate analytical predictions by detailed large scale spiking neural network simulations.

I-88. Synaptic volatility and the reorganization of electrical activity in neuronal networks

Gianluigi Mongillo[1] GIANLUIGI.MONGILLO@GMAIL.COM
Simon Rumpel[2] SIRUMPEL@UNI-MAINZ.DE
Yonatan Loewenstein[3] YONATAN@HUJI.AC.IL

[1]CNRS UMR8119, Descartes University
[2]Johannes Gutenberg University Mainz
[3]The Hebrew University of Jerusalem

The pattern and strengths of connections between neurons determine the network response to external inputs. However, there is substantial volatility of the excitatory connections in the cortex, even in the absence of any explicit learning. Comparable volatility in vitro, even when neural activity is blocked, indicates that dynamic remodeling is a fundamental feature of neocortical circuits. The objective of this study is to assess the effect of this network reorganization on the spatial organization of firing rates of neurons, as the latter can serve as a proxy to the effect of network reorganization on its response properties. We chronically imaged thousands cortical spines in the adult mouse and report substantial volatility: most spines present in the first imaging day were no longer present 20 days later, and most of the stable spines changed their size by at least a factor of two in that period of time. We used these excitatory-excitatory connectivity measurements in simulations of large networks of integrate and fire neurons. Surprisingly, we found that the substantial network volatility had only a small effect on the spatial distribution of the firing rates of the neurons. To understand these results, we developed an analytically-tractable mean-field theory that relates connectivity dynamics to changes in the pattern of spiking activity in balanced networks. This theory shows that in biologically-plausible balanced networks, inhibitory connections have a substantially larger contribution to the heterogeneity in the input to different neurons than excitatory connections. Therefore, the spatial organization of spontaneous firing rates and the network response properties are robust to the remodeling of the excitatory network but sensitive to changes in the inhibitory connections. These results also indicate that the often-overlooked inhibitory plasticity has a far larger potential in changing the functionality of cortical networks than the extensively-studied excitatory plasticity.

I-89. From structure to dynamics: origin of higher-order spike correlations in network motifs

Yu Hu[1] HUYUPKU@GMAIL.COM
Kresimir Josic[2] JOSIC@MATH.UH.EDU
Eric Shea-Brown[3] ETSB@UW.EDU
Michael A Buice[4] MICHAELBU@ALLENINSTITUTE.ORG

[1]Harvard University
[2]University of Houston
[3]University of Washington
[4]Allen Institute for Brain Science

An overarching goal in computational neuroscience is to understand the network origins of collective neural activity. Here, we take a statistical approach to both network structure and network dynamics. Higher order correlations between spike trains of neurons have attracted significant interest, both in improving statistical descriptions of population activity and in hypothesized functional roles in sensory coding and synaptic plasticity. Our work combines two recent advances: (1) methods for comput- ing the correlations between neurons in recurrent networks, and (2) methods for isolating contributions of small network motifs on global correlation structure. Our approach provides novel relationships between the statistics of connectivity with the statistics of correlations across a neural network: First, we compute the third order correlations in a general class of GLM type models, including a non-linearity in the firing rate function of the neurons, using the method of path integrals. Second, we develop expansions for such higher order correlations in terms of statistics of connectivity as measured by motif cumulants and show that these form a good approximation for the correlations for several network architectures. These results expand on previous work [Pernice et al 2011, Hu et al 2014, Jovanovic et al 2014] by including both the effects of nonlinearities and computing additional terms to the higher order correlations which arise due to recurrence and are valid for both excitatory and inhibitory interactions. We give numerical results that illustrate the rapid convergence of the motif expan- sion, enabling predictions based on a tractable number of network motifs. These results provide predictions which connect two large scale data sets soon to be collected at the Allen Institute for Brain Science, dense reconstruction of the synaptic connectivity and optical recordings of the same neural population in vivo.

I-90. A shotgun sampling solution for the common input problem in neural connectivity inference

Daniel Soudry DANIEL.SOUDRY@GMAIL.COM
Suraj Keshri SKK2142@COLUMBIA.EDU
Patrick Stinson PATRICKSTINSON@GMAIL.COM
Min-hwan Oh MO2499@COLUMBIA.EDU
Garud Iyengar GARUD@IEOR.COLUMBIA.EDU
Liam Paninski LIAM@STAT.COLUMBIA.EDU

Columbia University

Inferring connectivity in neuronal networks remains a key challenge in statistical neuroscience. The “common input” problem presents the major roadblock: it is difficult to reliably distinguish causal connections between pairs of observed neurons from correlations induced by common input from unobserved neurons. Since available recording techniques allow us to sample from only a small fraction of large networks simultaneously with sufficient temporal resolution, naive connectivity estimators that neglect these common input effects are highly biased. This work proposes a “shotgun” experimental design, in which we observe multiple sub-networks briefly, in a serial manner. Thus, while the full network cannot be observed simultaneously at any given time, we may be able to observe most of it during the entire experiment. Using a generalized linear model for a spiking recurrent neural network, we develop scalable approximate Bayesian methods to perform network inference given this type of data, in which only a small fraction of the network is observed in each time bin. We demonstrate in simulation

that, using this method: (1) Networks with thousands of neurons, in which only a small fraction of neurons is observed in each time bin, could be quickly and accurately estimated. (2) Performance is improved if we exploit prior information about the probability of having a connection between two neurons, its dependence on neuronal cell types (e.g., Dale's law) or its dependence on the distance between neurons. (3) the shotgun experimental design can eliminate the biases induced by common input effects.

I-91. Realistic structural heterogeneity in local circuits and its implications on the balanced state

Itamar Landau ITAMAR.LANDAU@MAIL.HUJI.AC.IL
Haim Sompolinsky HAIM@FIZ.HUJI.AC.IL

The Hebrew University of Jerusalem

Traditional analysis of cortical network dynamics has treated simplistic structure — most commonly model networks have been constructed with a uniform probability of connection between each pair of cell types. Here we characterize some of the salient features of the connectivity of real local cortical circuits and show that they have significant impact on their dynamical state. First, we present estimates from anatomical data in Layer 4 rat barrel cortex of the spatial variance in the number of local connections. We show that the variance in the number of incoming connections is large: it is of the same order as the square of the mean rather than the mean as expected for simple random networks. Using mean-field theory as well as simulations we show that networks with such large structural heterogeneity do not admit a global balance state solution in which the mean excitation and inhibition dynamically balance to yield fluctuation-driven firing. Rather, the input is locally unbalanced, the network fires extremely sparsely with exceedingly long-tailed rate distributions, and active cells fire with high temporal regularity. In order to recover realistic firing rate distributions, we construct a network model with a strong adaptation current in addition to heterogeneous structure. Using mean field theory, exact in the asymptotic large network limit, we find a novel dynamical state in which the structural heterogeneity is balanced locally by adaptation. This adaptation-driven local balance yields realistic distribution of firing rates. Applying this theory, we simulate the full layer 4 population with realistic connectivity structure, and find that with a moderate adaptation current this network is able to generate realistic rate distributions in both spontaneous and stimulus-evoked states. Our work shows that features of the realistic cortical connectivity as estimated from anatomical data have substantial qualitative impact on their dynamic properties.

I-92. Dynamics of recurrent networks in visual cortex with multiple inhibitory subpopulations

Ashok Litwin-Kumar[1] AK3625@COLUMBIA.EDU
Robert Rosenbaum[2] RROSENB1@ND.EDU
Brent Doiron[3] BDOIRON@PITT.EDU

[1]Columbia University
[2]University of Notre Dame
[3]University of Pittsburgh

Theoretical studies have typically treated inhibitory interneurons as a homogeneous population, but recent advances in the characterization of interneuron subtypes in mouse V1 reveal heterogeneous functional and connectivity properties. However, a consistent theoretical framework in which to investigate the dynamics of neuronal networks with multiple classes of inhibitory neurons has so far been lacking. We approach this problem by studying the dynamics of recurrent networks with connectivity properties constrained by experimental data from mouse V1. Our theory describes conditions under which the dynamics of such networks are stable and how perturbations of distinct neuronal subtypes recruit activity changes through recurrence. We apply these conclusions to study disinhibition, surround suppression, and subtraction or division of orientation tuning curves in a model network. Our calculations and simulations determine the conditions under which activity consistent with experiment is possible. They also lead to predictions concerning connectivity and network dynamics that can be tested via optogenetic manipulations. Finally, they illustrate that the recurrent dynamics of these networks must be taken into account to fully understand many effects reported in the literature.

I-93. What regime is the cortex in?

Yashar Ahmadian YAHMADIAN@GMAIL.COM
Kenneth Miller KEN@NEUROTHEORY.COLUMBIA.EDU

Columbia University

Cortical neurons fire irregularly and are largely uncorrelated. Asynchronous irregular activity arises when firing is driven by input fluctuations around a sub-threshold mean input. Defining threshold $= 1$, a key theoretical question is how both mean input and fluctuations can be $O(1)$, so that a reasonable amount of irregular activity arises. The seminal "balanced network" model suggested scaling synaptic strengths as $1/\sqrt{K}$, for K presynaptic neurons, and focused on the very large K limit. This yields $O(1)$ fluctuations but large, $O(\sqrt{K})$ mean recurrent excitatory or inhibitory inputs. Given a similarly large external input, the neurons' mean excitatory and inhibitory inputs precisely balance dynamically, leaving a net $O(1)$ mean input. However, this scenario has two problems. First, in this regime, network responses depend linearly on stimuli, whereas cortical responses involve many nonlinearities. Second, the large external drive is critical to achieving precise balance, but thalamic inputs to neurons in primary sensory cortices are $O(1)$. This results in a loose excitatory-inhibitory balance, yet net input remains subthreshold throughout the cortical dynamic range. We have recently developed the "stabilized supralinear network" (SSN) model, in which such a loose balance arises when external and network excitatory-inputs are $O(1)$. SSN responses exhibit a supralinear (sublinear) dependence on stimuli for very weak (moderate to strong) stimuli, and account for a range of nonlinear normalization phenomena observed in sensory cortices. Here, we show that empirical estimates of neural connection probabilities and strengths are consistent with this scenario, yet yield large enough estimated input fluctuations to self-consistently sustain subthreshold fluctuation-driven activity. Evidence suggests that cortical networks are inhibition-stabilized: effective recurrent excitation is sufficiently strong to render them unstable without recurrent inhibition. The SSN crossover from supra to sublinear behavior is accompanied by a transition to the inhibition-stabilized regime, which we estimate to occur as average responses exceed 0.5-7Hz.

I-94. Modelling retinal activity with Restricted Boltzmann Machines: a study on the inhibitory circuitry

Matteo Zanotto[1] MATTEO.ZANOTTO@IIT.IT
Stefano Di Marco[1] STEFANO.DIMARCO@IIT.IT
Alessandro Maccione[1] ALESSANDRO.MACCIONE@IIT.IT
Luca Berdondini[1] LUCA.BERDONDINI@IIT.IT
Diego Sona[1,2] DIEGO.SONA@IIT.IT
Vittorio Murino[1] VITTORIO.MURINO@IIT.IT

[1]Italian Institute of Technology
[2]Fondazione Bruno Kessler

Being the first stage of the visual system, the retina performs an extremely important task: encoding visual information and making it available to higher visual areas. In order to be effective, the encoding should be rich enough to capture the information of different stimuli, and stable enough to show small variations across repeated presentations of the same stimulus. This should arguably be achieved as a collective behaviour, capable of subduing the variability in the response of single Retinal Ganglion Cells (RGCs). In this work we study the statistical structure of the signal produced by a large population of RGCs, looking for prototypical activation modes of the retina subject to photo-stimulation. RGCs are modelled as log-Gaussian Cox Processes and a mean-covariance Restricted Boltzmann Machine (mc-RBM) is used to model the joint distribution of the firing rates of all the neurons in the recorded population. Due to its formulation, the mc-RBM allows to infer a set of activation modes of the retina defined by the configuration of the model's latent variables. These activation modes are obtained in a fully unsupervised way using no information about the input and thus reflect the regularities of RGCs signal. In our work we show that the activity modes found through the mc-RBM map reliably to different visual stimuli. Moreover, we show that the inferred modes can be used to evaluate the information content of the retinal signal. As a case study, we evaluate the information carried by the concurrent firing rates of RGCs of a retina in normal conditions and after pharmacologically blocking GABAC, at first, and then GABAC plus GABAA and GABAB receptors. As expected from physiology, blocking the inhibitory circuitry disrupts the spatiotemporal precision of retinal encoding, resulting in a reduced Mutual Information between the inferred activation modes and the presented visual stimuli.

I-95. Sequence generating recurrent networks: a novel view of cortex, thalamus, and the basal ganglia

Sean Escola SEAN.ESCOLA@GMAIL.COM
Larry Abbott LFA2103@COLUMBIA.EDU

Columbia University

Complex cognitions and behaviors can be viewed as sequences of simpler, more stereotyped components. For example, peeling a banana requires the combination of simple reaches and grasps. If we map each component of a sequence onto an activity "state" of a neural network, then sequence generation becomes the appropriately timed and ordered activation of network states. In this work, we present a novel multistate neural network rate model that can perform the computation of sequence generation and that can be mapped onto the connectivity of structures known to be involved in sequence generation: cortex, thalamus, and the basal ganglia. Mathematically, the state signal is provided to the network by perturbations to the synaptic weight matrix rather than by the standard approach of adding state-dependent inputs. The cortical-subcortical circuitry may implement this novel state signaling mechanism as follows. 1) The corticothalamocortical loop can by viewed as a perturbation to the direct intracortical connections. 2) The basal ganglia output nuclei (BGON) provide state-dependent tonic inhibitory patterns onto thalamus that select which thalamic cells participate in the corticothalamocortical loop for each state. Thus, when the inhibitory pattern changes, the perturbation to the cortical synaptic weight matrix changes. 3) The striatum listens to cortex and provides appropriately timed phasic signals to the BGON to switch their inhibitory patterns. In sum, when the striatum detects a transition opportunity, it drives the BGON to select a new set of thalamic cells to participate in the corticothalamocortical loop, which switches the cortical state. These novel hypotheses of subcortical functioning during sequencing yield experimental predictions. Additionally, this model has the computational advantage of near orthogonality between the network activity patterns of different states, which trivializes readout and results in a high capacity. This advantage provides an evolutionary argument for why sequence generation is computed via subcortical structures.

I-96. Symmetric matrix factorization as a multifunctional algorithmic theory of neural computation

Cengiz Pehlevan[1,2] CENGIZ.PEHLEVAN@GMAIL.COM
Dmitri B Chklovskii[3] DCHKLOVSKII@SIMONSFOUNDATION.ORG

[1]Janelia Farm Research Campus
[2]Howard Hughes Medical Institute

[3]Simons Center for Data Analysis

The cortex can perform different computational tasks using physiologically stereotypical hardware, neurons and synapses. Is it possible for the same neural algorithm to perform multiple tasks? If so, what determines which computational task is performed by the algorithm? To tackle these questions, we focus on two key unsupervised learning tasks that the cortex must perform: clustering and feature discovery. We show that these tasks can be unified algorithmically by symmetric matrix factorization (SMF) of the sample covariance matrix of the streamed data. We demonstrate that the SMF cost function can be minimized online by a biologically plausible single-layer network with local learning rules. Unconstrained SMF leads to a neural network that extracts the principal subspace of the streamed data. But when we introduce the biologically inspired nonnegativity constraint on the output the network becomes multi-functional: if the streamed data has clear cluster structure, the network performs soft clustering; if the streamed data is generated by a mixture of sparse features, e.g. natural images, the network discovers those sparse features. Interestingly, just like in neural circuits, nonnegative SMF can both reduce and expand the dimensionality of the input. The derived nonnegative SMF network replicates many aspects of cortical anatomy and physiology including unipolar nature of neuronal activity and synaptic weights, sparse heavy-tailed distribution of neuronal activity, local synaptic plasticity rules and the dependence of learning rate on cumulative neuronal activity. By proposing a biologically plausible algorithm performing two different tasks we make a step towards a unified algorithmic theory of neuronal computation.

I-97. Dendritic computations may contribute to contextual interactions in neocortex

Lei Jin[1] LEIJ@USC.EDU
Bardia F Behabadi[2] BARDIAFB@GMAIL.COM
Chaithanya A Ramachandra[3] CHAITHANYA@GMAIL.COM
Bartlett W. Mel[1] MEL@USC.EDU

[1]University of Southern California
[2]Qualcomm Inc.
[3]Eyenuk Inc.

A striking feature of sensory cortex is that the main feedforward projection rising "vertically" from the input layer (L4) to the next stage of processing in layer 2/3, which defines a layer 2/3 pyramidal neuron's (PN's) "classical" receptive field, accounts for only a small fraction of the excitatory innervations ($<$10%). The largest source of contacts (60-70%) arises from the massive network of horizontal connections through which cortical PNs exchange "contextual" information. Despite their large numbers, relatively little is known regarding the behavioral relevance of horizontal connections, the functional form(s) of the classical-contextual interactions they give rise to, or the biophysical mechanisms underlie their modulatory effects. To explore the role of nonlinear synaptic integration in mediating classical-contextual interactions, we have focused on the problem of contour detection in natural images. We compared two functions: (1) the probabilistic interaction between aligned edge elements in the classical and extra-classical receptive fields of a virtual V1 simple cell, computed from statistics of human-labeled natural contours, and (2) the NMDA-dependent interactions between proximal and distal excitatory inputs to PN basal dendrites, computed using compartmental simulations. We found these two functions, of very different origin, matched closely, suggesting that nonlinear multiplicative interactions between inputs to PN dendrites could contribute directly to the integration of classical and contextual information in the cortex. Given the variety of classical-contextual interactions that are likely to be needed in the cortex, we extended our biophysical studies to map out the spectrum of nonlinear interactions that PNs are capable of producing in their dendrites. Examples are shown to illustrate the variety of 2-pathway interactions that can be achieved. Our results support the prediction that basal and apical oblique dendrites of PNs play a central role in mediating classical-contextual interactions in the neocortex, owing to their flexible ecologically-relevant analog computing capabilities.

I-98. Stability of trained recurrent neural networks

Alexander Rivkind ARIVKIND@TX.TECHNION.AC.IL
Omri Barak OMRIB@TX.TECHNION.AC.IL

Technion - Israel Institute of Technology

Reservoir computing (RC) [1] is a popular paradigm for training recurrent neural networks (RNN). A network of neurons with random recurrent connectivity (referred to as the reservoir) has a linear readout with weights trained to produce a desired output, while keeping all other connectivity weights fixed. Output is fed back into the network, effectively changing its dynamics. Despite many successful applications of this paradigm in recent years, little is known about the conditions for stable readout training. Here, we advance the research of these conditions by examining training for fixed points and for periodic patterns. In our analysis we rely on relations between closed loop (output is fed back) and open loop (target is fed back) settings. We report the following results: i) Reservoir chaos must be suppressed in open loop for successful closed loop FORCE [2] learning; ii) learnability of fixed point solutions is shown to follow analytically from Mean Field Theory of input driven networks [3]. As a corollary, we show that a very weak modification of reservoir weights, implemented by a semi-Hebbian plasticity rule, suffices to transition reservoir dynamics from chaos to a stable fixed point; iii) under reasonable conjectures, the analysis used for (ii) is extendable for periodic patterns; iv) Contrary to common belief [1,2], readout weights obtained from least square open loop learning are suitable for closed loop pattern generation given a sufficiently large network. Our results offer insights on existing RC learning rules, hold hope to devise new rules, and set the stage for a theory of reservoir computing learnability. Reference: [1] Jaeger, H., Tech Report 2001; [2] Sussillo, D. and Abbott, L.F., Neuron 2009; [3] Rajan, K. et al., PRE 2010.

I-99. Supervised learning in spiking recurrent networks: balance is the key

Ralph Bourdoukan RALPH.BOURDOUKAN@GMAIL.COM
Sophie Deneve SOPHIE.DENEVE@ENS.FR

Ecole Normale Superieure

The brain has to constantly perform predictions of sensory inputs or motor trajectories, and thus, to learn temporal dynamics based on error feedback. However, it remains unclear how such supervised learning is implemented in biological neural networks. Learning in recurrent spiking networks is notoriously difficult because local changes in connectivity may have an unpredictable effect on the global dynamics. The most commonly used learning rules, such as temporal back-propagation, are not local and thus not biologically implausible. On top of that, reproducing the Poisson-like statistics of neural responses require the use of networks with balanced excitation and inhibition. Such balance is easily destroyed during learning. Using a top-down approach, we show how networks of integrate-and-fire neurons can learn arbitrary linear dynamical systems by feeding back their error as a feed-forward input. The network uses two type of recurrent connections: fast and slow. The fast connections learn to balance excitation and inhibition using a voltage-based plasticity rule. This enforces a maximally efficient spike-based representation of the underlying dynamical state while insuring that spike trains remain irregular and sparse. The slow connections are trained to minimize the error feedback using a current-based Hebbian learning rule. Importantly, the balance maintained by fast connections is crucial to ensure that global error signals are available locally in each neuron, in turn resulting in a local learning rule for the slow connections. This demonstrates that spiking networks can learn complex dynamics using purely local learning rules, using E/I balance as the key rather than an additional constraint. The resulting network implements a predictive coding scheme and can be considered as minimal, in terms of connectivity and activity, in order to implement a given function. This contrasts with approaches based on liquid computing, where very large, dynamically complex networks have to be used.

I-100. How entropy-producing networks can have precise spike times

Maximilian Puelma Touzel[1,2] MPTOUZEL@NLD.DS.MPG.DE
Michael Monteforte[1] MONTE@NLD.DS.MPG.DE
Fred Wolf[1] FRED@NLD.DS.MPG.DE

[1]MPI for Dynamics and Self-Organization
[2]BCCN Goettingen

Dynamic instabilities have been proposed as an explanation for the decorrelation of stimulus-driven activity observed in sensory areas such as the olfactory bulb [1]. However, these instabilities must be tamed somehow to be useful. Chaotic rate networks can be trained to stabilize the neighborhood around one of their trajectories [2]. How such a phenomenon could emerge in spiking networks is a topic of current research. Inhibitory circuit models of leaky integrate-and-fire neurons exhibit the irregular and asynchronous activity associated with chaos despite their stability, a phenomenon called 'stable chaos'. A recent finite-sized perturbation analysis nevertheless revealed an exponential instability accessed when perturbations are above a critical strength that scales characteristically with the size, density and activity of the circuit [3]. This demonstrates the existence of a large, finite set of locally-attracting irregular, asynchronous spike sequences, each contained in a 'flux tube', the latter set of which are mutually-repelling. Where stable chaos comes from and how far it extends away from simplified models remains unclear. We detail the effects of finite-sized perturbations on spiking trajectories and clearly reveal the mechanism underlying stable chaos in these circuits. We then track a tube in time and map its boundary dynamics, which we find are determined by the sensitivity horizon of the network: the characteristic end of influence of a perturbation in the causally connected future of the local phase space. From this, we analytically derive the probability that the network dynamics will erase an initial perturbation of a given size. The result gives the source of the scaling relations and leads us to quantitatively estimate the extent of stable chaos as more realistic single neuron features are introduced, highlighting those important for this phase space structure, which serves as an attractor reservoir that downstream networks might use to decode sensory input.

I-101. Learning universal computations with spikes

Raoul-Martin Memmesheimer R.MEMMESHEIMER@SCIENCE.RU.NL
Dominik Thalmeier D.THALMEIER@DONDERS.RU.NL
Marvin Uhlmann MARVIN.UHLMANN@STUDENT.RU.NL
Hilbert J Kappen B.KAPPEN@SCIENCE.RU.NL

Radboud University Nijmegen

Providing the neurobiological basis of information processing in higher animals, spiking neural networks must be able to learn a variety of complicated computations, including the generation of appropriate, possibly delayed reactions to inputs and the self-sustained generation of complex activity patterns, e.g. for locomotion. Many such computations require previous building of intrinsic world models. Here we show for the first time how spiking neural networks may solve these different tasks. Firstly, we derive constraints under which classes of spiking neural networks lend themselves to substrates of powerful general purpose computing. We then combine such networks with learning rules for outputs or recurrent connections. We show that this allows to learn even difficult benchmark tasks such as the self-sustained generation of desired low-dimensional chaotic dynamics or memory-dependent computations. Furthermore, we show how spiking networks can build models of external world systems and use the acquired knowledge to control them.

I-102. Temporal evolution of information in neural networks with feedback

Aram Giahi Saravani[1,2] ARAM@RICE.EDU
Xaq Pitkow[1,2] XAQ@RICE.EDU

[1]Rice University
[2]Baylor College of Medicine

Feedback constitutes an important attribute of information processing in the brain. Here we use a simple model system to identify how neural feedback transforms the internal representation of dynamic sensory variables. For analytical tractability, we investigate a linear dynamical system with additive gaussian noise, and relate the synaptic weight matrix to an effective Kalman filter that is embedded within it. The architecture of a recurrent network influences the information content of the network, as well as the dynamics of two experimental measurements often used to describe the computations of neural networks: the psychophysical kernel and choice probability. For this model, we compute the Fisher information of a general readout and compare its efficiency to that of an optimal leaky integrator. The optimal structure of the network is determined by both the persistence time of the stimulus and the Fisher information of sensory measurements. The psychophysical kernel correlates brief fluctuations of a dynamic stimulus to subsequent actions or perceptual estimate; the choice probability is the same measure applied to neural responses. Despite the similarity of these measures, we show they have quite different behaviors. The psychophysical kernel may be dominated by early or late stimuli, depending on the recurrence strength. Yet because the neural network accumulates signals, the choice correlations always increase, albeit with different shapes. This model thus provides clear guidance for what to expect from common experimental measurements of a recurrently connected network.

I-103. Assembly training and reinforcement through spike timing

Gabriel Ocker[1] GKO1@PITT.EDU
Changan Liu[2] PODLCA@MATH.UH.EDU
Kresimir Josic[2] JOSIC@MATH.UH.EDU
Brent Doiron[1] BDOIRON@PITT.EDU

[1]University of Pittsburgh
[2]University of Houston

Spike timing-dependent plasticity (STDP) provides a mechanism for the Hebbian theory of synaptic weight changes due to the precise timing of pre- and post-synaptic action potentials. STDP is commonly believed to underlie Hebbian plasticity in neural assemblies, linking network architecture to functional dynamics. Theories explaining the emergence of such macroscopic structure have, however, so far relied on plasticity rules that depend on the neurons' firing rates rather than precise spike timing, or have relied on external inputs to generate spike-time correlations. We have developed a theory for spike- timing dependent plasticity in recurrent networks of integrate-and-fire neurons. This theory separates the contributions of spike-time correlations from external inputs and different intra-network sources to synaptic plasticity. We show that internally generated spike-time correlations can reinforce a diverse array of macroscopic network structures through STDP. These include homogenous, clustered and feedforward architectures. Spatially correlated inputs, by inducing additional spike-time correlations in pairs of neurons, can guide the network structure into any of these macroscopic organizations, depending on the spatial profile of the inputs. In contrast to other studies, our network exhibits homogenous firing rates so that only precise spike-time correlations control plasticity. This reveals the potential richness of macro- and microscopic network structure available from STDP and internally generated spike-time correlations.

I-104. A ‘Neural Turing Machine’ model of working memory.

Greg Wayne GREGWAYNE@GOOGLE.COM
Alex Graves GRAVESA@GOOGLE.COM
Ivo Danihelka DANIHELKA@GOOGLE.COM
Google DeepMind

Working memory refers to the capacity for the rule-based manipulation of short-term memory stores [1] and is believed to be subserved by a loop between the prefrontal cortex and basal ganglia [2]. It is strongly implicated in most cognitive tasks; for example, in mental addition, one must store two numerical variables, perform the addition operation on them, and then store the result. Thus it is one of the closest processes in human cognition to conventional computation. We draw on this analogy by constructing a controller circuit that uses an attentional mechanism to modify a large memory system selectively; the combined system is analogous to a central processing unit that interacts with random access memory, so we call it a ‘Neural Turing Machine’ (NTM). The attentional mechanism can focus on information in memory based on its content (feature-based attention) or focus on information based on where it is stored (spatial attention). Previous models of working memory can be divided between those that hypothesize biophysical processes to sustain information [3-5] and those that consider purpose-built networks to solve specific tasks [6-8]. Our model does not directly target a biophysical level of explanation but is computationally very general. Among other tasks, we have trained the NTM, from examples, to reason about structured knowledge domains (e.g., to make inferences about relationships in family trees), to sort data, and to perform long addition. (We continue to explore new tasks.) The NTM often finds, to our eyes, very surprising dynamical rules to generate answers to novel queries. As a bonus, it also diversely outperforms the state-of-the-art Long Short-Term Memory [9] recurrent network in widespread use in machine learning. We are confident that the Neural Turing Machine will stimulate discussions in computational neuroscience about how brains do indeed perform computation and represent information.

I-105. Representation of syntactic information across human cerebral cortex

Lydia Majure LYDIA.MAJURE@GMAIL.COM
Alex Huth ALEX.HUTH@GMAIL.COM
Jack Gallant GALLANT@BERKELEY.EDU
University of California, Berkeley

Syntax is a critical part of human language, but little is known about precisely how syntactic information is represented in the human brain. Here we address this issue by means of a quantitative method for mapping the representation of information across the cerebral cortex. We used functional MRI to record brain activity in 6 normal humans while they listened to natural narrative stories. We then extracted syntactic features from the stories, and we used these as regressors to model brain activity in ~100,000 distinct volumetric pixels (i.e., voxels) distributed across the cerebral cortex. The syntactic features were obtained from a Hierarchical Hidden Markov Model (HHMM) trained on a large corpus of informal English text, and the resulting features reflected both word classes and phrasal structure. Regularized linear regression was used to determine how the presence of each syntactic feature in the stories influenced brain activity in each voxel, in each individual brain. We find that syntactic features are represented bilaterally in several regions of cortex, including the precuneus, dorsolateral prefrontal cortex and inferior temporal sulcus. We also compared the syntax results to those obtained using semantic features (derived from word co-occurrence statistics). We find that both syntactic and semantic information are represented in the inferior frontal gyrus, superior temporal sulcus, supramarginal gyrus, and interparietal sulcus. These results indicate that the representation of syntax while listening to natural speech is both more broad and less lateralized than suggested by previous studies.

II-1. An analysis of how spatiotemporal models of brain activity could dramatically improve MEG/EEG source

Camilo Lamus[1,2] LAMUS@MIT.EDU
Matti Hamalainen[2] MSH@NMR.MGH.HARVARD.EDU
Emery Brown[3] ENB@NEUROSTAT.MIT.EDU
Patrick Purdon[3] PATRICKP@NMR.MGH.HARVARD.EDU

[1]Massachusetts Institute of Technology
[2]Massachusetts General Hospital
[3]Harvard University

MEG and EEG are neuroimaging techniques that provide recordings of brain activity with high temporal resolution. The ability to estimate cortical activity from MEG/EEG data is limited by two factors. The first is that the number of source parameters that needs to be estimated at a given time instant is much larger than the number of sensors. The second, due to the biophysics of MEG/EEG, is that source currents from some regions of the brain, depending on their location and orientation, are more easily observed at the scalp than others. These factors limit the number of independent source parameters that can be recovered from an individual temporal measurement, and effectively restrict parameter estimates to cortical areas whose activity is easily observed at the sensors. In this paper we show that the incorporation brain connectivity spatial information and source activity temporal dynamics could dramatically improve MEG/EEG inverse solutions. In a statiotemporal dynamic framework, information about cortical activity at a given time instant is mapped to not only the immediate measurement, but also to measurements both forward and backwards in time. To analyze how such mapping occurs, we develop the concept of the dynamic lead field, which express how source activity at given time instant is mapped to the complete temporal series of observations. With this mapping, we show that the number of independent source parameters that can be recovered from MEG/EEG data can be increased by up to a factor of ~20, and that this increase is primarily represented by cortical areas that are more difficult to observe. Our result implies that future developments in MEG/EEG analysis that explicitly model connections between brain areas and temporal structure have the potential to dramatically increase spatiotemporal resolution by taking full advantage of the brain's inherent connectivity and dynamics.

II-2. Distinct dynamics of ramping activity in the frontal cortex and caudate nucleus in monkeys

Long Ding LDING@MAIL.MED.UPENN.EDU

University of Pennsylvania

The prefronto-striatal network is an important player for many cognitive functions, including perceptual decision making and reward-modulated behaviors. Consistent with strong feed-forward prefronto-striatal projections, neuronal responses on cognitive tasks, as assessed by mean firing rate averaged over many trials, frequently show similar patterns in the prefrontal cortex and striatum. Such similarities make it a challenging problem to tease apart distinct functional contributions of the two regions. Previously we showed that subsets of neurons in the frontal eye field (FEF) of the prefrontal cortex and caudate region of the striatum exhibit similarly modulated ramping activity on two tasks: in anticipation of reward-predictive targets on an asymmetric-reward saccade task and during stimulus viewing on a visual motion direction-discrimination task. Here I show that, despite these similar mean firing rate patterns, FEF and caudate responses differ in other temporal dynamics for both perceptual and reward-based tasks. Compared to simulation results, the temporal dynamics of FEF activity are consistent with accumulation of sensory evidence used to solve a perceptual task but not with accumulation of reward context information used for the development of a reward bias. In contrast, the temporal dynamics of caudate activity is consistent with accumulation of reward context information but less so with accumulation of sensory evidence. These results suggest that FEF and caudate neurons may have specialized functions for different tasks even with similar average activity.

II-3. Strategies for exploration in the domain of losses

Paul Krueger PAULK@PRINCETON.EDU
Robert Wilson RCW2@PRINCETON.EDU
Jonathan Cohen JDC@PRINCETON.EDU

Princeton University

The explore-exploit dilemma arises whenever we must choose between exploiting a known option or exploring the unknown in hopes of finding something better. Most previous work has focused on explore-exploit scenarios in the domain of gains where the goal is to maximize reward. In many real-world decisions, however, the primary objective is to minimize losses. While optimal decision strategies should be the same for both gains and losses, it is well known that human decision making can shift significantly when confronted with losses. In this study, we compared explore-exploit behavior of human participants under conditions of gain and loss, using a model that measures two distinct exploration strategies: a directed strategy, in which exploration is driven by information seeking, and a random strategy, in which exploration is driven by decision noise. We found that people use both types of exploration regardless of whether they are exploring in response to gains or losses and that there is quantitative agreement between the exploration parameters across domains. Our model also revealed an overall bias toward the lesser known option in the domain of losses, independent of directed exploration, indicative of uncertainty seeking. Based on these results, we built a second, reduced model that explicitly distinguished information seeking from uncertainty seeking that was better able to explain the data. Taken together, our results show that explore-exploit decisions in the domain of losses are driven by three independent processes: a baseline bias toward the uncertain option, and directed and random exploration.

II-4. Parietal Reach Region (PRR) inactivation affects decision and not attention process in reach choices

Vasileios Christopoulos[1] VCHRISTO@CALTECH.EDU
James Bonaiuto[2] JBONAIUTO@GMAIL.COM
Igor Kagan[3] IKAGAN@DPZ.EU
Richard Andersen[1] ANDERSEN@VIS.CALTECH.EDU

[1]California Institute of Technology
[2]Sobell Department of Motor Neuroscience
[3]German Primate Ctr. Göttingen, Germany

A growing body of neurophysiological studies challenges a long-held theory that views decision-making as a distinct cognitive process from the neural systems for perception and action. Recent findings suggest that many action decisions emerge via a continuous competition between neuronal populations within the same areas that plan and guide action execution (Cisek, 2012). The main line of evidence is the existence of decision-related neural activity in brain regions that traditionally have been associated with sensorimotor control (Platt and Glimcher 1999). However, it has been argued that this activity is not 'genuinely motor', but instead it is related to spatial attention or visual salience (Padoa-Schioppa, 2011). One approach to establish whether brain regions are causally involved in decision process is to temporarily inactivate these areas and observe the effects on decision making. We studied whether PPR inactivation causes deficits on attention and/or decision process (Wilke et al., 2012). We reversibly inactivated PRR by locally injecting the GABA-A agonist muscimol, while two macaque monkeys performed memory-guided reach or saccade movements either to a single target (instructed trials) or selected between two targets presented simultaneously in both hemifields (free-choice trials). We found that PRR inactivation led to a strong reduction of contralesional choices, but only for reaching. In contrast, the inactivation did not affect the saccade choices. We also found no effects on the reaching or saccade movements to single targets presented in either hemifield. These results cannot be explained as a spatial attention deficit, since the 'lesion' had an impact only on reaching and not on saccade choices. Hence, PRR seems to be causally involved in reach decisions. Finally, we developed a biologically plausible computational framework that explains how PRR

inactivation affects only reaching choices, leaving animals' response towards single targets largely intact.

II-5. Modeling motor commands as traveling waves of cortical oscillation

Stewart Heitmann[1] HEITMANN@PITT.EDU
Tjeerd Boonstra[2] T.BOONSTRA@UNSW.ED.AU
Pulin Gong[3] P.GONG@PHYSICS.USYD.EDU.AU
Michael Breakspear[4] MICHAEL.BREAKSPEAR@QIMRBERGHOFER.EDU.AU
Bard Ermentrout[1] BARD@PITT.EDU

[1]University of Pittsburgh
[2]University of New South Wales
[3]The University of Sydney
[4]Queensland Institute of Medical Research

Traveling waves of beta-band (15–30 Hz) oscillations have been observed in motor cortex during the preparation and performance of reaching movements. However the functional relationship between those oscillations and muscle activity is unknown. We speculate that motor commands can be encoded in the spatial properties of traveling waves which are then spatially filtered by the descending motor tract to selectively drive muscle activity. We propose that wave formation is governed by the lateral inhibitory connections in cortex. Spatial filtering of the waves is performed by the dendritic receptors of the primary output neurons of motor cortex, namely, layer 5 pyramidal tract neurons. We demonstrate this proposal using a simplified model of the descending motor system in which the orientation of waves in cortex govern muscle drive in a simulated biomechanical joint. Cortex was modeled by a sheet of coupled oscillators with anisotropic inhibitory surround connectivity. It produced planar waves of synchronized beta-band oscillations that were oriented along the dominant direction of the lateral inhibitory connections. The dendritic receptor fields of the pyramidal tract neurons were modeled as Gabor filters. Two populations of pyramidal tract neurons were constructed, one population for each muscle in the biomechanical joint. The two populations were tuned to orthogonal wave orientations with broadly overlapping tuning curves. This arrangement allowed graded co-activation of the opponent muscles. Arbitrary joint postures could thus be obtained by appropriate manipulation of the lateral connectivity in cortex. The proposed model demonstrates a putative mechanism by which oscillatory patterns in cortex are translated into steady limb postures. Furthermore, the model reproduces some physiological aspects of motor control. Namely the reduction of cortical beta power at the onset of movement and the weak levels of cortico-muscular coherence observed during steady motor output.

II-6. Expectation modulates activity in primary visual cortex

Ilias Rentzeperis[1] ILIAS@BRAIN.RIKEN.JP
Justin Gardner[2] JLG@STANFORD.EDU

[1]RIKEN Brain Science Institute
[2]Stanford Universtiy

Humans can exploit the prior probability of whether an event will occur to improve their interpretation of the world. For example, in signal detection theory, the criterion can be adjusted according to prior probability to maximize performance. We asked what are the cortical mechanisms that allow humans to adjust their criterion to accommodate different prior probabilities of stimulus occurrence. Key to understanding what occurs is to control for effects of attention which may covary with expectation. We therefore designed a signal detection task in which we could independently modify attention and prior expectation and measured cortical responses in early visual cortex with functional MRI. We found primary visual cortex responses that were modulated with expectation according to predictive coding theories (Summerfield et. al., 2008, Nature Neuroscience; Kok et. al., 2012, Neuron) in which an expected stimulus evoked a smaller neural response than stimuli that were unexpected.

II-7. Changes in prefrontal and striatal network dynamics during cocaine administration

Wesley Smith WES.CSMITH@GMAIL.COM
Leslie Claar LDCLAAR@GMAIL.COM
Justin Shobe JSHOBE@GMAIL.COM
Konstantin Bakhurin BAKHURIN@UCLA.EDU
Natalie Chong NATALIECCHONG@YAHOO.COM
Sotiris Masmanidis SMASMANIDIS@UCLA.EDU

University of California, Los Angeles

It is increasingly understood that drugs of abuse alter brain circuit function at the network scale. This places severe limitations on the traditionally reductionist approach to studying the neurobiological basis of addiction in single isolated brain regions. However, it remains unclear how network dynamics are altered at the resolution of single cells and spikes. We hypothesize that drug administration enhances fronto-striatal functional connectivity, and that the strength of this enhancement correlates with behavioral preference to drug-associated cues. To address this hypothesis we utilize novel 512 channel silicon microprobes to simultaneously record in two interconnected areas, the medial prefrontal cortex (mPFC) and nucleus accumbens (NAc), that are strongly implicated in mediating the addictive properties of cocaine. To study network activity and behavioral changes that accompany cocaine administration, we have established a conditioned odor preference test for head-restrained mice. When animals are tested with brief presentations of three odors after cocaine conditioning, they exhibit a conditioned response by running on a spherical treadmill preferentially after the cocaine-paired odor, but not a saline-paired or unpaired cue. The microprobes enable recordings of around 100 well isolated units per area, providing unique opportunities for studying how firing within and between the mPFC and NAc are correlated under different stages of drug craving. Here we will present preliminary results revealing changes in neural correlation coefficients, the strength of drug-paired cue encoding, and heretofore unseen aspects of cortical hypoactivity. We conclude by discussing the broader potential for our multi-region electrophysiological recording approach to shed light on reward circuit function during normal and drug-based learning.

II-8. Sparse encoding of complex action sequences

Andreas Thomik ANDREAS.THOMIK10@IMPERIAL.AC.UK
Aldo Faisal A.FAISAL@IMPERIAL.AC.UK

Imperial College London

A fundamental problem in neuroscience is to understand how the brain translates a symbolic sequence of action descriptors, or high-level motor intention, into the appropriate muscle commands. We use a data-driven approach to seek a generative model of movement capturing the underlying simplicity of spatial and temporal structure of behaviour observed in daily life. We take the view that the brain achieves this feat by mapping the necessary computation onto a finite and low-dimensional subset of control building blocks of movement, characterised by high correlation between a subset of the joints involved — kinematic primitives. These would be combined as required to achieve a given task. We investigate this possibility by collecting a large data set natural behavior capturing 90% of the perception-action loop using lightweight, portable and unobtrusive motion capture systems over a prolonged period of time. From this data we learn in an unsupervised fashion a dictionary of kinematic primitives (which we term eigenmotions) by analysing the local temporal correlation structure of the data. We show that the dictionaries learnt are broadly consistent across subjects with minor variations accounting for individuality of the subject and variations in the tasks executed. Using this dictionary we can compute a sparse representation of the data which is characterised by a very low-dimensional latent structure. Using this latent representation we can translate the time-series of joint movements into a symbolic sequence (“behavioural barcode”), which captures both spatial and temporal structure of the behavior. Sequences of different eigenmotions thus represent a “language of movement” which we can analyse to find its grammatical structure, yielding an insight into how the brain may generate natural behavior by temporally sparse activation of “eigenmotion neurons”, similar to grasp-type specific neurons found in the monkey premotor cortex.

II-9. Choice probabilities, detect probabilities, and read-out with multiple neuronal input populations

Ralf Haefner RALF.HAEFNER@GMAIL.COM
University of Rochester

Understanding the contribution of sensory neurons to perceptual decisions has been a long-standing goal of systems neuroscience (Parker & Newsome 1998). Recently we have presented the analytical relationship between choice probabilities (CP), noise correlations (NC) and read-out weights in the classical feedforward decision-making framework making it possible to rigorously interpret empirical CP results, and to infer the read-out weights in experiments that measure CPs and NCs at the same time (Haefner et al. 2013). For the derivation we assumed that behavioral reports are distributed evenly between the two possible choices. This assumption is often violated in empirical data — especially when computing so-called ‘grand CPs’ combining data across stimulus conditions. Here, we extend our analytical results to situations when subjects show clear biases towards one choice over the other, e.g. in non-zero signal conditions. Importantly, this also extends our results from discrimination tasks to detection tasks and detect probabilities (DP) for which much empirical data is available. We find that CPs and DPs depend monotonously on the fraction, p, of choices assigned to the more likely option: CPs and DPs are smallest for p=0.5 and increase as p increases, i.e. as the data deviates from the ideal, zero-signal, unbiased scenario. While this deviation is small, our results suggest a) an empirical test for the feedforward framework and b) a way in which to correct CP and DP measurements before combining different stimulus conditions to increase signal/noise. Furthermore, we apply this framework to a neuronal population that receives inputs from two separate input populations (e.g. MT from V1 and V2, Smolyanskaya et al. 2014) and show how CPs and DPs depend on both input correlations and on the firing-rate-dependent transformation of input correlations to response correlations (de la Rocha 2007).

II-10. Posterior parietal and prefrontal cortex involvement in rat auditory parametric working memory

Athena Akrami[1,2] AAKRAMI@PRINCETON.EDU
Ahmed El Hady[1] AHADY@PRINCETON.EDU
Carlos Brody[1,2] BRODY@PRINCETON.EDU

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

Parametric Working Memory (PWM) is the short-term storage of graded stimuli to guide behavior, and its neural correlates have been studied in primates1. The prefrontal and posterior parietal cortices (mPFC and PPC) have been proposed to be involved in working memory2-3, but no inactivation experiments probing whether these areas are necessary for parametric working memory have been performed. Here, we develop a parametric auditory delayed comparison task in rats, and for the first time show that activity in PPC and mPFC is necessary for a PWM behavior. Using a logistic regression analysis to analyze the specific task components impacted by inactivations, our data suggest that PPC may be required for the parametric memory itself, while mPFC may be required for more general aspects of the task.

II-11. Learning hierarchical structure in natural images with multiple layers of Lp reconstruction neurons

Zhuo Wang JSWANGZHUO@GMAIL.COM
Alan A Stocker ASTOCKER@SAS.UPENN.EDU
Daniel Lee DDLEE@SEAS.UPENN.EDU

University of Pennsylvania

High dimensional perceptual stimuli such as images or sounds are thought to be more efficiently represented in neural populations by redundancy reduction (Attneave 1954; Barlow 1961). Computational models for efficient coding optimize information theoretic objective functions such as maximum mutual information (MMI). In particular, MMI has been shown to be a promising principle to understand V1 simple cells due to its success in predicting edge-like filters for natural images (Bell and Sejnowski 1997). However, it is more difficult to apply the MMI principle iteratively to train additional layers without substantial modifications (Karklin and Lewicki 2003; Shan, Zhang and Cottrell 2006). Our work investigates the general principle of minimizing Lp reconstruction error to model multiple layers of noisy linear-nonlinear neurons. We show that both MMI (L0) and minimum mean squared error (MMSE, L2) are special cases of this generalized principle and optimal analytic solutions can be derived if the stimuli follows an elliptical distribution. In particular, we find that the optimal representation does not immediately eliminate correlations, but gradually reduces redundancy across the layers. As an application, we consider small (8x8) patches of natural images (van Hateren 1998). We demonstrate that the distribution of pixel intensities in these patches is near elliptical, and iteratively train multiple layers of MMSE neurons. We show detailed results for a two-layer model, where the response properties of the first and second layer neurons qualitatively match features of simple and complex cells respectively.

II-12. Cellular resolution functional imaging of cortical dynamics during integration of evidence in the rat

Benjamin Scott[1] BBSCOTT@GMAIL.COM
Christine Constantinople[1] CMC9@PRINCETON.EDU
Jeffrey Erlich[2] JERLICH@NYU.EDU
Carlos Brody[1,3] BRODY@PRINCETON.EDU
David Tank[1] DWTANK@PRINCETON.EDU

[1]Princeton University
[2]NYU Shanghai
[3]Howard Hughes Medical Institute

Investigation of the neural mechanisms of cognition will be facilitated by the use of animal models where techniques for neural circuit manipulations and precise behavioral control can be performed. In particular, rats are an increasingly attractive model to study cognition due to the availability of advanced brain recording techniques, such as in vivo optical imaging, and automated procedures for behavioral training. We recently developed a system that brings these two technologies together: an integrated two-photon microscope and automated behavioral training apparatus that allows cellular resolution functional imaging during voluntary head restraint. In proof of principle experiments rats were trained to perform a numerical comparison task involving gradual accumulation of evidence while head-fixed in a high throughput technician operated facility. Cortical dynamics were recorded in expert rats during performance of the task by two-photon imaging of calcium transients in GCaMP6f-labeled neurons. In ongoing experiments we are using linear regression analysis to characterize and compare the responses of neurons across three brain regions, the medial posterior secondary visual area (PM), the posterior parietal cortex (PPC) and the frontal orienting field (FOF) of the medial agranular cortex, to evaluate models of cortical processing during decision-making. These experiments suggest that automated systems for behavioral training combined with in vivo imaging during voluntary head restraint provide a useful approach to investigate the neural mechanisms of cognition.

II-13. The dynamics of memory search: a modelling study

Yifan Gu[1] YIGU8115@UNI.SYDNEY.EDU.AU
Pulin Gong[2] GONG@PHYSICS.USYD.EDU.AU

[1]University of Sydney
[2]School of Physics

The human mind has a remarkable ability to retrieve task relevant items from memory. It has been shown that there are clear regularities in such memory retrieval process. For instance, in the semantic fluency task, a paradigm used to study memory search, relevant items can be retrieved consecutively, with fast transitions between items in the same semantic group but slow transitions when they belong to different semantic groups; namely, the transition process depends on the semantic proximity of these items. The ratio of these transitions, however, is much less for patients with memory deficits such as Parkinson's disease than for healthy subjects. Despite the importance of this kind of transition process in memory retrieval, its dynamic nature and underlying neurophysiological mechanisms remain unknown. We develop a hierarchical spiking neural network capturing the similar structure as found in semantic networks. We show that for the hierarchical network with balanced excitation and inhibition (E/I balance), free retrieval of memorized items can happen when the system consecutively switches between different clusters in the network. We further demonstrate that the conditioned response probability of two clusters increases but its latency decreases as the semantic relatedness between them increases. We find that the transition process is a Levy flight process, with the distribution of transition distances following a power law function. However, when the E/I balance is broken in the network, there are serious deficits in the transition ability; at the behavioural level, this is consistent with the similar deficits found in patients with Parkinson's disease, and at the neurophysiological level, this is consistent with the observation that elevated excitation breaking the E/I balance can result in psychiatric diseases. Our work, therefore, suggests that the E/I balance is essential for one of the important computation tasks, i.e., memory retrieval in the brain.

II-14. Investigating the role of the striatum in working memory

Hessameddin Akhlaghpour HAKHLAGH@PRINCETON.EDU
Joost Wiskerke WISKERKE@PRINCETON.EDU
Jung Yoon Choi JUNGCHOI@PRINCETON.EDU
Jennifer Au JAU@PRINCETON.EDU
Ilana B Witten IWITTEN@PRINCETON.EDU

Princeton University

We aim to clarify the role of the striatum in working memory, using a delayed non-match to sample lever pressing task in rats. Optogenetic inactivations reveal that delay period activity in the striatum has a causal role in our task. Inactivations of the dorso-medial striatum during the longer delay periods led to significant impairments in accuracy. Inactivation of dorso-lateral striatum, however, did not lead to any significant impairment. Extracellular recordings from a total 115 units in the dorso-medial striatum of 9 rats, revealed that neurons exhibited sequential transient spiking activity tiling the duration of the trial. Peak firing rates were biased towards the beginning of the delay period. Roughly 60% of the neurons encoded the sample stimulus at some point of the task, and most of those neurons had most information about the stimulus within the first three seconds of the delay period. Our findings suggest that the dorso-medial striatum may have a more crucial role at the start of the delay period, which is consistent with theories that suggest that the striatum is responsible for gating information to working memory networks.

II-15. DMS inactivation disrupts model-based reinforcement learning in a two-step decision task

Thomas Akam[1] THOMAS.AKAM@NEURO.FCHAMPALIMAUD.ORG
Michael Pereira[2] MICHAEL.PEREIRA@NEURO.FCHAMPALIMAUD.ORG
Peter Dayan[3] P.DAYAN@UCL.AC.UK
Rui Costa[2] RUI.COSTA@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Neuroscience Programme
[2]Champalimaud Centre for the Unknown
[3]Gatsby Computational Neuroscience Unit, UCL

Mammalian decision making is thought to utilise at least two distinct computational strategies, model-free and model-based reinforcement learning (RL), instantiated in partially separate anatomical circuits. The former learns values of states and actions directly through reward prediction errors, while the latter utilises a forward model which predicts future state given current state and chosen action and supports computations akin to search in a decision tree. A key challenge in studying the neurobiology of these systems is developing tasks in which their contribution to behaviour can be differentiated. The classical approach of outcome devaluation has severe limitations for neuro-scientific applications as the strategies are dissociable only during a brief extinction test. Recently, a novel task has been developed for humans (Daw et al. 2011) in which an initial choice between two actions leads probabilistically to one of two states, from which further actions lead probabilistically to reward. Model-based and model-free RL can be differentiated either using model fitting or, as the two strategies have distinct action values updates, by analysing how events on one trial influence the subsequent choice. We have modified this task into a poke based format which mice learn in under 3 weeks with minimal shaping. Subjects perform hundreds of trials per session in a dynamic environment in which both action-state transition probabilities and the rewards available in different states change over time. Behaviour on the task is consistent with a mixture of model-based and model-free RL. Motivated by prior work implicating dorsomedial (DMS) and dorsolateral (DLS) striatum in model-based and model-free RL respectively, we are currently assessing how pharmacogenetic inactivation of these regions affects task performance. DMS inactivation changed behaviour in a manner consistent with a shift towards model-free control. DLS inactivation experiments are ongoing.

II-16. Influence of cognitive control on semantic representation: seeing things under a different light

Waitsang Keung WKEUNG@PRINCETON.EDU
Daniel Osherson OSHERSON@PRINCETON.EDU
Jonathan Cohen JDC@PRINCETON.EDU

Princeton University

Cognitive control is thought to guide activity along task-relevant pathways to achieve goal directed behavior. We use behavioral measures of similarity together with Representational Similarity Analysis and functional connectivity analysis of fMRI data to show how neural representations of semantic knowledge change to match the current goal demand. For instance, a horse may be more similar to a bear than to a dog in terms of size, but more similar to a dog in terms of domesticity. Here we present evidence that objects similar to each other in a given context are also represented more similarly in the brain and that these similarity relationships among representations are modulated by context specific activations in other parts of the brain. Subjects ranked 12 animals under two dimensions: domesticity and size. In the scanner, they compared among the animals under each dimension respectively. We compared, within subject, the similarity matrix computed from neural responses to each animal with that derived from their behavioral rankings. We found significant similarity between the two matrices for domesticity, and trending significance for size. We were able to replicate the effect in a separate set of subjects using fruits as the stimuli. We used taste and size as the dimensions and only found the effect in taste but not for size. We then used functional connectivity analysis to determine whether the neural similarity relationships were related to dimension-specific activity elsewhere in the brain. Using areas identified in the similarity analysis as seed regions, we examined correlations between spatial variance of these areas and activation in the rest of the brain during domesticity trials and during size trials. A significant number of subjects showed connectivity that was significantly higher during one dimension versus the other. The results were also replicated in the fruits group for both taste and size.

II-17. Optimal foraging as a framework for human task switching

Dominic Mussack MUSS0080@UMN.EDU
Robert Edge EDGERA@GMAIL.COM
Thomas Ringstrom RINGS034@GMAIL.COM
Paul Schrater SCHRATER@UMN.EDU

University of Minnesota, Twin Cities

One of the key problems that humans and animals solve is allocating time to different activities or goals, e.g. when to eat, work or play. Time allocation is a challenging control problem - we need to integrate the importance and availability of goals with our ability to complete them flexibly in response to changing demands and opportunities. Imagine a student with competing goals of playing a game and finishing a paper. The student works on the paper for a while before switching to the video game. They must decide both when to switch away from writing and what to switch to. By drawing on and extending optimal time allocation results in foraging theory, we develop a rational theory for task switching as an optimal time allocation process. Our work contrasts with standard task switching analyses, which view task switching as lapses in self-control or impulsivity that create performance costs due to switching, but not why people would voluntarily switch between tasks. We demonstrate the power of the approach by showing it provides novel explanations for the mysterious phenomena of task quitting near completion, and for when deadlines increase and decrease task completion. In particular we combine tasks that provide instantaneous rewards during the task and those that have goals on completion together in a unified framework. No prior work has been done in directly applying foraging ideas to human task switching behavior, and this approach synthesises usually disparate bodies of research. In addition we present a potential neurally plausible implementation for computing time allocation based on prediction error theories of dopaminergic action. Finally, incorporating intrinsic motivation ideas in artificial intelligence, we use the approach to predict human task engagement in a task switching game experiment.

II-18. The balance between evidence integration and probabilistic sampling in perceptual decision-making

Jozsef Fiser[1] FISERJ@CEU.HU
Jeppe Christensen[2] JEPPE.CHRISTENSEN@PSY.KU.DK
Máté Lengyel[3] M.LENGYEL@ENG.CAM.AC.UK

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

According to Probabilistic Sampling (PS, Fiser et. al 2010), to achieve any perceptual decision, the posterior distribution encoding features needs to be sampled through time. However, traditional evidence integration (EI) models also assume sequential integration of external sensory information over time (Gold & Shadlen, 2007). Which process shapes the trial-by-trial time course of human behavior? In a series of human behavioral experiments, we found that both processes are present during perceptual judgment and that their mutual influence on behavior is flexible. We used an estimation-based variant of the classical random dot motion (RDM) task, where in each trial, participants reported their best estimate of stimulus direction and their subjective uncertainty about

their decision simultaneously. We manipulated both global and local objective uncertainty and presented the RDM for durations between 100ms and 1.5sec. Confirming our analytical derivations, we found a significant and positive overall correlation between error and subjective uncertainty beyond 300-500msec as a function of time in all participants. As such positive correlations are the hallmark of PS and cannot be caused by EI, these results indicate that, indeed, probabilistic inference processes dominate the latter part of decision-making. Importantly, the transition between these segments shifted as a function of both local and global objective uncertainty. Thus, we propose that perceptual decision-making is not simply noisy evidence integration, but rather a probabilistic inference process. Moreover, this process in perceptual judgments follow a pattern that is related to that found during learning in an uncertain environment (e.g. McGuire et. al, 2014): When the global uncertainty is high, PS begins dominating earlier in time if local uncertainty is low compared to when local uncertainty is high. In contrast, when the global uncertainty is low — PS takes over at the same point in time regardless of the level of local signal uncertainty.

II-19. Generative models allow insight into perceptual accumulation mechanisms

Alexandre Hyafil[1,2] ALEXANDRE.HYAFIL@GMAIL.COM
Gustavo Deco[1] GUSTAVO.DECO@UPF.EDU
Jordi Navarra JORDI.NAVARRA@GMAIL.COM
Ruben Moreno-Bote[2] RMORENO@FSJD.ORG

[1]Universitat Pompeu Fabra
[2]Foundation Sant Joan de Deu

The mechanisms at play when the brain integrates sensory information over time to reduce uncertainty over percepts have been the focus of intensive study in the last two decades. Despite this effort, it has remained virtually impossible to experimentally tease apart distinct models for perceptual accumulation. This is largely due to the fact that in these models stochastic processes (either sensory or internal noise) play a large role and thus models only loosely constrain experimental data. Here we introduce a novel method based on generative models that reduces that part of stochasticity to its minimum and hence allows digging much more precisely into accumulation mechanisms. The method is applied in a context where, unlike the generic random-dot-motion task, perceptual information provided by each sensory sample can be quantified (Wyart et al., Neuron 2012). The method is based on parameter estimation through an Expectation-Maximization algorithm in generative models that combine three cognitive components: an accumulation stage (how sensory samples are integrated), a decision stage (including thresholds and noise) and a post-decision stage (modeling the time distribution from decision to the motor response). The method was applied to a speeded-reaction time task where subjects must judge the overall orientation of successive visual stimuli. Results very clearly arbitrate between different variants of the distinct model components both in terms of log-likelihood and fitting to experimental data characteristics. Notably we show that the decision threshold remains constant throughout stimulus presentation while decision noise increases sub-linearly, in line with diffusion-to-boundary hypothesis. Overall, the results open a promising path towards understanding of the refined mechanisms of perceptual accumulation. More generally, they illustrate the power of fitting generative models of behavior to human psychophysics data for unveiling cognitive mechanisms in action.

II-20. Contrasting roles of medial prefrontal cortex and secondary motor cortex in waiting decisions

Masayoshi Murakami MASAYOSHI.MURAKAMI@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Neuroscience Programme

Humans often succumb to an immediately available temptation while waiting for a better but delayed option. Here, we examined the role of medial prefrontal cortex (mPFC) and the secondary motor cortex (M2) in such waiting decisions. We trained rats to perform a waiting task in which on each trial the subject has a choice of either to wait for a randomly delayed tone to obtain a large reward or to give up waiting to obtain a smaller reward. The time rats were willing to wait was related to the mean tone delay but also varied substantially across trials. We then performed reversible pharmacological inactivations of mPFC and M2. While M2 inactivation had diverse effects on task performance, the effect of mPFC inactivation specifically affected the "patience" of the animal, reducing the average waiting time but not other performance parameters. We next compared recordings from neurons in M2 and mPFC. We previously reported that the activity of many M2 neurons was correlated with trial-by-trial fluctuations in waiting times. Here we performed recordings from mPFC. Thirty percent of mPFC neurons were activated during waiting period, consistent with a role of mPFC in patient waiting. However, surprisingly, neither amplitude nor the time course of these neurons was correlated with the strong trial-by-trial fluctuations in waiting time. Taken together, these results suggest that waiting behavior in this task is controlled by at least two different systems: One system, involving areas in the mPFC, is crucial for setting the average willingness to wait but does not drive precise decision timing. A second system, involving M2, translates these signals into precise trial by trial waiting times and is responsible for the variability from trial-to-trial.

II-21. Normative evidence accumulation in an unpredictable environment

Christopher Glaze CGLAZE@SAS.UPENN.EDU
Joseph Kable KABLE@PSYCH.UPENN.EDU
Joshua Gold JIGOLD@MAIL.MED.UPENN.EDU

University of Pennsylvania

Many decisions depend on the interpretation of noisy evidence that arrives sequentially over time. The temporal dynamics of this process can take many different forms, including perfect, leaky, and/or bounded evidence accumulation to identify stable signals and differentiation to detect abrupt signal changes. Existing models prescribe some of these dynamics under certain, restrictive conditions but have failed to account for if and how such diverse dynamics should be used more generally. Here we present a novel, normative model of decisions between two alternatives that can provide such an account. A key feature of the model is that the expected amount of instability in the environment governs decision dynamics. When perfect stability is expected, evidence is accumulated perfectly. Otherwise, evidence is accumulated with a leak and a bound that both depend on the expected level of instability. These dynamics reflect complementary roles for accumulation to identify signals and differentiation to identify changes and re-start accumulation. We show that human subjects can adjust their decision-making behavior according to these principles for two separate tasks. Both tasks required information accumulation in the presence of unpredictable change-points occurring at different rates but involved different timescales (tens of seconds versus hundreds of milliseconds). This work represents a new, empirically supported theoretical framework that provides a unified, normative account of multiple forms of decision dynamics based on expectations about environmental dynamics.

II-22. A causal role for dorsal hippocampus in a cognitive planning task in the rat

Kevin Miller[1] KJMILLER@PRINCETON.EDU
Matthew Botvinick[1] MATTHEWB@PRINCETON.EDU
Carlos Brody[1,2] BRODY@PRINCETON.EDU

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

Imagine you are playing chess. As you think about your next move, you consider the outcome each possibility will have on the board, and the likely responses of your opponent. Your knowledge of the board and the rules constitutes an internal model of the chess game. Guiding your behavior on the basis of model-predicted outcomes of your actions is the very definition of cognitive planning. It has been known for many decades that humans and animals can plan (Tolman, 1948), but the neural mechanisms of planning remain largely unknown. Recently, a powerful new tool for the study of planning has become available: the “two-step” task introduced by Daw et al. (2011). This task allows, for the first time, the collection of multiple trials of planned behavior within a single experimental session, opening the door to many new experimental possibilities. We have adapted the two-step task for use with rats, and developed a semi-automated pipeline to efficiently train large numbers of animals. Here, we show that the rodent two-step task reliably elicits planning behavior in rats, and that this behavior is impaired by inactivations of the dorsal hippocampus. Hippocampus is a structure long thought to play a role in navigational planning (O’Keefe & Nadel, 1974), but a causal role in planning has not yet been established. Our data demonstrate that dorsal hippocampus is critical for planning behavior in the rat. Inactivation experiments are ongoing in other structures, including orbitofrontal cortex and prelimbic PFC.

II-23. Separating population structure of movement preparation and execution in the motor cortex

Gamaleldin F Elsayed GAMALELDIN.ELSAYED@GMAIL.COM
Antonio Lara AHL2143@COLUMBIA.EDU
Mark Churchland MC3502@COLUMBIA.EDU
John Cunningham JPCUNNI@GMAIL.COM

Columbia University

Movement preparation decreases reaction time of voluntary movements. However, the neural computation performed during preparation remains debated. Many experiments explore preparation by introducing a delay period, during which subjects prepare a movement while holding their posture. A go cue then instructs subjects to execute the prepared movement. The use of a delay period allows the study of preparatory activity that is cleanly separated in time from subsequent movement activity; however, the extension of the study of preparatory activity to other contexts is challenging. A main difficulty is that preparatory and movement activities are mixed in single-neuron responses, making it difficult to study preparation beyond tasks that use an explicit delay period. Here, we investigate whether there exist exclusive preparation and movement neural dimensions (combination of activities of different neurons that preserves only preparatory or movement activity), allowing the separation of preparatory and movement activities at the population level. We test the hypothesis that motor cortex devotes different neural subspaces for movement preparation and execution (activities are separated at the population level), not a shared subspace (activities are mixed at the population level). We use PCA combined with a novel Monte Carlo method to identify preparation and movement subspaces and characterize their activities. This method utilizes data covariance to assess the significance of neural activity at a certain time in a given dimension within the high-dimensional neural space. Our results indicate that preparatory and movement activities are separated at the population level. Thus, not only is the preparatory subspace orthogonal to the inferred output dimensions (Kaufman et al. 2014), it is largely orthogonal to the entire movement activity subspace. This fact can be leveraged to study preparatory and movement activities independently, making it possible to assess the internal neural events that precede movement in a much greater variety of contexts.

II-24. Tonic vs. oscillatory multi-item memory maintenance: performances and capacity

Kenji Tanaka[1] TANAKA@MNS.K.U-TOKYO.AC.JP
Yasuhiko Igarashi[1] IGAYASU1219@MNS.K.U-TOKYO.AC.JP
Gianluigi Mongillo[2] GIANLUIGI.MONGILLO@PARISDESCARTES.FR
Masato Okada[1] OKADA@K.U-TOKYO.AC.JP

[1]The University of Tokyo
[2]Paris Descartes University

The amount of information that can be concurrently held active in working memory (WM) is limited. The origin of this limitation is a long-standing issue in memory research. According to one theoretical account, the network underlying memory maintenance possesses multiple preferred states of activity, each associated with enhanced spiking activity in the neuronal populations corresponding to the mnemonic representations of the stimuli. In another account, memory is maintained through activity-dependent, short-term changes in the properties of synaptic transmission, and can be held over long time scales through the periodic re-activation of the mnemonic representations. As the re-activation is short-lived, memory is maintained in the absence of significantly enhanced spiking. Here, we systematically compare multi-item performances in the two modes of active maintenance, i.e., tonic persistent activity (PA) vs. periodic population spikes (PSs), with particular regard to the mechanisms limiting capacity. Our main results are: (i) As long as short-term plasticity exhibits significant facilitation, WM implemented through periodic PSs is more robust against perturbations in the patterns of synaptic transmission than WM implemented through PA; (ii) WM implemented through periodic PSs reliably maintains the relative serial order of the items, unlike WM implemented through PA which only preserves the identity of the items; (iii) The two models make distinct experimentally testable predictions as to how the patterns of delay-period activity should change with increasing number of items, and as to the origin of the capacity limitations. In particular, multi-item WM implemented through periodic PSs predicts that spiking activity should become more temporally organized as the number of stored items increases. Important modifications should thus be apparent in the power spectrum, and possibly in the coefficients of variation of the spike trains, rather than in the average levels of spiking activity as predicted by multi-item WM implemented through PA.

II-25. Saccadic modulation of stimulus processing in primary visual cortex

James McFarland[1] JAMES.MCFARLAND@GMAIL.COM
Adrian Bondy[2] BONDYAL@MAIL.NIH.GOV
Bruce Cumming[2] BGC@LSR.NEI.NIH.GOV
Daniel Butts[3] DAB@UMD.EDU

[1]University of Maryland, College Park
[2]National Eye Institute, NIH
[3]University of Maryland

During natural vision primates actively sample the visual scene by making saccadic eye movements several times a second. Saccades have a profound impact on the spatiotemporal inputs to the visual system, and are critically important for visual perception and behavior. While perisaccadic 'corollary discharge' signals have been shown to modulate visual responses of neurons in higher cortical areas, previous studies have shown that V1 neurons respond similarly to a stimulus that is either flashed into their receptive field (RF) or introduced by a saccade, leading to the belief that V1 neurons encode the retinal stimulus independent of saccades. Here we revisit this important question leveraging two key innovations. First, we use an experimental design that decouples saccade-driven and stimulus-driven effects. Specifically, we presented uncorrelated sequences of 'one-dimensional' random bar patterns while animals made guided saccades parallel to the bar stimuli, thus leaving the spatiotemporal stimulus in a neuron's RF unaffected by saccades. Second, we utilize recently developed statistical models to analyze how saccades modulate the detailed stimulus processing of V1 neurons, rather than simply looking at their effects on mean firing rates. In contrast to previous studies, we find that saccades produce a robust biphasic modulation of V1 neuron firing rates that is composed of a divisive gain suppression followed by an additive firing rate increase. The net effect of these changes is a prolonged reduction in stimulus selectivity following saccades. This modulation is largely due to extra-retinal signals, and its timing (across neurons, across cortical lamina, and among the different components of each neuron's RF) suggests that it is inherited from the LGN. Microsaccades produced similar, but smaller, effects. These results thus highlight the importance of studying visual processing in the context of eye movements, and establish a foundation for integrating such effects into existing models of visual processing.

II-26. Comparing two models for activity-dependent tuning of recurrent networks

Julijana Gjorgjieva[1] GJORGJIEVA@BRANDEIS.EDU
Jan Felix Evers[2] JAN-FELIX.EVERS@COS.UNI-HEIDELBERG.DE
Stephen Eglen[3] S.J.EGLEN@DAMTP.CAM.AC.UK

[1]Brandeis University
[2]Ruprecht-Karls-Universität Heidelberg
[3]University of Cambridge

How do neural networks assemble to produce robust output? We address this question for recurrently connected networks of excitatory and inhibitory populations whose goal is to generate activity that propagates in a given direction, as is the case for many circuits in different nervous systems. Here, we investigate activity-dependent mechanisms for the tuning of recurrent networks that capture the propagation of peristaltic waves of muscle contractions that underlie crawling in Drosophila larvae. While mature Drosophila larvae exhibit rhythmic waves of coordinated muscle contractions, which pass along body segments during locomotion, this behavior is highly uncoordinated early in development. Noninvasive muscle imaging during Drosophila embryogenesis has shown that motor output gradually refines from spontaneous muscle contractions to neurally-controlled bursts of activity and peristaltic waves (Crisp et al., Development 2008). Moreover, manipulating the patterns of spontaneous activity has demonstrated that they influence maturation of the motor network (Crisp et al., J Neurosci 2011). This has implied that activity-dependent cues, contained in spontaneous activity, drive connectivity refinements in a network that ultimately produces crawling. We examine two plausible activity-dependent models which can tune weak connectivity in networks based on spontaneous activity patterns: (1) the Hebbian model, where coincident activity between neighboring neuronal populations strengthens the connections between them; (2) the homeostatic model, where connections are homeostatically regulated to maintain a constant level of postsynaptic activity. Our main results suggest the homeostatic model is the more likely candidate to produce networks in which locomotion patterns are generated robustly. Although the studied activity-dependent mechanisms modify connection strength based only on local activity patterns, they can generate a functional network which produces global propagating waves across the entire network, matching experimentally-observed properties of peristaltic waves in Drosophila larvae.

II-27. The neural representation of posterior distributions over hidden variables

Stephanie CY Chan SCYCHAN@PRINCETON.EDU
Kenneth A Norman KNORMAN@PRINCETON.EDU
Yael Niv YAEL@PRINCETON.EDU

Princeton University

The world is governed by unobserved variables, which generate the events that we do observe. For example, while sitting in a windowless lab we may not be able to observe the weather outside, but the weather influences whether we see people coming in to lab with umbrellas, or with snow in their hair, and so on. These observations allow us to make inferences about the hidden causes that generated them. In particular, we are able to maintain, approximately or exactly, a belief distribution over the hidden states, assigning varying levels of probability to each. Bayesian latent cause models provide a concrete way of describing this kind of inference about hidden variables (e.g., Gershman, Blei, & Niv, 2010). Using these models, the belief distribution can be computed as a posterior probability distribution over hidden states: P(hidden state | observations). We conducted an experiment in which participants viewed sequences of animals drawn from one of four “sectors” in a safari. They were tasked with guessing which sector the animals were from, based on previous experience with the likelihood of each animal in each sector. We used fMRI and representational similarity analysis (RSA) to investigate brain representations of the posterior distribution P(sector | animals). Results using RSA suggest that multivoxel patterns in the lateral orbitofrontal cortex, angular gyrus, and precuneus were best explained by representation of a posterior distribution over “sectors”. Further results suggest that a complementary set of areas were involved in the updating of the posterior distribution. These results are convergent with prior results implicating these areas in the representation of “state” from reinforcement learning (Wilson et al, 2014) and “schemas” or “situation models” (Ranganath & Ritchey, 2012).

II-28. Convergence of sleep and wake population activity distributions in prefrontal cortex over learning

Abhinav Singh[1] ABHINAV.SINGH@MANCHESTER.AC.UK
Adrien Peyrache[2] ADRIEN.PEYRACHE@GMAIL.COM
Mark Humphries[1] MARK.HUMPHRIES@MANCHESTER.AC.UK

[1]University of Manchester
[2]New York University

The inference-by-sampling hypothesis proposes that neural population activity at each point in time represents a sample from an underlying probability distribution. One key prediction is that the distribution during “spontaneous” activity (representing the prior) and during evoked activity (representing the posterior) converge over repeated experience. Just such a convergence has been observed in small populations from ferret V1 over development. Unknown is the extent to which this hypothesis is a general computational principle for cortex: whether it can be observed during learning, or in higher-order cortices, or during ongoing behaviour. We addressed these questions by analysing population activity from the prefrontal cortex of rats, learning rules in a Y-maze task. We focussed on sessions where the animal reached the learning criterion mid-session, allowing us to compare activity before and after learning. Our hypothesis was that the spontaneous activity in slow-wave sleep (SWS), in the absence of task-related stimuli and behaviour, constitutes the prior distribution. We find that the distributions were conserved across all epochs (SWS pre- and post-session and during task behaviour), in that the same population states appeared in each, consistent with our interpretation of SWS activity as a prior. Crucially, we find that the task-evoked distribution after learning was more similar to the distribution in post-session than in pre-session SWS epochs, consistent with convergence of the posterior and prior distributions over learning. We also find that the similarity between behaviour and post-session SWS distributions was larger for rewarded trials, but did not converge for pre-learning unrewarded trials, suggesting the distributions were directly updated by learning. Our results are thus evidence that inference-by-sampling can be observed over the course of learning, and is a potential general computational principle of cortex.

II-29. Combining behavior and neural data to model cognitive variables

Thomas Desautels[1] TADESAUTELS@GMAIL.COM
Timothy D. Hanks[2] THANKS@PRINCETON.EDU
Carlos Brody[2,3] BRODY@PRINCETON.EDU
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]Princeton University
[3]Howard Hughes Medical Institute

Many neurons in association areas are thought to reflect the results of internal computations that influence behavior. An archetypical example is in decision making, where the accumulation of temporal evidence appears to be reflected in both parietal and frontal areas. With limited neural data and many potential degrees of freedom, is it possible to identify the representation and dynamics of accumulation processes? Previous work exploited large volumes of behavioral data to estimate the parameters of a latent evidence accumulation process that could account for the choices made by an individual rat performing the Poisson clicks evidence accumulation task [1]. These parameters can be used to estimate the distribution over the evidence accumulation path on a single trial, and then to relate this distribution to concurrent electrophysiological data [2]. However, in this approach, information from the neural data is not used to identify either the parameters of the accumulator, nor single-trial variation. Thus, uncertainty in parameters and in internal state must be averaged, rather than resolved. We propose a method for using the neural data to aid in parameter and trajectory identification. We develop a generative model which connects stimulus and neural firing via a latent evidence accumulator of similar form to that used in previous studies [1 & 2] and a generalized linear model (GLM) with history-dependence, Poisson likelihood, and a log firing rate which is linear in the accumulator. We further develop a tractable iterative learning algorithm for this model using a Laplace approximation to marginalization over the latent. This method maximizes an approximation to the conditional log data likelihood or log posterior with respect to the model parameters, and thus can identify parameters of the latent accumulator. This method may be particularly useful for combining behavioral data with scarce and expensive neural data.

II-30. Deducing Hebbian adaption rules from the stationarity principle of statistical learning

Claudius Gros GROS07@ITP.UNI-FRANKFURT.DE
Rodrigo Echeveste ECHEVESTE@ITP.UNI-FRANKFURT.DE

Goethe University Frankfurt

Statistical learning rules, like Hebbian learning, deal with the problem of adapting (in an unsupervised manner) internal parameters, viz the synaptic weights. Statistical learning terminates generically when the post-synaptic neural activity becomes stationary, not changing any more its statistics for a given varying, but statistically stationary, input stream. At stationarity the post-synaptic neural activity hence becomes locally insensitive to further changes of the synaptic weights. The Fisher information is an information-theoretical functional measuring the sensibility of a given probability distribution function relative to changes in extrinsic parameters. The Fisher information of the distribution of the post-synaptic activity with respect to a suitable combination of the set of incoming synaptic weights can hence be regarded as an objective function for statistical learning. We propose to use the Fisher information with respect to the synaptic flux and show, that synaptic flux optimization leads to Hebbian learning rules for the afferent synaptic weights. By optimizing the Fisher information for the synaptic flux, Hebbian learning rules are derived having a range of interesting properties. They are intrinsically self-limiting and runaway synaptic growth does not occur. They perform a principal component analysis, whenever a dominant direction is present in the co-variance matrix of the afferent pre-synaptic neural activities. In the absence of a principal component, input directions with bimodal (bursting) neural activities are preferred, in the spirit of projection pursuit. The learning rules are inherently non-linear and capable to perform a non-linear independent component

analysis, without any modification, such as the bars problem.

II-31. An amygdalo-cortical circuit demonstrating hunger-dependent neural responses to food cues

Christian Burgess[1,2] CBURGESS@BIDMC.HARVARD.EDU
Rohan Ramesh[1] RAMESH@FAS.HARVARD.EDU
Kirsten Levandowski[1] KLEVAND1@BIDMC.HARVARD.EDU
Xiyue Wang[1] WXYUE@BU.EDU
Mark Andermann[1] MARKLAWRENCE_ANDERMANN@HMS.HARVARD.EDU

[1]Beth Israel Deaconess Medical Center
[2]Harvard University

Hunger can selectively enhance attention to food-associated cues. In modern society, where food cues are ubiquitous, this can lead to increased food intake and adverse health effects. Human neuroimaging studies demonstrate hunger-specific modulation of neural activity in temporal cortex when viewing pictures of food cues. However, the neural mechanisms that underlie this bias to motivationally relevant stimuli remain unclear. We used monosynaptic rabies tracing techniques to demonstrate strong reciprocal excitatory connections between postrhinal cortex (POR) and the lateral nucleus of the amygdala. We trained head-fixed, food-restricted mice in a go/no-go task to test how hunger selectively biases behavioral and neural responses to food-associated cues. Using two photon microscopy in a behaving animal, we separately imaged calcium activity in both POR neurons and lateral amygdala projections to POR. We imaged the same POR cell bodies or amygdala axons across days to investigate how changes in hunger state modulate neural responses to food cues. Preliminary results suggest that POR neurons are strongly visually responsive, and a subset of which these neurons are modulated by hunger. Amygdala responses are selective for food-associated visual stimuli, and appear more strongly modulated by hunger state. We propose that the amygdala, through reciprocal connections with cortex, may update and bias the value of food cue representations in POR in a hunger-dependent manner. As such, this amygdalo-cortical circuit might be part of a brain network for selective processing of motivationally-relevant sensory cues.

II-32. Multi-scale functional parcellation of cortex by unsupervised activation clustering

Liberty Hamilton[1] LIBERTY.HAMILTON@UCSF.EDU
Erik Edwards[2] ERIK.EDWARDS4@GMAIL.COM
Edward Chang[2] CHANGED@NEUROSURG.UCSF.EDU

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

The human neocortex consists of billions of neurons that interact in complex networks at multiple spatial scales during natural behaviors. During speech perception, for example, interactions at small scales may occur locally within auditory areas, and interactions at larger scales occur between more distant brain areas. Here, we develop a data-driven computational method to uncover such functional regions of interest at multiple spatial scales and use this to describe inter- and intra-areal functional interactions during natural speech perception. We apply this method to intracranial recordings from 256-channel high-density electrode arrays in human epilepsy patients. These regions are defined by soft clustering of functional activations, rather than by gross anatomical labeling or strict categorical boundaries. Moreover, we avoid the requirement of specifying one particular spatial scale, or the exact number of components to retain, leading to a “multi-scale” method requiring minimal user input and parameter choices. The method is robust, unsupervised, and computationally efficient. We present the essential aspects of our method and examples of the results.

II-33. Emergence of a reliability code in the barn owl's midbrain

Fanny Cazettes[1] FANNY.CAZETTES@PHD.EINSTEIN.YU.EDU
Brian Fischer[2] BFISCHER.SU@GMAIL.COM
Jose Pena[1] JOSE.PENA@EINSTEIN.YU.EDU

[1]Albert Einstein College of Medicine
[2]Seattle University

Perception can be seen as probabilistic inference. This implies that the brain generates an internal model of stimulus reliability. Theoretical work has suggested that reliability could be captured in either the magnitude (gain) or the selectivity of neural responses (width of tuning curves) but, so far, it has not been shown experimentally how reliability is represented. Here we demonstrate that neurons' selectivity can account for the coding of cue reliability in the owl's sound localization system. We focus on the interaural time difference (ITD), a primary cue used to localize sound in horizontal space. We estimated the reliability of the auditory cue over different environmental contexts. We show that the widths across the population of space-specific neurons, but not the gain, match the ITD reliability across horizontal sound-source location. We demonstrate that this effect results from the auditory system adapting its basic organizational principle, tonotopy, to represent reliability along with space. Additionally, by manipulating the reliability of ITD, we further confirm that the tuning curve widths increase as the reliability of ITD decreases. Using a model, we demonstrate that spatial-tuning width captures reliability dynamically by virtue of the location-dependent frequency tuning and the cross-correlation mechanism that generate the ITD selectivity. Finally, we tested the ability of different decoding models to predict the owl's behavior as ITD reliability varies. We show that only a decoding model that assumes reliability is encoded in the tuning widths could explain the owl's behavior. Thus we provide a case for a sensory system representing cue reliability, where it occurs and how it emerges.

II-34. Predictiveness and prediction in classical conditioning: a Bayesian statistical model

Kristin Volk[1] KRISTIN@GATSBY.UCL.AC.UK
Peter Dayan[2] DAYAN@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University College London

Classical conditioning is a rather pure form of prediction learning. Complex facets of conditioning are increasingly being seen as components of statistically-sound generative and recognition models of prediction rather than as algorithmic affectations or implementational infelicities. We focus on a critical facet that lacks a statistical treatment, namely that conditioned stimuli (CSs) not only make different predictions, but also can be differentially predictive (Mackintosh & Turner, 1971; Mackintosh, 1975). That is, some stimuli (and whole stimulus dimensions) are relevant as predictors in one or more contexts; other stimuli are not; determining and excluding the latter is statistically mandatory, with beneficial side-effects of speeding and simplifying learning. We show how this helps provide an ecumenical answer to the early (but continuing; Le Pelley et al., 2010) dispute in conditioning concerning which predictors should exhibit more plasticity: those whose predictions are currently good (Mackintosh, 1975) or bad (Pearce & Hall, 1979; 1980). We formalize the notion of predictiveness in a generative model in which each CS is awarded two (hidden) random variables. One is a binary indicator variable indicating the stimulus' predictiveness; the other a real-valued weight indicating the current association between the stimulus and the outcome (which is assumed to evolve according to the conventional dynamics underpinning the Kalman filter, thus providing the link to Pearce & Hall, 1980's theory; Dayan et al., 2000). The net prediction is then generated according to weights associated with those CSs that are both present and predictive. The model thus seamlessly integrates structure and parameter learning, all guided by the environment. We illustrate the workings of the model using the important paradigm of Latent Inhibition, and then demonstrate its capacity to reproduce many other standard conditioning paradigms including Blocking, Extinction, Overshadowing and Backwards Blocking.

II-35. Neural representation of complex associative structures in the human brain

Mona Garvert[1] MONA.GARVERT.11@UCL.AC.UK
Ray Dolan[2] R.DOLAN@UCL.AC.UK
Timothy EJ Behrens[3] BEHRENS@FMRIB.OX.AC.UK

[1]Wellcome Trust Centre for Neuroimaging
[2]University College London
[3]University of Oxford

Goal-directed behaviour requires a neural representation of the associations between objects, events and other types of information. The mechanisms underlying the association of pairs of objects are well characterized, and involve an increase in the similarity of the respective object representations. Much less is known about how the human brain stores multiple associations that form a more complex global structure. Does the cortex continue to store simple associative links, or is global knowledge about the relationship between objects that have not been directly associated nevertheless incorporated in the representation? Here, we address this question from a functional perspective and ask whether a signature of a global structure is apparent following an implicit learning paradigm. We presented human participants with sequences of objects with stimulus transitions drawn random walks along a grid. We find that reaction times in response to random stimulus transitions on a subsequent day reflect distances between objects on the grid, even for items never directly associated. This suggests that humans acquire implicit knowledge about global structure rather than storing associations alone. Functional magnetic resonance imaging (fMRI) repetition suppression analyses show this behavioural effect is mirrored by a distance-dependence of representational similarity for object representations in an hippocampal-entorhinal network. These behavioural and neural effects can be explained by a simple Hopfield network with auto-associative attractors and Hebbian plasticity between associated objects. This suggests that global knowledge about the relationship between non-associated objects emerges through increases in representational similarity for pairwise associations. Thus, our results demonstrate that global information about complex associative structures is represented in the hippocampus and can be interrogated using human fMRI repetition suppression.

II-36. Rapid sensorimotor transformation and song intensity modulation in Drosophila

Philip Coen PIPCOEN@GMAIL.COM
Mala Murthy MMURTHY@PRINCETON.EDU

Princeton University

Acoustic communication signals are prevalent throughout the animal kingdom. Many organisms modulate their signal amplitude in accordance with the distance of their partner to optimize for receiver intensity and conserve energy. Previously, this sensorimotor transformation has been identified in higher vertebrates, but never in lower organisms. While this may indicate a lack of capacity, it's also possible that technical challenges inherent to studying naturalistic acoustic communication have obscured such behaviors. Here, we use a novel assay to show that Drosophila perform this complex sensorimotor transformation, and we combine genetic tools with modeling to start dissecting the underlying neural computations for the first time in any system. Drosophila courtship song comprises two modes: pulse and sine. Pulse mode consists of pulse trains with an inter-pulse interval of approximately thirty-five milliseconds. We not only demonstrate that male flies actively increase pulse amplitude with distance from the female, but that this sensorimotor transformation is performed in milliseconds to modulate the production of ongoing acoustic signals on a pulse-by-pulse basis. Through genetic and physical manipulations, we establish that amplitude modulation at large distances is exclusively mediated by vision, but is independent of the canonical motion detection pathways. Further, we identify the indirect flight muscles as a critical component of the motor output. To examine song without confounding sensory cues, we thermogenetically activated males in the absence of any female. Even artificially activated flies increase amplitude with target distance, and global

changes in light intensity produced transient amplitude responses. Building on this success, we developed the first system for recording activated song while flies walk on an air-supported ball to map the stimulus-response transformation. Our results identify an unprecedented level of complexity in insect acoustic communication and establish Drosophila as a model system to study the neural interface between visual information and acoustic signal production.

II-37. Speech recognition using neural activity in the human superior temporal gyrus

David Moses[1] DAVID.MOSES@UCSF.EDU
Nima Mesgarani[2] NIMA@EE.COLUMBIA.EDU
Edward Chang[1] CHANGED@NEUROSURG.UCSF.EDU

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

The superior temporal gyrus (STG) plays a key role in human language processing. Previous studies have been performed that attempt to reconstruct speech information from brain activity in this region, but none of them incorporate the probabilistic framework and engineering methodology used in modern speech recognition systems. This work describes the first efforts towards the design of a neural speech recognition (NSR) system that performs phoneme recognition using the high gamma band power of neural activity in the STG as features. It implements a Viterbi decoder that incorporates phoneme likelihood estimations from a linear discriminant analysis model and transition probabilities from an n-gram phonetic language model. The system exhibited significant performance gains when using spatiotemporal representations of the activity instead of purely spatial representations, which is indicative of the importance of modeling the temporal dynamics of neural responses when analyzing their variations with respect to varying stimuli. In its current state, the system achieves a frame-by-frame accuracy of 21.23% and a phoneme error rate of 78.86% for one of the subjects. Although decoded speech is not yet intelligible, the NSR system could be further developed into a speech prosthetic that restores some communicative capabilities to impaired patients, such as those with locked-in syndrome.

II-38. Genetic networks specifying the functional architecture of orientation domains in V1

Joscha Liedtke[1,2] JOSCHA@NLD.DS.MPG.DE
Fred Wolf[1] FRED@NLD.DS.MPG.DE

[1]MPI for Dynamics and Self-Organization
[2]BCCN, Goettingen

Although genetic information is critically important for brain development and structure, it is widely believed that neocortical functional architecture is largely shaped by activity dependent mechanisms. The information capacity of the genome simply appears way too small to contain a blueprint for hardwiring the cortex. Here we show theoretically that genetic mechanisms can in principle circumvent this information bottleneck. We find in mathematical models of genetic networks of principal neurons interacting by long range axonal morphogen transport that morphogen patterns can be generated that exactly prescribe the functional architecture of the primary visual cortex (V1) as experimentally observed in primates and carnivores. We analyze in detail an example genetic network that encodes the functional architecture of V1 by a dynamically generated morphogen pattern. We use analytical methods from weakly non-linear analysis complemented by numerical simulation to obtain solutions of the model. In particular we find that the pinwheel density variations, pinwheel nearest neighbor distances and most strikingly the pinwheel densities are in quantitative agreement with high precision experimental measurements. We point out that the intriguing hypothesis that genetic circuits coupled through axonal transport shape the complex archi-

tecture of V1 is in line with several biological findings. (1) Surprisingly, transcription factors have been found to be transported via axons and to be incorporated in the nucleus of the target cells. (2) A molecular correlate was recently found for ocular dominance columns in V1. (3) We estimate that the speed of axonal transport is rapid enough to achieve appropriate timescales. This theory opens a novel perspective on the experimentally observed robustness of V1's architecture against radically abnormal developmental conditions such a dark rearing. Furthermore, it provides for the first time a scheme how the pattern of a complex cortical architecture can be specified using only a small genetic bandwidth.

II-39. Steps towards quantifying the vibrisso-tactile natural scene

Jennifer Hobbs JENNIFERHOBBS2008@U.NORTHWESTERN.EDU
Hayley Belli HAYLEYBELLI2016@U.NORTHWESTERN.EDU
Mitra Hartmann HARTMANN@NORTHWESTERN.EDU

Northwestern University

Analysis of natural scene statistics has been a powerful approach for understanding neural coding in the visual and auditory systems. In the field of somatosensation, it has been more challenging to quantify the natural tactile scene, in part because somatosensory signals are tightly linked to the animal's movements. Here we describe two steps towards quantifying the natural tactile scene for the rat vibrissal system. First, we simulated rat whisking motions to systematically investigate the probabilities of whisker-object contact in naturalistic environments. The simulations permit an exhaustive search through the complete space of possible contact patterns, thereby allowing for an estimation of the patterns that would occur during long sequences of natural exploratory behavior. We specifically quantified the probabilities of "concomitant contact," that is, given that a particular whisker makes contact with a surface during a whisk, what is the probability that each of the other whiskers will also make contact with the surface during that whisk? Probablities of concominant contact were quantified in simulations that assumed increasingly naturalistic conditions: 1), the space of all possible head poses; 2) the space of behaviorally-preferred head poses measured experimentally; and 3) the space of common head poses adopted in a cage and a burrow. As environments became more naturalistic, the probability distributions shifted from exhibiting a "row-wise" structure to a more diagonal structure. Second, we are now extending this work to incorporate vibrissal mechanics. Towards this goal, we have quantified a variety of whisker parameters, including base and tip diameters, arc length, density, and medulla size and shape. Scaling of these parameters across the vibrissal array will have significant effects on rat exploratory behavior, active-sensing, and the energetic costs associated with rhythmic whisking.

II-40. Behavioral correlates of combinatorial versus temporal features of odor codes

Debajit Saha SAHAD@SEAS.WUSTL.EDU
Chao Li LI.CHAO@WUSTL.EDU
Steven Peterson STEVEN.PETERSON@WUSTL.EDU
William Padovano WPADOVANO@WUSTL.EDU
Nalin Katta NK1@WUSTL.EDU
Barani Raman BARANI@WUSTL.EDU

Washington University in St. Louis

Most sensory stimuli evoke spiking responses that are distributed across neurons and are temporally structured. Whether neural circuits modulate the temporal structure of ensemble activity to facilitate different computations is not known. Here, we investigated this issue in the insect olfactory system. We found that an odorant can generate synchronous or asynchronous spiking activity across a neural ensemble in the antennal lobe circuit

depending on its relative novelty with respect to a preceding stimulus. Regardless of variations in temporal spiking patterns, the activated combinations of neurons robustly represented stimulus identity. Consistent with this interpretation, locusts robustly recognized both solitary and sequential introductions of trained odorants in a quantitative behavioral assay. However, the behavioral responses across locusts were predictable only to a novel stimulus that evoked synchronized spiking across neural ensembles. Hence, our results indicate that a combinatorial code encodes for stimulus identity, whereas the temporal structure selectively emphasize novel stimuli.

II-41. Visually evoked gamma rhythms in the mouse superior colliculus

Shinya Ito SHIXNYA@GMAIL.COM
Alan Litke ALAN.LITKE@CERN.CH

University of California, Santa Cruz

Understanding the detailed circuitry of functioning neuronal networks is one of the major goals of neuroscience. The recent maturation of in vivo neuronal recording techniques has made it possible to record the spiking activity from a large number of neurons simultaneously with sub-millisecond temporal resolution. Here we used a 64-channel high-density silicon probe to record, in each preparation, the activity from ~85 neurons in the superior colliculus (SC) of a mouse under anesthesia. To probe the correlation structure of many neurons, we employed a wavelet transform of the cross-correlogram to categorize the functional connectivity in different frequency ranges. We found that strong gamma rhythms, both in local field potentials (LFPs) and in the spiking activity, were induced by various visual stimuli such as drifting gratings, contrast modulated noise movies, checkerboard stimuli, and full-field flashes. The gamma rhythms started in the superficial SC, and then propagated down towards the deeper layers of the SC. The phase of the gamma rhythm between the superficial layers and the deeper layers was inverted after ~200 ms, a property of the mouse SC that has not been previously reported.

II-42. Characterization of very large ganglion cell populations in the mouse retina

Sahar Pirmoradian[1] SPIRMORA@INF.ED.AC.UK
Gerrit Hilgen[2] GERRIT.HILGEN@NCL.AC.UK
Martino Sorbaro[1] MARTOPIX@GMAIL.COM
Oliver Muthmann[3] JENSOLIVER@NCBS.RES.IN
Upinder S Bhalla[4,5] BHALLA@NCBS.RES.IN
Evelyne Sernagor[6] EVELYNE.SERNAGOR@NEWCASTLE.AC.UK
Matthias Hennig[1] MHENNIG@INF.ED.AC.UK

[1]University of Edinburgh
[2]University of Newcastle
[3]Manipal University
[4]National Centre for Biological Sciences
[5]Tata Institute of Fundamental Research
[6]University of Newcastle upon Tyne

Morphological and physiological analyses indicate the retinal population response consists of a number of separate channels, each represented by a different ganglion cell (RGC) type with distinct functional characteristics. To characterize these pathways in more detail, we recorded light responses from the mouse retina with a high-density multi-electrode array with 4096 channels (64x64 channels, pitch 42 μm). This allowed simultaneous monitoring of the activity of thousands of RGCs in about 1/3 of a mature retina. In these recordings, signals from single neurons are typically detectable on multiple, nearby channels. We present a new method to exploit this to improve the signal to noise ratio for spike detection, and to estimate a current source location for each spike. This yields a map of neural activity with much higher spatial resolution than provided by the array, where spikes from individual neurons form dense, isolated clusters. These were separated into single units using Mean Shift clustering. Direct comparison with raw data shows this is a new, highly efficient method for spike sorting requiring minimal manual intervention. We then quantified light responses using full field stimulation and linear models derived from white noise stimulation. Although broadly distributed, response kinetics had a clear dorsoventral gradient: RGCs in more ventral locations responded more slowly, and receptive fields of Off cells were larger in ventral than in dorsal locations. It is unclear whether this specificity reflects varying properties within certain cell classes, as for example receptive field sizes in the primate retina, or different cell classes in different locations. Moreover, unlike in other mammalian species, we found a larger number of On than Off cells. Overall our results demonstrate substantial region specificity and functional specialization in the retina, most likely reflecting ecological requirements.

II-43. Conjunctive population codes: insights from bat 3D head-direction cells

Arseny Finkelstein[1] ARSENYF@GMAIL.COM
Johnatan Aljadeff[2,3] ALJADEFF@UCSD.EDU
Alon Rubin[1] ALON.RUBIN@GMAIL.COM
Nachum Ulanovsky[1] NACHUM.ULANOVSKY@WEIZMANN.AC.IL
Misha Tsodyks[1] MISHA@WEIZMANN.AC.IL

[1]Weizmann Institute of Science
[2]University of California, San Diego
[3]Salk Institute for Biological Studies

A multi-dimensional stimulus space can be partitioned using different strategies, implying different shapes for neural tuning-curves. Recently we found that 3D head-direction - a two dimensional variable - is encoded in the bat by two neuronal subpopulations - pure (1D) azimuth and pitch cells vs. conjunctive (2D) azimuth-pitch cells - forming an overcomplete representation. Here, using analytical calculations and numerical simulations of a linear decoder, we show that these subpopulations are needed to efficiently represent the head-direction in different behaviorally-relevant regimes. Specifically, we found that conjunctive cells can be used to decode the stimulus more accurately compared to pure cells for short integration times. Our results suggest that optimal tuning-curve shape can strongly depend on a system's dynamic variables (e.g. integration time in the case of head-direction cells), which is likely to be relevant to other circuits encoding multi-dimensional stimuli.

II-44. Parallel remodeling of direct and indirect pathway neuronal activity during Pavlovian conditioning

Jesse Marshall[1] JESSEM1@STANFORD.EDU
Jones Parker[2] JONES.G.PARKER@GMAIL.COM
Biafra Ahanonou[1] BAHANONU@STANFORD.EDU
Benni Grewe[1] GREWE@STANFORD.EDU
Jin Zhong Li[1] JINZHONG@STANFORD.EDU
Mike Ehlers[2] MICHAEL.EHLERS@PFIZER.COM
Mark Schnitzer[1,3] MSCHNITZER@GMAIL.COM

[1]Stanford University
[2]Pfizer Inc.
[3]Howard Hughes Medical Institute

During learning, selective alterations in neuronal excitability drive adaptive responses to environmental stimuli. It is thought that within the nucleus accumbens (Acb), differential modulation of D1- and D2-dopamine receptor

expressing medium spiny neurons (MSNs) by stimulus-elicited dopamine release is critical for plasticity during reinforcement learning and the execution of goal-directed behavior. However the inability of existing techniques to differentiate the two cell types in vivo has prevented validation of this hypothesis. We imaged the simultaneous Ca2+ dynamics of hundreds of D1- or D2-MSNs during acquisition, extinction, and reinstatement of an appetitive Pavlovian association in freely moving mice. We found that both pathways were inhibited during reward consumption and that spatially coordinated ensembles of D1- and D2-MSNs were activated in response to conditioned stimulus presentation. After learning these responses diminished in overall amplitude, selectively responding on trials in which animals switched between passive and goal directed behavior. Systemic D1 receptor antagonism reversed the observed decrease in the amplitude of neural responses to the conditioned stimulus, suggesting that the observed remodeling of Acb network activity is dopamine-dependent. Taken together, our results suggest that conditioned approach is driven by spatio-temporally coordinated ensembles of D1- and D2-MSNs that act in concert to drive goal directed behavior, suggesting substantial revisions to existing models of the anatomical substrates of reinforcement learning in the basal ganglia.

II-45. Role of cortico-cerebellar communication in motor control.

Adam Hantman[1,2] HANTMANA@JANELIA.HHMI.ORG
Jian-Zhong Guo[1,2] GUOJ@JANELIA.HHMI.ORG
Allen Lee[1,2] ALLENLEETC@GMAIL.COM
Austin Graves[1,2] GRAVESA@JANELIA.HHMI.ORG
Kristin Branson[1,2] BRANSONK@JANELIA.HHMI.ORG

[1]Janelia Farm Research Campus
[2]Howard Hughes Medical Institute

Recurrent pathways from motor control to somatosensory centers have been posited to serve two main functions. 1) Motor-sensory feedback projections could selectively cancel self-generated sensory information, a critical step in error detection. 2) Recurrent pathways could facilitate motor planning by predicting the sensory consequences of self-generated movements. Efforts to test these proposed functions have been hindered by the inability to specifically manipulate feedback projections to sensory pathways. We have overcome this obstacle by gaining genetic access to the basal pontine nucleus, a major recurrent relay station between the cortex and cerebellum. We have examined the effects of basal pontine nucleus manipulation in a multistep motor task, the forelimb reach task. Mice subjected to acute silencing of the basal pontine nucleus exhibited abnormalities and inefficiencies in planning and execution of the forelimb reach task. The resulting deficits are consistent with a switch from a predictive to a reactive motor control strategy.

II-46. Tuning of the visual system to the curvature of natural shapes

Ingo Frund MAIL@INGOFRUEND.NET
James Elder JELDER@YORKU.CA

York University

Principles of efficient coding have been successful in predicting physiological properties of the visual system, but this success has largely been confined to early visual processing. Here we ask whether these principles can be extended to the coding of shape information in a higher cortical area. We introduce a generative model for shape and learn the curvature parameters of this model from the statistics of natural shapes. We find that natural curvature follows a power law distribution, reflecting underlying scale invariance properties. Using a shape discrimination task, we show that the human visual system is finely tuned to these statistics, and comparison with the population response of neurons in macaque Area V4 reveals an efficient code for this power law. Together, these results suggest that the efficient coding hypothesis, so successful in explaining functional properties of the

early visual system, can be extended to higher cortical areas encoding object shape.

II-47. Maximal causes for a masking based model of STRFs in primary auditory cortex

Abdul-Saboor Sheikh[1,2] SHEIKH@TU-BERLIN.DE
Zhenwen Dai[3] Z.DAI@SHEFFIELD.AC.UK
Nicol Harper[4] NICOL.HARPER@GMAIL.COM
Richard Turner[5] RET26@CAM.AC.UK
Joerg Luecke[2] JOERG.LUECKE@UNI-OLDENBURG.DE

[1]Technical University of Berlin
[2]University of Oldenburg
[3]University of Sheffield
[4]University of California, Berkeley
[5]University of Cambridge

Neural response properties in the primary auditory cortex of mammals (A1) are commonly characterized by spectro-temporal receptive fields (STRFs). Like simple cell responses in visual cortex, A1 responses have been linked to the statistics of their sensory input. Hypothesizing that A1 represents the structural primitives of auditory stimuli, in this study we employ a statistical model which can represent the non-linear interaction of auditory components in cochleagram formation. Such non-linear interactions give rise to psycho-acoustic masking effects, which have been successfully exploited in technical applications such as source separation. Underlying such masking effects and applications is a characteristic of cochleagrams which, for any time-frequency interval, allows only for a single component to be dominant under natural conditions. Applications to source separation exploit this property by assigning each time-frequency interval to the one sound source or component that exhibits maximal energy, a procedure sometimes referred to as log-max approximation. While a combination based on maximal sound energies seems to be a natural choice for cochleagram encoding, linear superposition assumption has remained a standard choice for STRF modeling. In this work we use a non-linear sparse coding model based on the log-max approximation and study its predictions for STRFs in A1. The used model is a version of maximal causes analysis (MCA), which combines spectro-temporal features using the maximum. After fitting model parameters using truncated EM, we estimate the predicted STRFs through regularized reverse correlation. The estimated STRFs are largely found to be localized in time, frequency or both including FM sweeps. Furthermore, their comparison with recent in vivo recordings of ferret A1 reveals structural similarities and consistencies in spectro-temporal modulation tuning.

II-48. A dual algorithm for olfactory computation in the insect brain

Sina Tootoonian ST582@ENG.CAM.AC.UK
Máté Lengyel M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

We study the early locust olfactory system in an attempt to explain its well-characterized structure and dynamics. We first propose its computational function as recovery of high-dimensional sparse olfactory signals from a small number of measurements. Detailed experimental knowledge about this system rules out standard algorithmic solutions to this problem. Instead, we show that solving a dual formulation of the corresponding optimisation problem yields structure and dynamics in good agreement with biological data. Further biological constraints lead us to a reduced form of this dual formulation in which the system uses independent component analysis to continuously adapt to its olfactory environment to allow accurate sparse recovery. Our work demonstrates the challenges and rewards of attempting detailed understanding of experimentally well-characterized systems in a

principled computational framework.

II-49. A model of spinal sensorimotor circuits recruited by epidural stimulation of the spinal cord

Emanuele Formento EMANUELE.FORMENTO@EPFL.CH
Marco Capogrosso MARCO.CAPOGROSSO@EPFL.CH
Eduardo Martin Moraud EMARTINMORAUD@GMAIL.COM
Miroslav Caban MIROSLAV.CABAN@EPFL.CH
Gregoire Courtine GREGOIRE.COURTINE@EPFL.CH
Silvestro Micera SILVESTRO.MICERA@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Electrical spinal cord stimulation (ESCS) can improve motor control after various neurological disorders. However the basic mechanisms involved in the regain of loco-motor abilities remain unclear. In a previous study, in rodents, we demonstrated that the main structures recruited by ESCS in the lumbar region are large proprioceptive afferents. Here we exploit these results to understand the effect on the main sensorimotor circuits involved in the generation of movement. We designed a realistic biological-inspired model of the spindles (Group Ia and II) reflex network of a couple of agonist-antagonist muscles encompassing a realistic model of alpha-Motoneurons, Ia inhibitory interneurons, group II excitatory interneurons, group Ia and group II afferents. The network receives inputs from ESCS computing the induced firing rate in the recruited afferents. In parallel, the natural firing rate of the afferent during stepping, is estimated by the use of a realistic biomechanical model of the rat hindlimb providing the variation of the muscle fibers during stepping. The model output suggests that the stretch reflex alone can explain the well-known modulation of ESCS-induced spinal reflexes during locomotion, showing a strong dependency of the stimulation output from the actual state of the sensory system. This property explains the stereotyped behavior of spinal-rats during ESCS induced locomotion on treadmill where the stepping adapts to belt velocities and other sensory inputs. We also demonstrate that given the agonist-antagonist alternate and reciprocal inhibition that modulates the stimulation output, phase detection of gait should be exploited in the design of real time strategies for closed loop neuromodulation of the spinal cord to modulate independently flexion and extension phases in real-time.

II-50. Self-organized mechanisms of the head direction sense

Adrien Peyrache ADRIEN.PEYRACHE@GMAIL.COM
Marie Lacroix MARIELACROIX.ENS@GMAIL.COM
Peter Petersen PETERSEN.PETER@GMAIL.COM
Gyorgy Buzsaki GYORGY.BUZSAKI@NYUMC.ORG

New York University

The head direction (HD) system functions as a compass with member neurons robustly increasing their firing rates when the animal's head points in a specific direction, a signal believed to be critical for navigation. Although computational models assume that HD cells form an attractor, experimental support for such mechanism is lacking. We addressed the contributions of stimulus-driven and self-generated activity by recording ensembles of HD neurons in the antero-dorsal thalamic nucleus and the postsubiculum of mice in various brain states. The temporal correlation structure of HD neurons is preserved during sleep, characterized by a 60 degree-wide activity packet, both within as well as across these two brain structures. During REM, the spontaneous drift of the activity packet was similar to that observed during waking and was accelerated ten-fold during slow wave sleep. These findings demonstrate that peripheral inputs impinge upon a self-organized network, which provides amplification and enhanced precision of the head-direction signal.

II-51. Highly non-uniform information transfer in local cortical networks

Sunny Nigam SUNNY.NIGAM.INDIANA@GMAIL.COM
Masanori Shimono MNS@P.U-TOKYO.AC.JP
Olaf Sporns OSPORNS@INDIANA.EDU
John Beggs JMBEGGS@INDIANA.EDU

Indiana University Bloomington

The performance of complex networks depends on how they route their traffic. Despite the importance of routing, it is virtually unknown how information is transferred in local cortical networks, consisting of hundreds of closely-spaced neurons. To properly address this, it is necessary to record simultaneously from hundreds of neurons at a spacing that matches typical axonal connection distances, and at a temporal resolution that matches synaptic delays. We used a 512 electrode array (60μm spacing) to record spontaneous activity sampled at 20 kHz from up to 700 identified neurons simultaneously in slice cultures of mouse somatosensory cortex (n = 15) for 1 hr at a time. We used a previously validated version of Transfer Entropy to quantify Information Transfer (IT) between pairs of neurons. Similar to in vivo reports, we found an approximately lognormal distribution of firing rates. Pair-wise IT strengths also were nearly lognormally distributed, similar to reports of synaptic strengths. We observed that IT strengths coming into, and going out of, cortical neurons were correlated, consistent with the predictions of computational studies which have shown that such correlations are necessary for obtaining a lognormal distributions of firing rate from a lognormal distribution of synaptic weights. Neurons with the strongest total IT (out/in) were significantly more connected to each other than chance, thus forming a "rich club" network. The members of the rich club accounted for 70% of the total IT in the networks and also made Information routing more efficient by lowering the overall path lengths of the network. This highly unequal distribution of IT has implications for the efficiency and robustness of local cortical networks, and gives clues to the plastic processes that shape them

II-52. Structure of mammalian grid-cell activity for navigation in 3D as predicted by optimal coding

Andreas Herz[1,2] HERZ@BIO.LMU.DE
Alexander Mathis[3] AMATHIS@FAS.HARVARD.EDU
Martin Stemmler[1,2] STEMMLER@BIO.LMU.DE

[1]Bernstein Center for Computational Neuroscience, Munich
[2]Ludwig-Maximilians-Universität, Munich
[3]Harvard University

Lattices abound in nature — from the beautiful crystal structure of minerals to the astounding honey-comb organization of ommatidia in the compound eye of insects. Such regular arrangements provide solutions for dense packing, efficient resource distribution and cryptographic schemes — and highlight the importance of lattice theory in research fields as distinct as mathematics and physics, biology and economics, and computer science and coding theory. Are geometric lattices also of relevance for how the brain represents information? To answer this question, we focus on higher-dimensional stimulus domains, with particular emphasis on neural representations of the physical space explored by an animal. Using information theory, we ask how to optimize the spatial resolution of neuronal lattice codes. We show that the hexagonal activity patterns of 'grid cells' found in the hippocampal formation of mammals navigating on a flat surface lead to the highest spatial resolution in a two-dimensional world. For species that move freely in their three-dimensional environment, firing fields should be arranged along a face-centered cubic (FCC) lattice or a equally dense non-lattice variant thereof known as a hexagonal close packing (HCP). This quantitative prediction could be tested experimentally in flying bats, arboreal monkeys, or marine mammals. New results from the Ulanovsky lab [1] provide first evidence that grid cells in bats navigating in three-dimensional habitats indeed exhibit the predicted symmetries. More generally, our results suggest that populations of grid-cell-like neurons whose activity patterns exhibit lattice structures at multiple, nested scales [2] allow the brain to encode higher-dimensional sensory or cognitive variables with high efficiency.

[1] G. Ginosar, A. Finkelstein, L. Las, N. Ulanovsky: 3-D grid cells in flying bats. Annual meeting of the Society for Neuroscience. Poster 94.22/SS42 (2014) [2] See also A. Mathis, M.B. Stemmler, A.V.M. Herz: Probable nature of higher-dimensional symmetries underlying mammalian grid-cell activity patterns. arXiv:1411.2136 (2014)

II-53. A race model for singular olfactory receptor expression

Daniel Kepple DKEPPLE@CSHL.EDU
Ivan Iossifov IOSSIFOV@CSHL.EDU
Alexei Koulakov KOULAKOV@CSHL.EDU

Cold Spring Harbor Laboratory

In vertebrates, olfactory sensory neurons select only one olfactory receptor (OR) to produce out of thousands of possible choices. Singular OR gene expression may optimize the detection of odorants by the ensemble of olfactory receptor neurons. The mechanism for how this singular OR expression occurs is unknown. Here we explore the possibility that singular OR expression results from the gene-gene interaction network having a winner-takes-it all solution. In one mechanism, the network selecting single gene is based on a large number of transcription factors (TFs) that mediate specific inhibitory interactions between genes. Instead, here we propose that the network of TF binding to OR genes is not specific and inhibition is provided by the competition between OR genes for the limited pool of TF. In this model, OR genes race between each other to reach the target number of bound TFs. A gene that recruits a target number of TFs is selected for expression. Our model can ensure stable single gene expression by OR cells if the binding of TFs by the OR promoters is cooperative which translates into an accelerated race between promoters. To support this model, we have analyzed the probability of OR choice represented by the levels of transcripts contained in previously published sequencing data. We correlated the probability of choice with the number of repeats of various TF binding motifs within the OR promoters. We find that the number of repeats of certain motifs within OR promoters is predictive of the levels of OR transcripts, suggesting the causal effect of promoter composition on the OR choice probability. Our model suggests that a small number of TFs can control the selection of a single gene out of ~2000 possibilities.

II-54. A feasible probe of the detailed microcircuit architecture of grid cells

John Widloski[1] WIDLOSKI@PHYSICS.UTEXAS.EDU
Ila Fiete[2] ILAFIETE@MAIL.CLM.UTEXAS.EDU

[1]University of Texas, Austin
[2]University of Texas at Austin

Populations of grid cells with similar period have been shown to exhibit key signatures of low-dimensional continuous attractor dynamics1,2. Unfortunately, existing data reveal little about the neural circuit responsible for this striking dynamical property. Structurally distinct grid cell models can produce low-dimensional dynamics: those based on 2D hexagonal patterning in the neural sheet, with periodic or aperiodic boundaries3,4; those with periodic boundaries in which the population pattern is a single bump5; those in which feedforward inputs exhibit 1D patterning and are simply summed6. To distinguish these models, we propose a simple novel probe of neural activity that 1) is based on simultaneously recording the spatial tuning of a handful of cells; 2) does not require topographic organization in the cortical sheet; 3) should provide direct evidence of recurrently driven 2D patterning within the neural sheet if it exists; 4) can reveal how many distinct bumps are in the population pattern, and thus 5) can help determine whether the network likely has periodic or aperiodic boundaries. The proposal is based on the theoretical observation that when the period of an underlying population pattern is slightly modified, the pairwise relationships in the spatial tuning of cells will shift in a way that is characteristic of the underlying pattern. The desired perturbations of the population period can be induced though changes in the temperature or global inhibition strength in the circuit. Feedforward vs. different recurrent architectures (single-bump vs. multi-bump; periodic

vs. aperiodic boundary) are predicted to behave differently under circuit perturbation. For aperiodic networks, the number of bumps in the pattern can be inferred from the distribution of relative phase shifts. We show that the proposed experiment can, at data volumes consistent with current technologies, reveal the circuit mechanisms of grid cells more directly and with much more detail than known at present.

II-55. Nonlinear coding and gain control in the population of auditory receptor neurons in Drosophila

Jan Clemens CLEMENSJAN@GOOGLEMAIL.COM
Cyrille Girardin GIRARDIN25@GMAIL.COM
Mala Murthy MMURTHY@PRINCETON.EDU

Princeton University

Successful acoustic communication relies on the ability to robustly recognize temporal patterns. Sound is used in Drosophila courtship during which females evaluate the temporal pattern of the male song to arbitrate mating decisions. The distance between male and female varies greatly during the courtship ritual leading to large fluctuation in sound intensity. Gain control mechanisms in the auditory system create intensity-invariant codes and thereby support robust pattern recognition. Previous studies looking at responses of antennal mechanosensory receptor neurons (Johnston's organ neurons — JONs) have shown that active amplification increases the fly's sensitivity to soft, sinusoidal sounds. Gain control is so far believed to be implemented using static compression. However, little is known about how complex temporal patterns with large amplitude fluctuations as found in natural song are represented in JONs. We recorded antennal movement, compound action potential responses (CAP) and Calcium responses (Gcamp6f) to artificial and natural stimuli and used computational models to describe how JONs encode temporal patterns with strong amplitude modulations. Specifically, we ask how the JON code supports robust song pattern recognition. We find that quadratic encoding models — but not linear models — reproduce CAP responses. The quadrature pair filters found in JONs are known from higher-order visual or auditory neurons in vertebrates, suggesting that JONs represent the instantaneous stimulus energy. We further find that JONs performs dynamic gain normalization over a wide range of behaviorally relevant intensities. Interestingly, adding gain control to the model is necessary for reproducing responses to courtship song — gain control is thus heavily engaged during coding song. Finally, information theory reveals that gain control facilitates pattern recognition over a large range of intensities. We have revealed novel, dynamical aspects of coding in JONs and show that intensity invariance is established at the input to the auditory system of Drosophila, likely facilitating song recognition during courtship.

II-56. Aligning auditory and motor representations of syllable onsets in songbird vocal learning

Emily Mackevicius ELM@MIT.EDU
Michale Fee FEE@MIT.EDU

Massachusetts Institute of Technology

A powerful strategy for learning a complex motor sequence is to split it into simple parts. The present study reveals a potential neural correlate for parsing complex sequences into manageable pieces during songbird vocal learning. We focus on a songbird brain region necessary for flexible sequencing of song syllables, nucleus interface (NIf) (Hosino & Okanoya, 2000). NIf is widely recognized as a node of interaction between the auditory and motor systems (Akutagawa & Konishi, 2010; Bauer et al., 2008; Cardin & Schmidt, 2004; Lewandowski et al. 2013). NIf plays a primary role during vocal development: lesions of NIf have minimal effect on adult zebra finch song, but inactivating NIf in juvenile birds causes loss of emerging song structure (Naie & Hahnloser, 2011). Furthermore, inactivating NIf while a juvenile bird is being tutored interferes with tutor imitation (Roberts et al.

2012). It was unknown what signals NIf conveys to the song motor system during vocal learning. Here we recorded NIf neurons in singing juvenile zebra finches. We also recorded during tutoring. Our recordings include antidromically identified single-units projecting to the premotor nucleus HVC. In all birds, NIf neurons exhibited premotor firing patterns, bursting prior to each syllable onset. These results provide a potential neural correlate for the sequence-initiator neurons that play a key role in existing models of HVC (Jun & Jin 2007, Bertram et al. 2014, Amador et al. 2013, Gibb et al. 2009). Notably, NIf neurons also fire at tutor syllable onsets, and some align just prior to tutor syllable onsets by firing in gaps between syllables. This activity during tutoring closely resembles NIf premotor activity during singing. By aligning to syllable onsets, NIf could "chunk" songs into syllables each of which is accessible as a discrete unit in both auditory and motor representations.

II-57. The dorsolateral striatum constrains the execution of motor habits through continuous integration of contextual and kinematic information

David Robbe[1,2] DAVID.ROBBE@INSERM.FR
Pavel Rueda-Orozco[1] PAVEL.RUEDA@GMAIL.COM

[1]Institut de Neurobiologie de la Méditerranée
[2]INSERM U901

The striatum is required for the acquisition of procedural memories but its contribution to motor control once learning has occurred is unclear. Here we created a task in which rats learned a difficult motor sequence characterized by fine-tuned changes in running speed adjusted to spatial and temporal constraints. Specifically, we customized a motorized treadmill and trained rats to obtain rewards according to a spatiotemporal rule. After an extensive training (1-3 months of daily sessions), rats succeeded in this task by performing a stereotyped motor sequence that could be divided in 3 overlapping phases: 1) passive displacement from the front to the rear portion of the treadmill, 2) stable running around the rear portion of the treadmill and 3) acceleration across the treadmill to enter the stop area. Tetrode recordings of spiking activity in the dorsolateral striatum (DLS) of well-trained animals revealed continuous integrative representations of running speed, position and time. These representations were weak in naive rats hand-guided to perform the same sequence and developed slowly after learning. Finally, DLS inactivation in well-trained animals preserved the structure of the sequence while increasing its trial-by-trial variability and impaired the animals capacity to make corrections after incorrect trial. We conclude that after learning the DLS continuously integrates task-relevant information to constrain the execution of motor habits. Our work provides a straightforward mechanism by which the basal ganglia may contribute to habit formation and motor control.

II-59. Behavioral and neural indices of the extraction of scale-free statistics from auditory stimuli

Brian Maniscalco[1,2] BMANISCALCO@GMAIL.COM
Biyu He[3,4] BIYU.JADE.HE@GMAIL.COM

[1]National Institute of Neurological Disorders
[2]National Institutes of Health
[3]NIH
[4]NINDS

Many features in both natural stimuli and brain activity exhibit power-law scaling relationships between power and temporal or spatial frequency ($P \propto 1/f\hat{}\beta$). Here, we investigate how well human observers can extract such statistical information from environmental stimuli, and the neural mechanisms underlying such processing. We presented subjects with 10-second long auditory sequences composed of 300-ms tones of varying pitches, and parametrically modulated the autocorrelation of tone pitch over time, mimicking autocorrelations in pitch fluctuation found in natural auditory stimuli such as music (Voss & Clarke, 1975). The final tone in each sequence was set to one of six possible pitches; depending on the preceding tones, the statistical likelihood of the final tone to have occurred in that sequence varied. Subjects judged the degree of pitch autocorrelation in each sequence and rated the probability of the final tone, given the preceding tone sequence. Crucially, subjects' ratings of final tone probability tracked changes in the expected value for final tone pitch governed by the statistical structure of the preceding tones, suggesting that subjects were able to extract the autocorrelation statistics in the sequence necessary for making valid predictions about upcoming tones. MEG data reveal that weaker sequence autocorrelation is associated with stronger tone ERFs, particularly in sensors located near auditory cortex, suggesting a possible neural signature of prediction error that higher-order brain regions might use to calibrate prediction of the upcoming stimulus. Furthermore, the ERF for the penultimate tone in the stimulus sequence is modulated by the expected value of the final tone pitch, suggesting a neural correlate of sequence prediction. The sensors modulated by sequence autocorrelation and final tone prediction exhibit substantial topographical overlap, suggesting that in this task similar neural sites over auditory cortices may participate in computations of prediction and prediction error.

II-60. Optogenetic disruption of posterior parietal cortex points to a role early in decision formation

Michael Ryan[1] MICHAEL.RYAN@UCSF.EDU
Anne Churchland[2] CHURCHLAND@CSHL.EDU
David Raposo[3] DRAPOSO@CSHL.EDU

[1]University of California, San Francisco
[2]Cold Spring Harbor Laboratory
[3]Champalimaud Neuroscience Institute

The posterior parietal cortex (PPC) is necessary for accurate decisions that integrate visual evidence over time. However, it is not clear when in these decisions PPC is required. Pharmacological inactivation points to an early, feedforward role for PPC, but electrophysiological responses are tuned over the entire 1000 ms decision duration. To define the window during which PPC is required, we disrupted PPC activity during decision formation in restricted time periods in 2 rats trained to make visual/auditory decisions about the event rate of a series of flashes/clicks. Activity was disrupted by unilaterally elevating firing rates of PPC neurons that expressed ChR2. Because we have previously shown that PPC neurons can increase or decrease in response to the decisions measured here, pan-neuronal elevation can potentially push the population into an unnatural state and disrupt behavior. To demonstrate that ChR2 elevation disrupts behavior, we confirmed that stimulation (40Hz, 5-10mW) presented throughout the 1000 ms trial increased decision uncertainty. An increase in uncertainty for visual decisions on stimulation trials led to worse performance on those trials. This was individually significant in 6/7 sites. We observed a smaller change for auditory decisions (individually significant in 3/7 sites). Next, we examined the effect on visual decision uncertainty when stimulation was presented during restricted 250 ms intervals spanning the trial. We observed the largest effects when stimulation took place from 0-250 ms of decision formation. There, the magnitude of the effect was comparable to the full 1000 ms stimulation. Effects were weaker when stimulation was presented mid-trial, and no effect was evident when we stimulated at decision end, from 750-1000ms. Our results suggest that PPC activity is influential early in decision formation. Our results argue that visual signals are integrated elsewhere, and that the sustained electrophysiological responses observed previously reflect feedback from other areas.

II-61. Cortical population nonlinearities reflect asymmetric auditory perception in mice

Thomas Deneux[1] THOMAS.DENEUX@UNIC.CNRS-GIF.FR
Louise Francois[2] FRANCOIS.LOUISE@GMAIL.COM
Suncana Sikiric SUNCANASIKIRIC@GMAIL.COM
Emmanuel Ponsot EMMANUEL.PONSOT@IRCAM.FR
Brice Bathellier[3] BRICE.BATHELLIER@UNIC.CNRS-GIF.FR

[1]UNIC, CNRS UPR 3293
[2]Etudiant
[3]UNIC - CNRS

Natural sounds display strong temporal variations of intensity, which are part of the features influencing perception and recognition of sounds. Strikingly, for example, humans perceive a tone that increases in intensity as louder than same tone with a decreasing intensity profile, although both sounds have the same energy and frequency content. The underlying neuronal mechanisms of this perceptual asymmetry are still elusive. To test if the direction of intensity variations is asymmetrically processed by the auditory system, we have measured the activity of large populations of neurons in the auditory cortex of awake mice using GCAMP6 two-photon calcium imaging. Pooling a large number of recordings together, we observed that the time integral of cortical population firing rate is much larger for sounds ramping-up than for sounds ramping-down. This asymmetry demonstrates that cortical population response is strongly non-linear. To test for perceptual consequences of this non-linearity, we performed behavioral experiments in which the saliency of a sound is measured through associative learning speed. We observed that increasing ramps are more rapidly associated to a correct behavior than decreasing ramps, showing that the asymmetry of cortical population responses reflects an asymmetry in perceived saliency. Moreover, finer analysis of cortical data indicate that ramps produce complex population activity sequences in which distinct patterns emerge depending on the direction of sound intensity variations. Based on simple population models, we show that the asymmetry of population firing rate and the complexity of activity sequences could be explained by competing populations of neurons encoding different aspects of the temporal profile (e.g. sound onset, offset). Beyond proposing a mechanism for a perceptual asymmetry that may emphasize approaching sound sources, our results suggest that non-linear interactions between temporal feature detectors could be one of the bases of sound recognition.

II-62. How did the evolution of color vision impact V1 functional architecture?

Manuel Schottdorf[1] MANUEL@NLD.DS.MPG.DE
Wolfgang Keil[2] WKEIL@MAIL.ROCKEFELLER.EDU
Juan Daniel Florez Weidinger[1] CHEPE@NLD.DS.MPG.DE
David M Coppola[3] DCOPPOLA10@VERIZON.NET
Amiram Grinvald[4] AMIRAM.GRINVALD@WEIZMANN.AC.IL
Koji Ikezoe[5] IKEZOE@FBS.OSAKA-U.AC.JP
Zoltan F Kisvarday[6] KISVARDAY@ANAT.MED.UNIDEB.HU
Tsuyoshi Okamoto[7] OKAMOTO@DIGITAL.MED.KYUSHU-U.AC.JP
David B Omer[4] DAVID.OMER@WEIZMANN.AC.IL
Leonard White[8] LEN.WHITE@DUKE.EDU
Fred Wolf[1] FRED@NLD.DS.MPG.DE

[1]MPI for Dynamics and Self-Organization
[2]The Rockefeller University
[3]Randolph-Macon College
[4]Weizmann Institute of Science
[5]Osaka University
[6]University of Debrecen

[7]Kyushu University
[8]Duke University

Color vision was lost in mammals during the nocturnal bottleneck when our ancestors were small, dark-dwelling animals between 205 to 65 Million years ago (Ma). Among modern mammals only old world monkeys and great apes (re-)invented trichromacy 30-40Ma. The newly developed color vision inserted new pathways into cortical functional architecture, most obviously potentially perturbing the orientation domains through non-orientation selective CO-Blobs. How much impact color vision had on the overall functional visual cortical architecture remains unclear. Here, we investigate this question focusing on orientation domains, a key characteristic of V1 functional architecture allowing quantitative analysis. Orientation domains are arranged around pinwheel singularities, whose spatial distribution in ferrets, shrews, and galagos is quantitatively indistinguishable. At least for dichromats, there exists a common design, characterized by the statistical identity of (i)pinwheel density, (ii)pinwheel density fluctuations as a function of subregion size, and (iii)nearest neighbor distance distributions. Against a background of normal(N=82) and dark-reared(N=21) ferret, shrew(N=25), galago(N=9), and cat(N=13) we compared macaque(N=9) OPMs exhibiting 1183 pinwheels and found that their layout adheres to the common design. Most notably, the pinwheel density ρ=3.19[3.04,3.39], is extremely close to the mathematical constant π verifying the prediction of a universal solution set of a large symmetry defined class of self-organization models. This class also predicts the measured exponent of pinwheel density fluctuations γ=0.40[0.33,0.43] and the mean distance d=0.35[0.33,0.37]. Our quantitative results indicate that the evolutionary invention of the color vision machinery induced only a minor perturbation of the system of orientation domains. The selective forces that favor the common design might thus be so powerful as to preserve it under major transformations of the retinocortical pathway.

II-63. Distally connected memory-like network gates proximal synfire chain

Jacopo Bono J.BONO13@IMPERIAL.AC.UK
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK

Imperial College London

Recent experimental results (Branco and Hausser, 2011) showed that local synaptic integration on a thin dendritic branch is non-linear, and the way excitatory post-synaptic potentials (EPSPs) are integrated temporally depends on the synaptic location on a dendrite. Using computational modelling, we investigated the functional implication of different synaptic integration in the development of connectivity structure formed by synaptic plasticity. Input arriving at distal synapses on a dendritic branch, where temporal integration of EPSPs is efficient, is more suitable for rate codes and easily forms strong bidirectional connections through synaptic plasticity. On the other hand, proximal synapses on the same branch, where no efficient temporal integration is present, are better suited for synchrony detection and temporal codes, and can more easily form strong unidirectional connections. These results suggest that differences in synaptic locations can lead to different functions and that one neuron could use different codes simultaneously. In particular, a memory-like bidirectional network distally connected to a synfire chain of proximally unidirectionally connected neurons, can serve as a gate for information transfer in this chain. This gate in single neurons can be a computational mechanism for the integration of top-down and bottom-up inputs.

II-64. Intracellular, in vivo, characterization and control of thalamocortical synapses in cat V1

Madineh Sedigh-Sarvestani MADINEH@UPENN.EDU
M Morgan Taylor TAYM@MAIL.MED.UPENN.EDU
Larry A Palmer PALMERL@MAIL.MED.UPENN.EDU
Diego Contreras DIEGOC@MAIL.MED.UPENN.EDU

University of Pennsylvania

Neurons in input layers of primary visual cortex (V1) exhibit response properties that are absent one synapse away in the lateral geniculate nucleus (LGN). How these properties arise in V1 is a fundamental problem in neuroscience. It is estimated that 30 LGN cells synapse onto a cortical cell, thus input integration dynamics is important in understanding the contribution of thalamocortical (TC) input. Furthermore, TC synapses make up only 10% of simple cell synapses, so understanding the functional contribution of the TC drive is critical. We use two novel techniques to study TC synapses and to isolate the contribution of the TC drive to response properties of L4 simple cells. Using paired recordings of monosynaptically connected LGN (extracellular) and V1 (intracellular) neurons in cats in vivo, we characterize the TC synapse, quantify visually evoked cortical EPSPs, identify rules of connectivity, and quantify input integration. We find that LGN action potentials lead to cortical EPSPs with an average efficacy of 42 $\pm$ 16%. Monosynaptic EPSPs have an average amplitude of 0.85 $\pm$ 0.18 mV. We find clear instances of short-term synaptic depression and facilitation at the TC synapse. Using the dynamic clamp technique, we simulate monosynaptic TC synapses between many LGN neurons and one V1 neuron. We find that using as few as six LGN cells, we can create simple cell like receptive fields in L4 complex cells, exhibiting elongated ON and OFF subregions and an F1/F0 ratio greater than 1. We are using this paradigm to investigate the contribution of TC population dynamics to functional properties of L4 cells. Given the high resolution of its visual system, the cat is an ideal system for the study of thalamocortical visual encoding. Our studies in the cat are needed for contextualizing future studies in the mouse and primate.

II-65. Perceptual adaptation: Getting ready for the future

Xue-Xin Wei WEIXXPKU@GMAIL.COM
Pedro Ortega PEDRO.ORTEGA@GMAIL.COM
Alan A Stocker ASTOCKER@SAS.UPENN.EDU

University of Pennsylvania

Perceptual systems continually adapt to changes in their sensory environment. Adaptation has been mainly thought of as a mechanism to exploit the spatiotemporal regularities of the sensory input in order to efficiently represent sensory information. Thus, most computational explanations for adaptation can be conceptualized as a form of Efficient coding. We propose a novel and more holistic explanation. We argue that perceptual adaptation is a process with which the perceptual system adjusts its operational regime to be best possible prepared for the future, i.e. the next sensory input. Crucially, we assume that these adjustments affect both the way the system represents sensory information (encoding) and how it interprets that information (decoding). We apply this idea in the context of a Bayesian observer model. More specifically, we propose that the perceptual system tries to predict the probability distribution from which the next sensory input is drawn. It does so by exploiting the fact that the recent stimulus history is generally a good predictor of the future and that the overall long-term stimulus distribution is stationary. We assume that this predicted probability distribution reflects the updated prior belief of the Bayesian observer. In addition, we assume that the system is adjusting its sensory representation according to the predicted future stimulus distribution via Efficient coding. Because this sensory representation directly constrains the likelihood function, we can define an optimal Bayesian observer model for any predicted distribution over the next sensory input. We demonstrate that this model framework provides a natural account of the reported adaptation after-effects for visual orientation and spatial frequency, both in terms of discrimination thresholds and biases. It also allows us to predict how these after-effects depend on the specific form of the short-

and long-term input histories.

II-66. Linearity in visual cortex: correlations can control neurons' sensitivity to coincident input

Mark Histed[1] HISTED@UCHICAGO.EDU
Oliver J Mithoefer[2] OLIVER.MITHOEFER@GMAIL.COM
Lindsey Glickfeld[3] LLGLICK@GMAIL.COM
Alex Handler[2] THEALEXHANDLER@GMAIL.COM
John Maunsell[1] MAUNSELL@UCHICAGO.EDU

[1]University of Chicago
[2]Harvard University
[3]Duke University

Brain computations occur when neurons transform their complex barrage of synaptic inputs into output spiking. We studied input-output transformations in mouse visual cortex by optogenetically stimulating excitatory or inhibitory cells to change neurons' spontaneous rate, while presenting the same visual input repeatedly and measuring spiking responses. We find that neurons' responses to visual stimuli are nearly independent of spontaneous firing, a hallmark of a linear input-output relationship. Using a spiking model, we show that the statistical structure of background cortical input can shape neurons' input-output relationships. Supralinearity, which confers sensitivity to input timing, occurs when background inputs are independent, or when excitatory inputs are correlated. Correlation between inhibitory inputs, however, allows model neurons to produce a linear response that accurately describes the experimental observations. The strength of coupling between inhibitory neurons may control cortical computations by regulating cortical neurons' sensitivity to input timing.

II-67. Characterizing local invariances in the ascending ferret auditory system

Alexander Dimitrov[1] ALEX.DIMITROV@VANCOUVER.WSU.EDU
Stephen V David[2] DAVIDS@OHSU.EDU
Jean Lienard[1] JEAN.LIENARD@GMAIL.COM

[1]Washington State University Vancouver
[2]Oregon Health Sciences University

The sense of hearing requires a balance between competing processes of perceiving and ignoring. Behavioral meaning depends on the combined values of some sound features but remains invariant to others. The invariance of perception to physical transformations of sound can be attributed in some cases to local, hard-wired circuits in peripheral brain areas. However, at a higher level this process is dynamic and continuously adapting to new contexts throughout life. Thus the rules defining invariant features can change. In this project, we test the idea that high-level, coherent auditory processing is achieved through hierarchical bottom-up combinations of neural elements that are only locally invariant. The main questions we address in the context of an auditory system are: 1. What kinds of changes in sound do not affect initial stages of auditory processing? 2. How does the brain manipulate these small effects to achieve a coherent percept of sounds? Local probabilistic invariances, defined by the distribution of transformations that can be applied to a sensory stimulus without affecting the corresponding neural response, are largely unstudied in auditory cortex. We assess these invariances at two stages of the auditory hierarchy using single neuron recordings from the primary auditory cortex (A1) and the secondary auditory cortex (PEG) of awake, passively listening ferrets. Our results show that stimulus invariance to frequency and time dilations are present at every tested stage and increase along the hierarchical auditory processing. At least in the early stages, parametric models having invariance properties by design are well-suited to describing biological functions. We were further able to characterize meaningful relationships among

receptive field shapes. Preliminary observations indicate that joint time/frequency receptive fields are oriented toward central frequencies; receptive field widths are proportional to the best frequency; and late-onset neurons are also exhibiting the most sustained activity.

II-68. Low-rank decomposition of variability in neural population activity from songbird auditory cortex

Lars Buesing LARS@STAT.COLUMBIA.EDU
Ana Calabrese AMC2257@COLUMBIA.EDU
John Cunningham JPCUNNI@GMAIL.COM
Sarah Woolley SW2277@COLUMBIA.EDU
Liam Paninski LIAM@STAT.COLUMBIA.EDU

Columbia University

Simultaneously recorded spike trains of multiple neurons from a local population exhibit shared variability, often characterized by so-called noise-correlations. Previous studies have shown that the magnitude of such shared variability depends on (amongst others) external stimuli as well as the distance between neurons. Here we analyze shared variability in multi-electrode recordings from the primary auditory cortex of awake songbirds under auditory stimulation with conspecific song stimuli and artificial stimuli as well as under no stimulation. We find that shared variability is strongly modulated by stimulus class as well as recording location within the auditory cortex, on a sub-millimeter scale. In spite of these seemingly complex dependencies on covariates, shared variability in this data has a surprisingly simple structure. We show here that neural variability under auditory stimulation with different stimulus classes can be captured by a single, probabilistic factor model. Using this model, we show that neural variability is accurately captured by one or two latent factors, yielding a precise, low-rank description of noise-correlations. Furthermore, the factor loadings identified from the data exhibit a clear co-variation with recording position allowing the model to account for the location-dependence of variability. As a technical contribution, we provide methods for fitting latent variable models with Poisson observations and non-canonical link function.

II-69. Constraining the mechanisms of direction selectivity in a fruit fly elementary motion detector

Jonathan Leong LEONGJCS@STANFORD.EDU
Benjamin Poole POOLE@CS.STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Thomas Clandinin TRC@STANFORD.EDU

Stanford University

Elementary motion detectors (EMDs) must combine information from two or more points in space-time to transform local, non-direction selective inputs into direction selective output. Still, it remains unknown how EMDs of the fruit fly visual system integrate spatiotemporal input to become direction selective. We used in vivo two-photon calcium imaging to monitor visually evoked activity in the cell type T4, a previously identified fruit fly EMD, while displaying motion stimuli and spatiotemporal noise. While activity in T4 dendrites exhibited little directional tuning of their mean responses to drifting gratings, “1F” tuning of the amplitude of the phase-locked response exhibited a salt-and-pepper functional organization characterized by scattered puncta, each having a strong and uniform direction selectivity. 1F tuning is a hallmark of linear mechanisms of direction selectivity, and is not predicted by a nonlinear mechanism such as the Hassenstein-Reichardt correlator. Estimated from the responses of direction selective puncta to spatiotemporal noise, linear receptive fields (RFs) were bilobed, exhibiting excitatory and inhibitory subfields offset in both space and time. Excitation was faster than inhibition. The spatiotemporal orga-

nization of the linear RF could predict direction preference, as well as the contrast frequency tuning and spatial frequency tuning of T4 axons. A two-layer model circuit predicted our experimental observations. The first layer consisted of two neurons constructed from the known spatial filters, temporal filters, and nonlinearities of cell types Tm3 and Mi1, the main inputs to T4. Subtraction in the second layer reproduced the observed spatiotemporal organization of the linear RF, and also the contrast frequency tuning and spatial frequency tuning of T4. Taken together, these functional observations serve to further constrain the mechanisms of direction selectivity in a fruit fly EMD.

II-70. A linear-non-linear model of heat responses in larval zebrafish

Martin Haesemeyer[1] HAESEMEYER@FAS.HARVARD.EDU
Drew Robson[2] ROBSON@ROWLAND.HARVARD.EDU
Jennifer Li[1] JLI@ROWLAND.HARVARD.EDU
Alexander Schier[1] SCHIER@FAS.HARVARD.EDU
Florian Engert[1] FLORIAN@MCB.HARVARD.EDU

[1]Harvard University
[2]Rowland Institute at Harvard

Studying how animals transform sensory input into behavioral actions yields important insight into how nervous systems process information. To avoid noxious heat larval zebrafish robustly change their swim kinematics in response to changes in water temperature. How larval zebrafish integrate information about temperature in order to select appropriate behavioral actions is largely unknown. Here we set out to map the temporal receptive field of heat responses in larval zebrafish. To this end we extend reverse correlation approaches that have been widely used to map receptive fields of neurons, to the organism level relating sensory input to behavioral output. We use a custom-built setup to heat freely swimming zebrafish with high spatio-temporal precision using an infrared laser. This allows us to probe swim-bout initiation in response to white noise heat stimuli and derive a LNP model describing a transfer function from sensory input to behavioral output. We found that larval zebrafish mainly integrate temperature information over a time window of about 600 ms before bout initiation. The associated non-linearity revealed that the swim probability increases with increasing temperature. Furthermore subdividing bouts according to speed or turn magnitude allowed us to identify how sensory information directs not only swim initiation but also the type of bout a fish will likely perform. We could validate the obtained models by accurately predicting responses of larval zebrafish to white noise and filter playback stimuli. In addition our model can partly predict avoidance of noxious heat observed in freely behaving zebrafish. Taken together we were able to derive a linear-nonlinear model that transforms sensory input, in this case heat, to behavioral output, in our case bout initiation. The results of this study put important constraints on how neuronal circuits may integrate temperature information in order to generate appropriate behavioral output.

II-71. Modular neural circuit architecture for optic flow processing in the larval zebrafish

Eva Naumann[1,2] EVA.AIMABLE@GMAIL.COM
Timothy Dunn[2] DUNN.TW@GMAIL.COM
Jason Rihel J.RIHEL@UCL.AC.UK
Florian Engert[2] FLORIAN@MCB.HARVARD.EDU

[1]University College London
[2]Harvard University

In order to understand the necessary computations of neural circuits underlying visual motion processing, we examined directed turning behavior in larval zebrafish responding to whole-field motion. High-speed behavioral

analysis revealed several principles of information integration across eyes as well as general visuo-motor transformations that are used by fish for appropriately directed locomotion. These findings predict multiple, separate and overlapping neural circuit modules fed by distinct and specific information channels for visual motion emerging from each eye. These channels, revealed via presentation of monocular stimuli in a closed-loop assay, integrate binocular information linearly but modulate swim bout frequency and orientation change separately. Importantly, they also serve to stabilize orienting behavior in ambiguous visual environments. In-vivo two-photon imaging of neural activity at single-cell resolution throughout the brain of transgenic zebrafish expressing a genetically encoded calcium indicator (GCaMP5) provides a descriptive overview of the underlying neural circuit. Within this circuit, lateralized midbrain nuclei with sharp directional tuning integrate motion binocularly and enable necessary reciprocal suppression; the resulting locomotor instructions are then distributed to specific premotor areas that modulate turn and bout frequency separately. These findings, combined with a cluster analysis across tens of thousands of neurons, have motivated a feed-forward rate code model composed of modular circuit elements that captures the behavioral output accurately. Furthermore, evaluation of minimal models based on measured neuron response profiles has provided testable predictions for further characterizing the functional brain architecture underlying this behavior.

II-72. Non-linear stimulus integration for computing olfactory valence

Joe Bell JOE_BELL@HMS.HARVARD.EDU
Rachel Wilson RACHEL_WILSON@HMS.HARVARD.EDU
Harvard University

All animals must use patterns of olfactory receptor neuron (ORN) activity to compute appropriate behavioral responses to odors, but the form of this computation is unknown. In both flies and mammals, each ORN expresses one type of odorant receptor that determines its odor response profile, and all ORNs expressing the same receptor project to the same compartment in the brain, or glomerulus. These glomeruli constitute parallel processing channels that relay olfactory information to other areas, and are often co-activated by odors. One popular hypothesis holds that activity in each processing channel contributes to a central representation of valence, or pleasantness, by a fixed linear weight. Alternatively, non-linear interactions between specific channels could confer increased selectivity or sensitivity for ethologically important odors. We investigate these alternative models by optogenetically activating olfactory channels in freely walking Drosophila, and use video tracking to analyze the behavioral responses of thousands of individual flies. Some channels produce robust attraction when activated individually. Combining stimulation of pairs of channels produces unpredictable results: some pairs sum to produce attraction greater than that elicited by either component, but others produce the same behavior as the more attractive component alone. Surprisingly, we find no reliably repulsive channels, but some potently reduce attraction in combinations. Based on which channels summate, we develop a simple model that establishes a lower bound on the dimensionality of internal olfactory representations. Although we provide evidence of several distinct pools of channels, detailed analysis of walking trajectories suggests that flies respond to these stimuli by graded recruitment of a single behavioral program. These data provide a quantitative and systematic description of how a sensory system integrates signals from many sources to generate a behavioral output, and provide a model example of how psychophysics can suggest constraints on the functional architecture of sensory circuits.

II-73. Optogenetic modulation of transient and sustained response to tones in auditory cortex of awake mice

Pedro Goncalves[1] PEDROG@GATSBY.UCL.AC.UK
Jannis Hildebrandt[2] JANNIS.HILDEBRANDT@UNI-OLDENBURG.DE
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK
Jennifer Linden[3] J.LINDEN@UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University of Oldenburg
[3]University College London

The roles of cortical interneurons in sensory processing have been the subject of many recent optogenetic experiments, typically involving brief and temporally precise manipulations of activity in targeted classes of interneurons. Such experiments have helped to unveil the instantaneous impact of the targeted interneurons on cortical processing of sensory inputs. However, much less is known about the roles of cortical interneurons in shaping the slower dynamics of cortical spontaneous activity and its interaction with sensory input. Understanding the impact of prolonged changes in interneuron activity on cortical dynamics may provide key insight into how sensory processing is altered by disease and by behavioural phenomena such as attention or learning. Here, we studied the involvement of parvalbumin-positive (PV+) interneurons in modulating neural activity in the auditory cortex of awake mice. Continuous low-level activation of PV+ interneurons with stabilised step-function opsin (SSFO) produced a decrease in LFP power in the high-gamma range during spontaneous activity and during the sustained response to 500 ms tones. However, the prolonged SSFO activation of PV+ interneurons increased transient LFP responses to the onsets and offsets of tones without substantially affecting overall tuning of tone-evoked activity. These results are consistent with the hypothesis that auditory cortical dynamics are shaped by synaptic depression: a decrease in spontaneous activity releases intracortical synapses from depression, increasing transient cortical responses to sensory inputs. In further agreement with this hypothesis, we observed that tone-evoked LFP responses are negatively correlated with preceding spontaneous LFP power in the high-gamma range. Overall, these results suggest that PV+ interneuron activity has a profound impact on auditory cortical dynamics, modulating the relative strength of transient versus sustained responses to sensory stimuli.

II-74. Future actions determine the currently active motor memory during skill learning

Daniel Wolpert[1] WOLPERT@ENG.CAM.AC.UK
Ian Howard[2] IAN.HOWARD@PLYMOUTH.AC.UK
David Franklin[1] DWF25@CAM.AC.UK

[1]University of Cambridge
[2]University of Plymouth

To learn a motor skill over a prolonged period of time the motor memory of the skill must be stored, protected from interference by intervening tasks, and reactivated for modification when the skill is practised. In many ball sports we are taught to follow-through consistently, despite the inability of events after contact or release to influence the outcome. Here we examine the potential role that such future movements may have on the activation of current motor memories and show that the specific motor memory active at any given moment critically depends on the movement that will be made in the near future. We examined a motor skill that is known to be long-lasting but also subject to interference—learning to reach in the presence of a dynamic (force-field) perturbation generated on the hand by a robotic interface When two opposing force-fields are presented alternately, there is substantial interference preventing learning of either. We first show that linking such skills that normally interfere to different follow-through movements activates separate motor memories for each, thereby allowing both skills to be learned without interference. This implies that when learning a skill, a variable follow-through would activate multiple motor memories across practice, whereas a consistent follow-through would activate a single motor memory, resulting in

faster learning. We confirm this prediction and using a dual-rate state-space model show that the variable follow through leads to lower retention of the motor skill from trial to trial, consistent with activating multiple memories. Finally, we show that such follow-through effects influence adaptation over time periods (days) associated with real-world skill learning. Our results suggests that there is a critical period both before and after the current movement that determines motor memory activation and controls learning.

II-75. Isolating short-time directional stability in motor cortex during center out reaching

Josue Orellana[1] JOSUE@CMU.EDU
Steven Suway[2] SBS45@PITT.EDU
Robert Kass[1] KASS@STAT.CMU.EDU
Andrew Schwartz[2] ABS21@PITT.EDU

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

Activity in motor cortex is strongly modulated by reach direction. Typically, a large time window of peri-movement activity is selected to estimate the preferred direction (PD) under a cosine-tuning model. Yet multiphasic activity is common during straight center-out reaches, suggesting the classical model may be incomplete. Here we use a novel method to separate this multiphasic activity into components that exhibit directional stability for periods on the order of 200-500 milliseconds. Using multi-electrode activity recorded from M1 and PMv from two rhesus monkeys performing hand-controlled center out reaching in 2/3D space, we applied a bootstrap method to assess PD stability. Only 36% and 24% of the cells, respectively, for monkey C (80 units) and F (119 units) were found to have stable PDs across the entire movement task. However, 90% (monkey C) and 91% (monkey F) of the cells had at least one stably tuned period over a shorter time window. We observed up to three separable components per cell and found 64% and 75% of all isolated components from all cells to be directionally stable. These results suggest that directional tuning may be stable at short time scales.

II-76. Context-dependent information decoding of sensory-evoked responses

He Zheng[1] NEURON@GATECH.EDU
Biyu He[2,3] BIYU.JADE.HE@GMAIL.COM
Garrett Stanley[1] GARRETT.STANLEY@BME.GATECH.EDU

[1]Georgia Institute of Technology
[2]NIH
[3]NINDS

The cortical response to an external sensory stimulus is embedded in the spontaneous activity that is constantly ongoing and dynamic, resulting in a highly variable representation of the sensory input that yet can produce relatively consistent stimulus perception. To characterize the sensory-evoked response, the traditional approach has been to model the variability as the linear superposition of a stereotypical, trial-averaged response and the constantly changing spontaneous activity. However, without accounting for the potential interaction between the ongoing spontaneous activity and the sensory-evoked activity on single-trial basis, the trial-averaged response cannot provide a complete picture of the dynamic processes in the brain. We hypothesize that given access to the spontaneous activity, the ideal observer can reduce uncertainty about the stimulus. To model the potential interaction between the spontaneous activity and sensory-evoked response, we recorded spontaneous and sensory-evoked brain activity with simultaneous local field potential (LFP) and genetically-encoded voltage-sensitive fluorescent protein imaging (ArcLight) in the same cortical column of the primary somatosensory cortex in the anesthetized rat. We employed an auto-regressive model to predict spontaneous activity after stimulus on-

set and found that, under a linear summation model, the spontaneous activity (denoted O) and stimulus-evoked component (denoted E) that sum up to the observed peak activity (denoted P, where P=O+E) are anti-correlated ($r(O,E)<0$). In the non-adapted state, the variance of the stimulus-evoked component is not significantly different from that of the observed peak activity. Thus, the knowledge of spontaneous activity does not reduce uncertainty about the stimulus. However, following sensory adaptation, the relationship between spontaneous and evoked components becomes less anti-correlated, resulting in a reduction of variance in the evoked component and a reduction of uncertainty about the stimulus. Modulating the energy of the adapting stimulus by changing adapting stimulus frequency and/or velocity is found to induce a gradient of this effect.

II-77. Why do receptive field models work in recurrent networks?

Johnatan Aljadeff[1,2] ALJADEFF@UCSD.EDU
David Renfrew[3] DTRENFREW@MATH.UCLA.EDU
Tatyana Sharpee[2,1] SHARPEE@SALK.EDU

[1]University of California, San Diego
[2]Salk Institute for Biological Studies
[3]University of California, Los Angeles

Advances in experimental techniques emphasize the rich connectivity structure of biological neural networks. However, only in a small number of well-studied circuits the recurrent nature of the neural substrate can be directly related to the “receptive-field” view of the network’s activity. To address these issues in a general network model, we studied the dynamics of random recurrent networks with a synapse-specific gain function. These networks become spontaneously active at a critical point that is derived here, directly related to the boundary of the spectrum of a random matrix model that has not been studied previously. Given the gain function we predict analytically the network’s leading principal components in the space of individual neurons’ autocorrelation functions, thereby providing a direct link between the network’s structure and some of its functional characteristics. In the context of analysis of single and multi-unit recordings our results (a) offer a mechanism for relating the recurrent connectivity in sensory areas to feed-forward receptive field models; and (b) suggest a natural reduced space where the system’s trajectories can be fit by a simple state-space model.

II-78. Using expander codes to construct Hopfield networks with exponential capacity

Rishidev Chaudhuri RCHAUDHURI@AUSTIN.UTEXAS.EDU
Ila Fiete ILAFIETE@MAIL.CLM.UTEXAS.EDU

University of Texas at Austin

Noise is ubiquitous in the brain, limiting how networks of neurons represent and propagate information. Consequently, the brain must encode signals redundantly and recover them from degraded input. Current models yield either weak increases in representational capacity with network size or exhibit poor robustness to noise. The grid cell code has exponential capacity and can in principle be robust to noise, but only with an appropriate decoder whose complexity is not yet well-characterized. Here we show that undirected graphs with Hopfield dynamics can store exponentially many states, and moreover, that errors in a finite fraction of nodes will be corrected by the dynamics, so the information rate of the code is finite. By contrast, the best such existing codes have rates decreasing as 1/log(N). The present advance is achieved by mapping the decoding constraints of expander codes, a subclass of low-density parity-check codes, onto the energy function of a higher-order Hopfield network (in which edges connect more than two nodes). In general, the constraints of error-correcting codes cannot be decoded through the relaxation dynamics of a Hopfield-like energy function. Expander codes are sparse bipartite graphs where relatively few nodes co-occur in multiple constraints. This allows for simple, local decoding and in particular,

implementation in Hopfield networks. The higher-order variable dependencies of higher-order Hopfield networks can be converted into pairwise-only weights by adding hidden units. This suggests an equivalent construction of our network, where each higher-order edge is replaced by a sparsely-connected Boltzmann machine imposing the appropriate constraints; non-constraint nodes are the representational units. Importantly, the sparsity of expander code constraints means that hidden units occupy a constant fraction of the network. Our results demonstrate that the structures of certain important error-correcting codes, wherein sparse constraints produce high-dimensional systems with large capacity and robustness, might be applicable to neural architectures.

II-79. Encoding and discrimination of multi-dimensional signals using chaotic spiking activity

Guillaume Lajoie[1,2] GLAJOIE@NLD.DS.MPG.DE
Kevin K. Lin[3] KLIN@MATH.ARIZONA.EDU
Jean-Philippe Thivierge[4] JEAN-PHILIPPE.THIVIERGE@UOTTAWA.CA
Eric Shea-Brown[5] ETSB@AMATH.WASHINGTON.EDU

[1]MPI for Dynamics and Self-Organization
[2]Bernstein Center for Computational Neuroscience
[3]University of Arizona
[4]University of Ottawa
[5]University of Washington

Networks of neurons throughout the brain show highly recurrent connectivity. While this recurrence presumably contributes to information coding and processing, it also has a seemingly contrary effect: across many models, recurrent connectivity leads to chaotic dynamics [Sompolinsky et al., 1988, van Vreeswijk and Sompolinsky, 1998, Monteforte and Wolf, 2010] This implies that the spike times elicited by a given stimulus can and do depend sensitively on initial conditions, and suggests that precise temporal patterns of spikes are too fragile and unreliable to carry any useful information [van Vreeswijk and Sompolinsky, 1998, London et al., 2010]. However, recent work has shown that despite being chaotic, the variability of network responses to temporal inputs can be surprisingly low, a fact attributed to low-dimensional chaotic attractors [Lin et al., 2009, Lajoie et al., 2013, Lajoie et al., 2014] and compatible with experimental observations of spike-time repeatability in recurrent networks [Reinagel and Reid, 2000]. Here, we investigate the implications of low-dimensional chaos on the ability of large recurrent neural networks to encode and discriminate time-dependent input signals, focusing on networks operating in a fluctuation-driven balanced state regime [van Vreeswijk and Sompolinsky, 1998]. We find that one can easily discriminate between network responses to temporal inputs with up to 95% correlations, despite chaotic activity. This discrimination can be achieved either based on the “sub-threshold” activity (voltage traces) or based on spike outputs, including using a trained tempotron [Gutig and Sompolinsky, 2006] classifier on output spike patterns. Furthermore, this discrimination capacity persists even when observing only a few neurons in a network that do not themselves directly receive discriminable signals. Thus, recurrent connections distribute signals throughout networks in ways that can enhance the classification power of the network, despite the chaotic behaviour that often results from recurrent connections.

II-80. Neural oscillations as a signature of efficient coding

Matthew Chalk MATTHEWJCHALK@GMAIL.COM
Boris Gutkin BORIS.GUTKIN@ENS.FR
Sophie Deneve SOPHIE.DENEVE@ENS.FR

Ecole Normale Superieure

Cortical networks typically exhibit ‘global oscillations’, in which neural spike times are entrained to a global oscil-

latory rhythm, but where individual neurons fire irregularly, on only a small fraction of oscillation cycles. While the network dynamics underlying global oscillations have been well characterised, their computational role is debated. We show that global oscillation are predicted as a consequence of efficient coding in a recurrent spiking network. To avoid firing unnecessary spikes (i.e. that communicate redundant information), neurons must share information about the global state of the network. Previously, we have shown that maximally efficient coding is achieved in a balanced network, where membrane potentials encode a global prediction error. After each spike, recurrent inhibition ensures that all neurons maintain a consistent representation of the error, and do not fire redundant spikes. As a result, spikes are aligned to global fluctuations in prediction error, while single neuron responses are irregular and sparse. In a network with realistic synaptic delays, inhibition does not always arrive fast enough to prevent multiple neurons firing together. This results in rhythmic fluctuations in the prediction error and population activity. Oscillations are mediated by balanced fluctuations in excitation and inhibition, with excitation leading inhibition by several milliseconds. To investigate the impact of oscillations on coding, we varied the intrinsic noise, spike threshold and connection strengths, to alter the network synchrony while keeping firing rates constant. Regardless of the manipulation we performed, coding performance was maximised in the asynchronous-rhythmic regime where: (i) neural spike times are entrained to global oscillations (implying a consistent representation of the error); (ii) few neurons fire on each cycle (implying high efficiency), and (iii) excitation and inhibition are tightly correlated. In contrast, the network performed sub-optimally when the population activity was arhythmic, or when many neurons fired on each cycle.

II-81. Adaptation and homeostasis in a spiking predictive coding network

Gabrielle Gutierrez GABRIELLE.GUTIERREZ@ENS.FR
Sophie Deneve SOPHIE.DENEVE@ENS.FR

Ecole Normale Superieure

We study a balanced, spiking network of LIF neurons in which neurons fire spikes to reduce the error of the decoded signal. Here, metabolic costs on spikes ensure an efficiently coded representation of the input signal. Using a network of leaky integrate and fire neurons, we have found that imposing metabolic costs that both reduce overall spiking and distribute spikes among network neurons leads to tuning curves that adapt to the statistics of the input (Fig. 1). Importantly, decoding performance is not impeded by this adaptation. The present study reproduces a phenomenon that has been observed in biological experiments, such as in visual cortex, where neuronal tuning curves are shaped by the statistics of the stimulus presentation (Carandini et al, 2013). Thus real neuronal networks appear to be producing the most accurate representations possible as efficiently as possible by confronting a trade-off between reducing decoding error and also minimizing metabolic cost. The second part of this study attempts to relate the abstract components in this model back to the biophysical mechanisms at work on the individual neuron level. Homeostatic regulation of ionic conductances may implement a cost on spiking that leads to adaptation. The predictive coding framework implemented with LIF neurons relies on symmetrical and tightly tuned connectivity. We hypothesize that homeostasis may necessarily replace some of the symmetrical recurrent connectivity that is needed for shaping the dynamics of the LIF predictive coding network (Fig. 2). As a whole, this study begins with an abstract model of predictive coding and goes from demonstrating adaptation as a consequence of efficient coding to predicting that the abstract costs are linked to the biophysical mechanics of individual neurons.

II-82. Balanced spiking networks preserve and transform input information

Ingmar Kanitscheider[1] INGMAR.KANITSCHEIDER@UNIGE.CH
Ruben Coen-Cagli[2] RUBEN.COENCAGLI@UNIGE.CH
Ruben Moreno-Bote[3] RMORENO@FSJD.ORG

[1]University of Texas at Austin
[2]University of Geneva

[3]Foundation Sant Joan de Deu

The activity of cortical neurons in vivo is highly variable, a fact thought to limit how reliably neural populations encode sensory stimuli. A widely believed origin of cortical spiking variability is the effective stochastic dynamics of balanced neuronal networks, but external sources of variability like stimulus uncertainty and noise are likely to contribute as well. To what extent externally versus internally originated variability impacts sensory encoding is currently unknown. Here we test the hypothesis that, in the realistic case that stimuli are noisy and therefore carry limited information, internally generated variability has negligible effects on coding reliability. Previous work (Moreno et al, 2014) has shown that some classes of spiking networks can preserve input information and that simple linear decoders using population spike counts can extract most of the information. However, whether more realistic networks with leaky integrate-and-fire (LIF) neurons, random connectivity and non-linear tuning can also preserve input information has not been studied. First, we find that for large network size, the output information of LIF networks saturates to the asymptotic input information and, therefore, no information is lost. Second, for small network sizes, the information loss of the network is well-predicted by a linear non-linear Poisson (LNP) model, where the non-linearity is given by the mean-field approximation to the rates of the spiking network. This result implies that only correlations induced by the input fluctuations affect output information, while those induced by recurrent connectivity are irrelevant, even if they are large. Finally, we find that due to non-linear dynamics, the network is able to convert information encoded in non-linear form into linearly decodable information. Our results provide support for our hypothesis that internally generated variability does not affect coding reliability and suggest a prominent role of the dynamics of spiking neuronal networks for non-linearly reformatting information.

II-83. Causal inference by spiking networks

Ruben Moreno-Bote[1] RMORENO@FSJD.ORG
Jan Drugowitsch[2] JDRUGO@GMAIL.COM

[1]Foundation Sant Joan de Deu
[2]University of Geneva

Because sensory stimuli are noisy and ambiguous, the brain must implement probabilistic inference to solve problems such as object recognition. The neuronal algorithms that achieve these probabilistic computations are currently unknown. These algorithms operate on spiking codes, with the advantage that inter-neuronal communication is both sparse over time and energetically very efficient. However, a potential disadvantage of spike-based codes is their possible sensitivity to shifts in the timing of individual spikes caused by inherent sources of variability and noise in the brain. Here we demonstrate that a family of non-linear, high-dimensional optimization problems, which correspond to causal inference problems, can be solved efficiently and exactly by networks of spiking integrate-and-fire neurons. We show that these spiking networks encode an input vector as a linear combination of stored vectors ('features') weighted by their firing rates, and that this encoding minimizes the reconstruction error of the input vector with non-negativity constraints. The network infers the set of most likely causes ('objects') given the set of observations. These causes are encoded in the firing rate of the neurons in the network. The network solves this inference problem using 'explaining away' as the underlying computational principle, implemented by highly tuned inhibition. The algorithm features high performance even when the network intrinsically generates variable spike trains, or the timing of spikes is scrambled by external sources of noise, thus demonstrating its robustness to noise and spike-timing variability. Specifically, most of the input information can be recovered from the slow covariations of firing rates across cells. This type of networks might underlie tasks such as odor identification and classification, and provide new vistas on the type of robust computations that can be performed with spiking networks.

II-84. Training spiking neural networks with long time scale patterns

Dongsung Huh DONGSUNGHUH@GMAIL.COM
Peter Latham PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

Neurons in the brain often exhibit complex activity patterns, with fluctuations on time scales of several seconds. The generation of complex patterns is critical for directing movements, and is likely to be involved in processing time-varying input (such as speech). However, it is not yet understood how networks of spiking neurons, with time constants of only a few milliseconds, could exhibit such slow dynamics. This should be contrasted with rate-based neural networks, which can be trained to generate arbitrary complex activity patterns by an iterative training method (FORCE learning, Susillo and Abbott 2009). So far, however, FORCE learning has not led to successful training of spiking neural networks. Here we show that by modifying the connectivity away from the standard, fully random networks, FORCE learning can be used to train networks of spiking neurons. In particular, we consider networks of inhibitory neurons whose recurrent connectivity is modified according to an anti-Hebbian-like rule: neurons that tend to fire at the same time have sparser connectivity. The reliability of the network activity, when driven by external stimulus, depends critically on the window of the anti-Hebbian plasticity. In particular, there exists an optimal window that maximizes the reliability. The plasticity window also affects the autonomous dynamics of the network's firing rate to generate slow, chaotic fluctuations. The modified network was successfully trained to generate complex patterns with long periods, up to several seconds. Moreover, the feedback from the readout neuron readily stabilized the rate activity of the network without disrupting the irregular spiking pattern. This result may be crucial both for other types of learning in spiking networks and for understanding the complex neural dynamics of the brain.

II-85. Predicting the dynamics of network connectivity in the neocortex

Yonatan Loewenstein[1] YONATAN.LOEWENSTEIN@MAIL.HUJI.AC.IL
Uri Yanover[1] URI.YANOVER@GMAIL.COM
Simon Rumpel[2] SIMON.RUMPEL@IMP.AC.AT

[1]The Hebrew University of Jerusalem
[2]Johannes Gutenberg University Mainz

In contrast to electronic circuits, the neural architecture in the neocortex exhibits constant remodeling. The functional consequences of these modifications are poorly understood, in particular because the determinants of these changes are largely unknown. We studied dendritic spines of pyramidal neurons in the mouse auditory cortex as proxies for excitatory synapses onto these neurons. Using chronic in vivo two-photon imaging, we longitudinally followed several thousand dendritic spines and analyzed their morphologies and life-times. In order to identify those modifications that are predictable from current network state, we applied non-linear regression to quantify the predictive value of spine age and several morphological parameters to the survival of a spine. We found that spine age, size and geometry are parameters that can provide independent contributions to the prediction of the longevity of a synaptic connection. Using these parameters, we found that a considerable fraction of the measured connectivity changes in the dataset can be predicted for multiple intervals into the future. Understanding the dynamics of neuronal circuits is of particular importance for the functional interpretation of connectivity matrices typically obtained from single time point electron microscopy data that will become increasingly available in the future. Therefore, we used this framework to emulate a serial sectioning electron microscopy experiment and demonstrated how incorporation of morphological information of dendritic spines from a single time-point allows estimation of future connectivity states. Our finding that elements in dynamics of network connectivity are predictive is a strong indication that many of the synaptic changes observed in vivo do not result from the acquisition of new memories. Moreover, our model may help us distinguish those network changes that are associated with learning and memory from those that are not.

II-86. Superlinear precision and memory in simple population codes

David Schwab[1] DSCHWAB@PRINCETON.EDU
Ila Fiete[2] ILAFIETE@MAIL.CLM.UTEXAS.EDU

[1]Northwestern University
[2]University of Texas at Austin

A widely used tool for quantifying the precision with which a population of noisy sensory neurons encodes the value of an external stimulus is the Fisher Information (FI). Maximizing the FI is also a commonly used objective for constructing optimal neural codes. The primary utility and importance of the FI arises because it gives, through the Cramer-Rao bound, the smallest mean-squared error (MSE) achievable by any unbiased estimator. However, it is well-known that when neural firing is sparse, the bound is not tight and optimizing the FI can result in codes that perform poorly when considering the resulting MSE, a measure with direct biological relevance. Here we construct optimal population codes by directly minimizing MSE. We study the scaling properties of the resulting sensory network, focusing on tuning curve width, then extend our analysis to persistent activity networks. In particular, we consider a neural population that encodes a one-dimensional periodic stimulus through unimodal tuning curves. Neurons fire Poisson spikes at rates determined by the location of the stimulus relative to their tuning curves. The MSE contains two terms: local errors, which are well-described by the inverse FI; and more global threshold errors. We derive the optimal tuning width that minimizes MSE, as a function of the number of neurons. The MSE-optimal tuning curves in an N-neuron population enable superlinear reductions in MSE by a factor log(N)/N^2. Our analytic argument is confirmed by simulations of maximum likelihood decoding with systems of various sizes and tuning widths. We then study ring attractor memory networks and measure bump diffusivity in systems of varying size and interaction width. We find similar scalings for optimal tuning width and diffusivity as in the sensory system. Finally, we discuss how minimization of threshold error risk may drive tuning curves to much wider than MSE-optimal values.

II-87. Pattern decorrelation arising from inference

Simon Barthelme[1] SIMON.BARTHELME@UNIGE.CH
Agnieszka Grabska Barwinska[2] AGNIESZKA.GRABSKABARWINSKA@GMAIL.COM
Peter Latham[3] PEL@GATSBY.UCL.AC.UK
Jeff Beck[4] JEFF.BECK@DUKE.EDU
Zachary Mainen[5] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG
Alexandre Pouget[1] ALEX.POUGET@UNIGE.CH

[1]University of Geneva
[2]Deep Minds
[3]Gatsby Computational Neuroscience Unit, UCL
[4]Duke University
[5]Champalimaud Neuroscience Programme

Before olfactory information reaches the cortex it is first processed by the olfactory bulb (OB). Despite intensive study there is still a great deal of uncertainty about the computational role of the OB. An important experimental finding (Friedrich and Laurent, 2001) is that odour representations in the bulb tend to differentiate over time, a phenomenon known as “pattern decorrelation”. Odours that are encoded similarly in an early phase become progressively easier to tell apart. The original interpretation of decorrelation in the bulb is that decorrelation is the computational role of the bulb; i.e., that the bulb preprocesses odour signals to facilitate discrimination by the cortex. Here we argue that decorrelation instead arises as a side-effect of statistical inference. The demixing hypothesis (Beck et al., 2012; Grabska-Barwinska et al., 2013; Rokni et al., 2014) suggests that the role of olfactory processing is to solve a cocktail party problem: at any given time, many odours are present, and animals need to know which and at what concentration. Demixing can be implemented using a predictive model in which: (a) the cortex is trying to explain incoming signals and (b) activity in mitral cells (MCs) in the OB can be understood as an error signal or residual. We show that decorrelation arises naturally in such a model, which could imply that the OB is not "preprocessing" olfactory data in a feedforward manner, but rather that it is involved in a dynamical process of inference that involves two-way communication between the OB and olfactory cortices

II-88. Dual homeostatic mechanisms cooperate to optimize single-neuron coding

Jonathan Cannon CANNON@BRANDEIS.EDU
Paul Miller PMILLER@BRANDEIS.EDU
Brandeis University

Multiple coexistent mechanisms of single-neuron homeostatic firing rate regulation have been identified, including the scaling of intrinsic neuronal excitability and the scaling of incoming synapse strength [1]. However, the nature of the balance and interaction between distinct homeostatic mechanisms has not been characterized. The harmonious coexistence of multiple homeostatic mechanisms is not trivial: in a simple model, two mechanisms attempting to achieve slightly different firing rates may work against each other and "ramp up," scaling homeostatic variables to extreme limits. Given this complication, it is difficult to see how an organism would benefit from such redundant homeostasis. We perform an analytical and computational study of the interaction of intrinsic and synaptic homeostatic scaling in a family of simple models, including rate models and spiking models. Using tool from stochastic and deterministic dynamical systems, we show analytically how the relationship between the target rates and homeostatic response functions of the two mechanisms determines the existence and stability of a homeostatic equilibrium, and provide conditions on these targets and functions that guarantee the existence of a stable equilibrium. We address the question of why redundant homeostatic mechanisms might be evolutionarily preserved by demonstrating that a neuron implementing both homeostatic mechanisms (but not either one without the other) can regulate both the mean and the variance of its output in response to changes in it input statistics. We show in simulation that this capacity allows a neuron to maximize mutual information between its input and output across a wide range of input statistics; that it allows a feedforward neuronal network to preserve input information across multiple layers; and that it allows a recurrent network to produces an approximation to a line attractor/neural integrator.

II-89. Inferring readout of Distributed population codes without massively parallel recordings

Kaushik J Lakshminarasimhan[1] JLAKS@CNS.BCM.EDU
Dora Angelaki[1,2] DANGELAKI@CNS.BCM.EDU
Xaq Pitkow[1,2] XAQ@CNS.BCM.EDU
[1]Baylor College of Medicine
[2]Rice University

Information about task-relevant variables is often distributed among neurons across multiple cortical areas. Neuronal responses are rarely independent of each other, but are correlated to some degree due to common input as well as recurrent message-passing. Consequently, determining how these neurons collectively drive behavioral changes requires not only examining how individual neurons are correlated with behavior, but also estimating the correlated variability among neurons. Precisely estimating the structure of correlated variability requires massively parallel recordings, which remains very difficult with current technology. Fortunately, it has recently been shown that the expansion in neural representation from sensory periphery will lead to a predictable pattern of correlations that ultimately limits the information content in brain areas downstream. We examined the implications of these so-called information-limiting correlations for the readout of distributed population codes in a simple discrimination task. Surprisingly we found that both the behavioral precision, as well as the correlation of individual neurons with behavioral choice (choice correlation) were determined largely by the relative magnitudes of neuronal weights in the different brain areas and not on their specific pattern. We also found that, in the presence of information-limiting correlations, the choice correlations of neurons within an area should all scale by the same factor following inactivation of other potentially task-relevant brain areas. Together, our results lead to a novel framework for inferring how different brain areas contribute to behavioral response. Specifically, we show that the contribution of a brain area can be inferred simply by observing how the magnitude of choice correlations of individual neurons within the area and the behavioral precision are affected by inactivating other areas, thus obviating the need for large-scale recordings.

II-90. Robust non-rigid alignment of volumetric calcium imaging data

Benjamin Poole POOLE@CS.STANFORD.EDU
Logan Grosenick LOGANG@STANFORD.EDU
Michael Broxton BROXTON@STANFORD.EDU
Karl Deisseroth DEISSEROO@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU

Stanford University

Improvements in genetically-encoded calcium indicators have led to ever broader application of calcium imaging as a means to record population activity with cellular resolution in targeted neural populations. However, during in vivo experiments, brain motion relative to the sensor results in motion artifacts both within and between images. Such artifacts render it difficult to accurately estimate calcium transients using fixed spatial regions of interest, making image alignment a critical first step in analyzing calcium imaging datasets. Previous alignment algorithms for calcium imaging require the specification of a fixed template image that can be matched to each frame. These algorithms are ineffective when applied to the latest generation of optical activity reporters, which have both negligible background signal and large changes in neural activity over time. Our work is based on RASL, an image alignment technique that learns a low-rank matrix that represents a set of templates that can be adaptively combined to match each frame, and a sparse matrix that represents deviations from the adaptive template (Peng et al., 2010). This allows us to leverage low-dimensional dynamics in neural activity to robustly align functional data without a fixed template. We extend RASL to account for three-dimensional translations, rotations, and non-rigid deformations such as those caused by scanning artifacts. To scale RASL to datasets containing millions of voxels and thousands of frames, we introduce several extensions including randomized decompositions and online alignment. We validate our technique on images and volumes from two-photon and light field microscopy (Grosenick et al., 2009; Broxton et al., 2013), showing improved accuracy and reduction of motion artifacts compared with existing techniques. An implementation of our algorithm is released as a Python package using the Apache Spark distributed computing framework and integrated with Thunder (Freeman et al., 2014).

II-91. Marginalization in Random Nonlinear Neural Networks

Rajkumar Vasudeva Raju[1] RAJDBZ@GMAIL.COM
Xaq Pitkow[1,2] XAQ@RICE.EDU

[1]Rice University
[2]Baylor College of Medicine

Computations involved in tasks like object recognition and causal reasoning in the brain require a type of probabilistic inference known as marginalization. In probability theory, marginalization corresponds to averaging over irrelevant variables to obtain the probability of the variables of interest. This is a fundamental operation whenever an animal is uncertain about variables that affect the stimuli but are not task relevant. Animals often exhibit behavior consistent with marginalizing over some variables, but the neural substrate of this computation remains

unknown. It has been previously shown (Beck et al., 2011) that marginalization can be performed optimally by a deterministic nonlinear network that implements a quadratic interaction of neural activity with divisive normalization. Here we show that a simpler network can perform essentially the same computation. These Random Nonlinear Networks (RNN) are feedforward networks with a single hidden layer, sigmoidal activation input-output functions and normally-distributed weights connecting the input to the hidden layer. We train the output weights connecting the hidden units to an output population, such that the output model accurately represents a desired marginal probability distribution without significant information loss compared to optimal marginalization. Simulations for the case of linear coordinate transformations show that the RNN model has good marginalization performance, except for extremely uncertain inputs that have very low population spike counts. The proposed model is much more generic and suggests that a larger, less constrained class of nonlinearities can be used for marginalization. Behavioral experiments, based on the results obtained, could then be used to identify if animals exhibit biases consistent with this approximation rather than exact marginalization.

II-92. Relating reachability to classifiability of trajectories in E-I networks with high dimensional input

Maxwell Wang MAXWANG@WUSTL.EDU
ShiNung Ching SHINUNG@ESE.WUSTL.EDU

Washington University in St. Louis

We characterize the space of time-varying activation patterns (trajectories) that could be induced in E-I networks by high dimensional afferent inputs. This space - the set of reachable trajectories - is determined by the network structure and dynamics, as well as the input itself. The supposition is that in order to support meaningful trajectory classification, such a space should be sufficiently large (i.e., many unique trajectories could be created), but also insensitive (i.e., trajectories should not change drastically with small input perturbations). In this spirit, we develop several analyses, motivated by formal systems theory, to characterize fundamental tradeoffs between a network's expressiveness (the size of the reachable set) and its sensitivity to distraction and noise. More specifically, we formulate E-I networks consisting of Izhikevich-type neurons receiving combinatorial excitation from a high dimensional feature-space. A particular input realization creates a corresponding (but, not necessarily unique) output trajectory. We use dimensionality reduction techniques to characterize the relative size of the reachable space as compared with the dimensionality of the inputs. Further, we use a sensitivity analysis to ascertain the perturbation of trajectories relative to input variation. Finally, we relate these properties to putative information processing by performing a classifiability analysis in the full trajectory space using an inner product distance. In all cases, a convex dependence on key physiological quantities such as E-I ratio, is exhibited, indicating how these parameters may trade off for balance between expressiveness and sensitivity.

II-93. Extracting spatial-temporal coherent patterns from large-scale neural recordings using dynamic mode

Bingni Brunton BBRUNTON@UW.EDU
Lise Johnson JOHNSON.LISE@GMAIL.COM
Jeffrey Ojemann JOJEMANN@U.WASHINGTON.EDU
J. Nathan Kutz KUTZ@UW.EDU

University of Washington

There is a growing need in the neuroscience community to understand and visualize large-scale recordings of neural activity, big data acquired by tens or hundreds of electrodes simultaneously recording dynamic brain activity over minutes to hours. Such dynamic datasets are characterized by coherent patterns across both space and time, yet existing computational methods are typically restricted to analysis either in space or time separately. We will

describe the adaptation of dynamic mode decomposition (DMD), an algorithm originally developed for the study of fluid physics, to large-scale neuronal recordings. DMD is a modal decomposition algorithm that describes high-dimensional data using coupled spatial-temporal modes; the resulting analysis combines key features of performing principal components analysis (PCA) in space and power spectral analysis in time. The algorithm scales easily to very large numbers of simultaneously acquired measurements. We validated the DMD approach on sub-dural electrode array recordings from human subjects performing a known motor activation task. Next, we leveraged DMD in combination with machine learning to develop a novel method to extract sleep spindle networks from the same subjects. We suggest that DMD is generally applicable as a powerful method in the analysis and understanding of large-scale recordings of neural activity.

II-94. Model-based reinforcement learning with spiking neurons

Johannes Friedrich JF517@CAM.AC.UK
Máté Lengyel M.LENGYEL@ENG.CAM.AC.UK
University of Cambridge

Behavioural and neuroscientific data on reward-based decision making point to a fundamental distinction between habitual and goal-directed action selection. An increasingly explicit set of neuroscientific ideas has been established for habit formation, whereas goal-directed control has only recently started to attract researchers' attention. While using functional magnetic resonance imaging to address the involvement of brain areas in goal-directed control abounds, ideas on its algorithmic and neural underpinning are scarce. Here we present the first spiking neural network implementation for goal-directed control that selects actions optimally in the sense of maximising expected cumulative reward. Finding the optimal solution is commonly considered a difficult optimisation problem, yet we show that it can be accomplished by a remarkably simple neural network over a time scale of a hundred milliseconds. We provide a theoretical proof for the convergence of the neural dynamics to the optimal value function. We extend the applicability of the model from discrete to continuous state spaces using linear function approximation. Further, we also show how the environmental model can be learned using a local synaptic plasticity rule in the same network. After establishing the performance of the model on various benchmark tasks from the machine learning literature, we present a set of simulations reproducing behavioural as well as neurophysiological experimental data on tasks ranging from simple binary choice to sequential decision making. We also discuss the relationship between the proposed framework and other models of decision making.

II-95. A flexible and tractable statistical model for in vivo neuronal dynamics

Simone Carlo Surace[1,2] SURACE@PYL.UNIBE.CH
Jean-Pascal Pfister[2] JPFISTER@INI.UZH.CH
[1]University of Bern
[2]Institute of Neuroinformatics, Zurich

The increasing amount of intracellular recordings of spontaneous activity as well as the increasing number of theories which critically rely on a characterization of spontaneous activity calls for a proper quantification of spontaneous intracellular dynamics. Here we propose a statistical model of spontaneous activity which is very flexible and remains tractable. More specifically, we propose a doubly stochastic process where the subthreshold membrane potential follows a Gaussian process and the spike emission intensity depends nonlinearly on the membrane potential as well as the previous spiking history. Moreover, the separation of sub- and suprathreshold dynamics is done in terms of a spike shape kernel, which captures the stereotypic shape of the action potential and is also learned during model fitting.

II-96. Short term persistence of neural signals

Grant Gillary GGILLARY@GMAIL.COM
Rudiger von der Heydt VON.DER.HEYDT@JHU.EDU
Ernst Niebur NIEBUR@JHU.EDU

The Johns Hopkins University

Persistent post-stimulus activity is an ubiquitous phenomenon in cortical systems, often associated with working memory found in central structures. However, even early sensory areas often have persistent stimulus representations which decay over the course of seconds. Since they respond quickly to a stimulus onset, these networks exhibit two different time constants. Such dynamics, in which learning a memory's contents is much faster than its decay, is obviously useful for most memory structures. We use a model based on derivative feedback with dynamically changing synapses to show how cortical networks might exhibit fast response to a stimulus onset followed by a post-stimulus slow decay on a time scale consistent with firing rate changes in cortex. Such networks effectively decouple the post-stimulus persistence of a representation from its synaptic strength.

II-97. The mesoscale mouse connectome is described by a modified scale-free graph

Rich Pang[1] RPANG@UW.EDU
Sid Henriksen[2] S.HENRIKSEN@NCL.AC.UK
Mark Wronkiewicz[1] WRONK@UW.EDU

[1]University of Washington
[2]Laboratory of Sensorimotor Research, NEI, NIH

Graph theory provides a mathematical framework for studying the structure of networks. We present a graph theoretic analysis of the Allen Institute for Brain Science's mouse connectivity atlas and show that the mouse connectome exhibits an inverse relationship between node degree and clustering coefficient. Standard random graphs, such as small-world and scale-free graphs, do not account for this property. We propose a simple binary random graph in which nodes are localized in space and connections are formed with greater probability between high degree nodes (preferential attachment, as in scale-free graphs) and nearby nodes (spatial proximity). This biophysically inspired network model accounts for the inverse relationship between degree and clustering coefficient, as well as the distributions of these two quantities. Further, the model's response to simulated lesions more closely resembles the response of the connectome than standard random graphs. Our work suggests that a network growth process based on spatial proximity of nodes and preferential attachment can capture many characteristics of the mouse connectome. (All authors contributed equally to this work.)

II-98. Memory savings through unified pre- and postsynaptic STDP

Rui Ponte Costa[1] RUI.COSTA@CNCB.OX.AC.UK
Robert Froemke[2] ROBERT.FROEMKE@MED.NYU.EDU
P. Jesper Sjostrom[3] JESPER.SJOSTROM@MCGILL.CA
Mark van Rossum[4] MVANROSS@INF.ED.AC.UK

[1]University of Oxford
[2]New York University
[3]Montreal General Hospital
[4]University of Edinburgh

While the expression locus of long-term synaptic plasticity has been debated for decades, there is increasing evidence that it is both pre- and postsynaptically expressed (Padamsey and Emptage, Philos Trans R Soc 2014). However, most models are agnostic about expression locus and the functional role of this segregation remains mysterious. We introduce a novel phenomenological model of spike- timing-dependent plasticity (STDP) that unifies pre- and postsynaptic components. Our unified model captures the presynaptic aspects of STDP and the co-modification of short and long-term synaptic plasticity, consistent with a wide range of experimental results from rat visual cortex (Sjostrom et al., Neuron 2001 and 2003). Functionally, this unified STDP rule develops receptive fields with improved reliability, as has been observed in rat auditory cortex in vivo (Froemke et al., Nat Neuro 2013). In addition, this unified model enables fast relearning of previously stored information, in keeping with the memory savings theory (Ebbinghaus, Leipzig: Duncker & Humblot 1885), which refers to rapid relearning through hidden storage of forgotten but previously acquired memories. Thus our work shows that unified pre- and postsynaptic STDP leads both to improved discriminability and more flexible learning.

II-99. Synaptic consolidation: from synapses to behavioral modeling

Lorric Ziegler[1] LORRIC.ZIEGLER@GMAIL.COM
Friedemann Zenke[1] FRIEDEMANN.ZENKE@EPFL.CH
David B Kastner[2] DBKASTNER@GMAIL.COM
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH

[1]Ecole Polytechnique Federale de Lausanne
[2]University of California, San Francisco

Synaptic plasticity, a key process for memory formation, manifests itself across different time scales, ranging from a few seconds for plasticity induction, up to hours or even years for consolidation and memory retention. We developed a novel three-layered model of synaptic consolidation that accounts for data across a large range of experimental conditions including cross-tagging, tag-resetting and depotentiation, as well as meta-plasticity. Consolidation occurs in the model through the interaction of the synaptic efficacy or weight with a scaffolding variable by a read-write process, mediated by a tagging-related variable. Plasticity inducing stimuli modify the weight, but the state of tag and scaffold can only change if a write protection mechanism is overcome. Only a strong or sustained stimulus is able to remove the write protection from weight to tag. Consolidation requires the removal of a second layer of write protection, from tag to scaffold, not possible without the presence of a neuromodulator such as dopamine. Second, our model connects synapses to behavior. It makes a link from depotentiation protocols in vitro to behavioral results regarding the influence of novelty on inhibitory avoidance memory. It shows that experimental data on tagging and consolidation is sufficient to capture the essence of behavioral phenomena.

II-100. Neuronal avalanches in behaving monkeys

Tiago Lins Ribeiro[1] TIAGOLINSRIBEIRO@GMAIL.COM
Shan Yu[1] YUSHAN.MAIL@GMAIL.COM
Samantha Chou[1] SAMANTHA.CHOU@NIH.GOV
Dietmar Plenz[2] PLENZD@MAIL.NIH.GOV

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

Ongoing neuronal dynamics has been shown to be maintained at a phase transition between an inactive (sub-critical) and a self-sustained state (supercritical), through a precise balance of excitation/inhibition. At this critical point, spatiotemporal clusters of synchronized firing (as measured through large negative deflections of local-field potentials, or nLFPs) emerge as a scale-free activity, called neuronal avalanches [1], and characterized by power-

law size distributions. Those were studied in many different experimental setups, but mostly restricted to ongoing dynamics [1-3] (avalanches in freely-behaving animals were studied in the context of spiking activity [4]). Here we analyze neuronal avalanches in monkeys (adult Macaca mulatta) subjected to a working memory task: they had to reach for a reward in a feeder either at the left or at the right side, depending on which cue was presented. The LFPs were recorded from 10x10 multielectrode arrays implanted in the left arm representation region of the premotor cortex while the monkeys performed multiple trials. We found that there was a modulation of the rate of nLFPs during the cue presentation only for trials in which the correct feeder was the left one (we call these left trials), as shown in Fig. 1A (black curves). Together with that increase in the activity rate for left trials, we observed a decrease in the Fano factor, a measure of the variance in the signal over different trials (Fig. 1A, colored curves), in agreement with what is found in the literature [5]. Despite non-stationary rates and sizes, patterns in size distribution invariably obeyed a power law over the duration of all trials (Fig. 1B&D) as predicted for avalanche dynamics. This was abolished using a shift-predictor, which randomizes individual recording sites across trials (Fig. 1C&D). We conclude that the critical state is, on average, maintained even though there is an evoked response.

II-101. Optimal feature integration in critical, balanced networks

Udo Ernst UDO@NEURO.UNI-BREMEN.DE
Nergis Tomen NERGIS@NEURO.UNI-BREMEN.DE
University of Bremen

Recent experimental and theoretical work established the hypothesis that cortical neurons operate close to a critical state which describe a phase transition from chaotic to ordered dynamics. This state is suggested to optimize several aspects of neuronal information processing. However, although critical dynamics have been demonstrated in recordings of spontaneously active cortical neurons, the relations between criticality and active computation remain largely unexplored. In our study we focus on visual feature integration as a prototypical and prominent example for cortical computation. In particular, we construct a network of integrate-and-fire neurons with balanced excitation and inhibition which performs contour integration on a visual stimulus. In dependence on synaptic coupling strength, the network undergoes a transition from subcritical dynamics, over a critical state, to a highly synchronized regime. We show that for different measures, contour detection performance is always maximized near or at the critical state. In particular, spontaneous synchronization contains far more information about the presence of a target in a stimulus than coding schemes based on trial-averaged firing rates. At the same time, our paradigm provides a novel unifying account for stylized features of cortical dynamics (i.e. high variability) and contour integration (i.e. high performance and robustness to noise) known from psychophysical and electrophysiological studies. Acknowledgement: This research project was funded by the BMBF (Bernstein Award Udo Ernst, grant no. 01GQ1106).

II-102. Detecting effective connectivity in spiking data: implications for coding space in rat hippocampus

Cristina Savin[1] CSAVIN@IST.AC.AT
Margret Erlendsdottir[2] MARGRET.ERLENDSDOTTIR@GMAIL.COM
Jozsef Csicsvari[1] JOZSEF.CSICSVARI@IST.AC.AT
Gasper Tkacik[1] GTKACIK@IST.AC.AT
[1]IST Austria
[2]Yale University

Correlations in spiking activity arise in part due to local network interactions and, as such, can provide insights into underlying circuit computation. While the use of correlations to infer effective (or even anatomical) connectivity has a long history in neuroscience [Aertsen1989], detecting such couplings remains challenging in practice. The reason is that total correlations reflect not only local interactions, but also shared stimulus preferences and global modulatory influences, e.g. oscillations. Thus, in order to make inferences about connectivity one needs to first factor out other shared sources of variability. To this aim, we developed a new maximum entropy null model that takes alternative sources of correlations into account in absence of any effective coupling and used it to compare against neural data. After validating the method on simulated data with known statistics, we applied it to tetrode recordings from awake behaving rats. We found a subset of CA1 neurons to be functionally coupled, consistent with previous reports of collective responses in CA1 [Harris2003]. Moreover, this connectivity reflects stimulus preferences: connected neurons tend to have strongly (anti-)correlated place fields. To investigate their importance we constructed a network model in which we could systematically vary the strength of these interactions, while preserving overall firing rates. We found that data-like connectivity improves the precision of the spatial representation, suggesting that collective behaviour in CA1 optimizes the encoding of the animal's position. More generally, the work highlights the utility of maximum entropy models for making sense of neural population responses, by creating a hierarchy of precise controls against which rigorous statistical tests are possible.

II-103. Hebbian and non-Hebbian plasticity orchestrated to form and retrieve memories in spiking networks

Friedemann Zenke[1] FRIEDEMANN.ZENKE@EPFL.CH
Everton J Agnes[2] EVERTON.AGNES@GMAIL.COM
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH

[1]Ecole Polytechnique Federale de Lausanne
[2]Universidade Federal do Pampa

Synaptic plasticity, the putative principle underlying learning and memory formation, manifests itself in various forms and across different timescales. Hebbian forms of plasticity can be induced on the timescale of seconds to minutes. To avoid rapid run-away effects in simulated neural networks, synapse growth has to be either restricted [1] or subjected to compensatory mechanisms [2, 3] which have to act on a timescale of seconds [4]. In many existing models the latter is achieved through rapid homeostatic scaling [2, 3]. This, however, is in conflict with the timescale of homeostatic mechanisms observed in nature (hours to days [5]) and renders synaptic plasticity non-local [6]. Here we show that the interplay of Hebbian homosynaptic plasticity with forms of non-Hebbian heterosynaptic plasticity alone is sufficient for assembly formation and memory recall in a spiking recurrent network model. First, receptive fields emerge from external afferent connections during repeated sensory stimulation. Cells with the same stimulus preference then form assemblies characterized by strong recurrent excitatory connections. Even days after formation and despite ongoing network activity and synaptic plasticity, these structures are stable and can be recalled through selective delay activity when a distorted cue is fed into the network. Blocking individual components of plasticity prevents stable functioning as a memory network. Our modeling results suggest that the diverse forms of plasticity in the brain are orchestrated toward common functional goals. [1] G. Mongillo, E. Curti, S. Romani, and D. J. Amit. 2005. [2] A. Lazar, G. Pipa, and J. Triesch. 2009. [3] A. Litwin-Kumar and B. Doiron. 2014. [4] F. Zenke, G. Hennequin, and W. Gerstner. 2013. [5] G. G. Turrigiano, K. R. Leslie, N. S. Desai, L. C. Rutherford, and S. B. Nelson. 1998. [6] E. Bienenstock, L. Cooper, and P. Munro. 1982.

II-104. Synaptic plasticity as Bayesian inference

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

Graz University of Technology

We present a model that explains how a Bayesian view of synaptic plasticity as probabilistic inference could be implemented by networks of spiking neurons in the brain through sampling. Such Bayesian perspective of brain plasticity has been proposed on general theoretical grounds (Pouget et al. 2013). But it is open how this theoretically attractive model could be implemented in the brain. We propose that apart from stochasticity on the level of neuronal activity (neural sampling), also plasticity should be understood as stochastic sampling from a posterior distribution of parameters ("synaptic sampling"). This model is consistent with a number of puzzling experimental data, such as continuing spine motility in the adult cortex. In addition it provides desirable new functional properties of brain plasticity such as immediate compensation for perturbations and integration of new tasks. Furthermore it explains how salient priors such as sparse synaptic connectivity and log-normal distributions of weights could be integrated in a principled manner into synaptic plasticity rules.

II-105. Normalization of excitatory-inhibitory balance by cortical spike-timing-dependent plasticity

James D'Amour JAMESDMR@GMAIL.COM
Robert Froemke ROBERT.FROEMKE@MED.NYU.EDU

New York University

Synapses are plastic and can be modified by changes of spike timing. While most studies of long-term synaptic plasticity focus on excitation, inhibitory plasticity may be critical for controlling information processing, memory storage, and overall excitability in neural circuits. Here we examine spike-timing-dependent plasticity (STDP) of inhibitory synapses onto layer 5 neurons in slices of mouse auditory cortex, together with concomitant STDP of excitatory synapses. Pairing pre- and postsynaptic spikes potentiated inhibitory inputs irrespective of precise temporal order within ~10 msec. This was in contrast to excitatory inputs, which displayed an asymmetrical STDP time window. These combined synaptic modifications both required NMDA receptor activation, and adjusted the excitatory-inhibitory ratio of events paired together with postsynaptic spiking. Finally, subthreshold events became suprathreshold, and the time window between excitation and inhibition became more precise. These findings demonstrate that cortical inhibitory plasticity requires interactions with co-activated excitatory synapses to properly regulate excitatory-inhibitory balance.

III-1. Proprioceptive feedback modulates motor cortical tuning during brain-machine interface control

Danielle Rager[1] DRAGER@ANDREW.CMU.EDU
John Downey[2] JDOWNEY388@GMAIL.COM
Jennifer Collinger[2] COLLINGR@PITT.EDU
Douglas Weber[2,3] DJW50@PITT.EDU
Michael Boninger[2] BONINGER@UPMC.EDU
Robert Gaunt[2] RAG53@PITT.EDU
Valerie Ventura[1] VVENTURA@STAT.CMU.EDU

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

Loss of proprioception is known to severely impair motor control, but the neural mechanisms by which proprioception aids in the planning and execution of visually guided movements are not well understood. We investigate the impact of providing proprioceptive feedback to a human subject with tetraplegia and intact sensation who was implanted with two 96-channel microelectrode arrays in primary motor cortex (M1). Passive proprioceptive feedback was provided either by manually moving the subject's arm in conjunction with the brain-machine interface (BMI)-controlled robotic arm or by moving the subject's arm in an exoskeleton under BMI control. Performance of a BMI-assisted reaching task degrades when we allow a visually-trained decoder to leverage the subject's own proprioceptive signals, indicating that proprioceptive feedback alters M1 tuning structure. We show that velocity tuning is stable across trials in the visual feedback only (V) condition and in the visual and proprioceptive feedback (VP) condition, but is not stable across the two conditions in greater than 70% of recorded M1 channels. We then identify M1 channels that are preferentially tuned in the V or VP feedback condition across days. Finally, we show that proprioceptive input is correlated with a population decrease in mean firing rate but a population increase in spiking correlations. These findings suggest that M1 does not encode movements with an invariant set of preferred directions, but rather with a latent tuning structure that is dependent on the sensory feedback modalities available. Therefore, new decoders may need to be developed for closed-loop BMI control with natural or surrogate somatosensory feedback.

III-2. Corticospinal neuroprosthetic technologies to restore motor control after neuromotor disorders

Tomislav Milekovic[1] TOMISLAV.MILEKOVIC@EPFL.CH
David A. Borton[2] DAVID_BORTON@BROWN.EDU
Marco Capogrosso[1] MARCO.CAPOGROSSO@EPFL.CH
Eduardo Martin Moraud[1] EDUARDO.MARTINMORAUD@EPFL.CH
Jean Laurens[1] JEAN.LAURENS@EPFL.CH
Nick Buse[3] NICHOLAS.BUSE@MEDTRONIC.COM
Peter Detemple[4] PETER.DETEMPLE@IMM.FRAUNHOFER.DE
Tim Denison[3] TIMOTHY.DENISON@MEDTRONIC.COM
Jocelyne Bloch[5] JOCELYNE.BLOCH@CHUV.CH
Silvestro Micera[1] SILVESTRO.MICERA@EPFL.CH
Erwan Bezard[6] ERWAN.BEZARD@U-BORDEAUX2.FR
Gregoire Courtine[1] GREGOIRE.COURTINE@EPFL.CH

[1]Ecole Polytechnique Federale de Lausanne
[2]Brown University
[3]Medtronic Neuromodulation
[4]Fraunhofer ICT-IMM
[5]Lausanne University Hospital

[6]University of Bordeaux

Continuous electrical spinal cord stimulation (ESCS) can improve motor control after various neurological disorders. However, current technologies stimulate the spinal cord continuously, irrespective of the subject's intentions. Development of time varying stimulation protocols synchronized with the movement intentions of the subject may improve the therapeutic effects while reducing the period ESCS is active. Here, we introduce a neuroprosthetic platform capable of spatially selective ESCS controlled by subject's movement intentions decoded from motor cortex neuronal activity. A rhesus macaque monkey was implanted with (i) a 96-microelectrode Blackrock cortical array in the lower limb area of left MI, (ii) an 8-channel electromyogram (EMG) system into eight right leg muscles spanning four joints of the lower limb and (iii) a 16-electrode epidural ESCS array placed over the lumbar spinal cord. All three implants were equipped with modules for wireless data transfer which allowed us to simultaneously record wideband (30kHz) neuronal data and high fidelity EMG signals (2kHz) and initiate temporally and spatially selective ESCS protocols while the monkey walked freely on a horizontal treadmill at 1.5km/h. Based on mapping of muscle responses to single spatially-specific ESCS pulses, we designed “Flexion” and “Extension” ESCS protocols that selectively induced flexion and extension of the right leg. In real time, a linear discriminate analysis (LDA) algorithm predicted Foot Off and Foot Strike events based on 96 channels of multi-unit activity (MUA) and, upon prediction, initiated the Flexion and Extension ESCS protocols during the swing or stance gait phases, respectively. In this way, we modified the locomotion without disrupting the natural rhythmic alternation of movements. Here, we demonstrated a new method to manipulate spinal circuits that can be used for basic neuroscience research and translational medicine. Our results provide a substantial step for the development of Brain-Spinal Interfaces in order to reestablish locomotion in paralyzed individuals.

III-3. Learning shifts visual cortex from bottom-up to top-down dominant states

Hiroshi Makino HMAKINO@UCSD.EDU
Takaki Komiyama TKOMIYAMA@UCSD.EDU

University of California, San Diego

Sensory perception arises by the interplay between external bottom-up sensory information and internal top-down modulations from higher brain areas. Although theoretical models postulate a dynamic transition in sensory cortices from bottom-up to top-down dominant states through learning, circuit mechanisms underlying this process are poorly understood. Here using two-photon calcium imaging in the primary visual cortex (V1) of head-restrained mice, we examined the activity of layer (L) 2/3 excitatory neurons and their major inputs, bottom-up inputs from L4 excitatory neurons and top-down projections from the retrosplenial cortex (RSC) during passive sensory experience and associative learning over days. In both passive and learning conditions, L4 neuron activity gradually decreased, while RSC inputs arriving at L1 became stronger. These asymmetrical changes in the population activity between the two distinct pathways were accompanied by an emergence of anticipatory responses for an associated aversive event in the L2/3 population, which initially responded faithfully to a visual cue. This learning-specific change was present in RSC inputs but not in L4. Furthermore, learning led to a reduction in the activity of somatostatin-expressing interneurons (SOM-INs) that mainly inhibit distal dendrites of L2/3 excitatory neurons and potentially gate top-down inputs arriving at L1. Reduction of the top-down influence after learning by anesthesia, RSC silencing and optogenetic enhancement of SOM-IN activity was individually sufficient to reverse the learning-induced change in the activity pattern of L2/3 neurons. These results reveal a dynamic shift in the balance between bottom-up and top-down pathways in V1 during learning and uncover a role of SOM-INs in controlling this process.

III-4. Medial prefrontal activity during the delay period contributes to learning of a working memory task

Chengyu Li[1,2] TONYLICY@ION.AC.CN
Ding Liu[1,2] LIUDINGABC@ION.AC.CN
Xiaowei Gu[1,2] XWGU@ION.AC.CN
Jia Zhu[1,2] ZHUJIA@ION.AC.CN

[1]Institute of Neuroscience
[2]Chinese Academy of Sciences

Working memory (WM) is essential for cognition by allowing active retention of behaviorally relevant information over a short duration, known as the delay period. Previous studies have shown that the prefrontal cortex (PFC) is crucial for WM, because perturbation of PFC activity impaired WM and WM-related activity was observed during the delay period in neurons of dorsal-lateral PFC (DL-PFC) in primates and medial PFC (mPFC) in rodents. Nevertheless, the functional role of PFC delay-period activity in WM remains unclear. Memory retention and attentional control are leading candidates. However, PFC is also critical for other brain functions and has been suggested to be important for inhibitory control, decision making, or motor selection. These roles cannot be distinguished by a delayed-response task, in which decision making precedes the delay period. In addition, traditional methods for perturbing neural activity, including transcranial magnetic stimulation and electrical stimulation, do not provide the temporal resolution and cell-type specificity required for delineating the functional role of PFC delay-period activity in WM. These issues were addressed in the present study (Liu, et al., 2014, Science) by using a WM task with a delay period designed to temporally separate memory retention from other functions and optogenetic approaches to bidirectionally manipulate mPFC activity of excitatory and inhibitory neurons during the delay period. We optogenetically suppressed or enhanced activity of pyramidal neurons in mouse mPFC during the delay period. Behavioral performance was impaired during the learning phase but not after the mice were well-trained. Delay-period mPFC activity appeared to be more important in memory retention than in inhibitory control, decision making, or motor selection. Furthermore, endogenous delay-period mPFC activity showed more prominent modulation that correlated with memory retention and behavioral performance. Thus, properly regulated mPFC delay-period activity is critical for information retention during learning of a WM task.

III-5. Decision making with ordered discrete responses

Ulrik Beierholm[1,2] BEIERH@GMAIL.COM
Adam Sanborn[3] A.N.SANBORN@WARWICK.AC.UK

[1]University of Birmingham
[2]CN-CR
[3]University of Warwick

Analyses of decision-making behavior have often compared people against a normative standard, assuming that people are attempting to maximize their expected rewards. According to the Bayesian formulation people's prior beliefs are combined with the likelihood of the observed stimulus via Bayes rule to produce posterior beliefs, which in turn are integrated with the loss function to determine which response has the highest expected value. When people are asked to choose from among a small set of discrete responses, the most straightforward approach is to treat the responses as unrelated, allowing the use of a multinomial prior distribution. However, when asked to make a continuous response, order is very important and responses are often modeled with normal distributions for the prior and likelihood. Loss functions also differ in that discrete response models tend to use an all-or-none loss function while continuous response models often use the mean of the posterior. Lying between these two well-explored cases is decision making with ordered discrete responses. This kind of decision-making is prevalent outside the laboratory, e.g. counting numbers of objects requires an ordered estimation and a discrete response. We investigated how people make decisions with a small number of ordered discrete responses, using a numerosity task in which participants are asked to count the number of dots that were shown in a very briefly

presented display. We characterized the kinds of prior distributions, likelihood distributions, and loss functions that people used in this task. People's choices were not well described by either common discrete or common continuous models of normative decision-making. The likelihoods and loss functions reflected the ordered nature of the responses, but people learned complex prior distributions that reflected the discrete nature of the task. Hence the best explanation of people's decision making with ordered discrete responses involved aspects of both approaches.

III-6. Mnemonic manifolds: convergent low-dimensional dynamics yield robust persistent representations

Joel Zylberberg[1] JOELZY@UW.EDU
Robert Hyde[2] RAH38@CASE.EDU
Ben Strowbridge[2] BENS@CASE.EDU

[1]University of Washington
[2]Case Western Reserve University

To make good decisions, animals maintain memories of the recent past, and use them to guide behavior. How are memories robustly represented, despite the variability in neuronal responses, and those responses' evolution over time? For the formation of persistent ("mnemonic") representations, noise is problematic because long response durations allow ample time for noise to accumulate. To understand robust memory function, we used a recently-developed hippocampal slice preparation that generates persistent responses to transient electrical stimulation. For $>$20s after stimulation, responses are indicative of the stimulus that was applied, even though the responses exhibit different dynamical trajectories on each trial. Using the Isomap manifold identification algorithm, we found that, following stimulation, the neural responses converge onto nearly 1-D manifolds, with a different "mnemonic manifold" associated with each stimulus. The neural responses' dynamical evolution, and trial-to-trial variability, shift the responses along these manifolds rather than between them, reducing the impact of those sources of variability on the mnemonic representations. The observed configuration of neural responses allows for the representation to be read-out (by downstream neural structures) at any time post-stimulation: by estimating which manifold the response lies on, the applied stimulus can be deciphered. For contrast, behavioral experiments frequently record activities during delay periods of fixed duration, thereby failing to address the question of "any-time-addressable" persistent representations.

III-7. The successor representation as a mechanism for model-based reinforcement learning

Evan Russek EMR443@NYU.EDU
Nathaniel Daw DAW@CNS.NYU.EDU

New York University

Behavioral and neural data suggest that animals and humans learn to make decisions by a combination of model-based and model-free reinforcement learning (RL) methods. Whereas a prominent hypothesis suggests that model-free temporal-difference (TD) learning is supported by dopaminergic projections to striatum, the neural mechanisms supporting model-based decision making are largely mysterious. A puzzle in this regard is that model-based algorithms, as typically envisioned, require structurally quite different computations than those used for TD learning: for instance, model-based learning does not make use of a TD prediction error of the sort associated with dopamine. However, rodent lesion studies suggest that model-based and model-free decisions are carried out by adjacent cortico-striatal loops that are relatively homologous in structure and receive similar dopaminergic innervation. We suggest that this seeming paradox could be resolved if model-based RL were accomplished by a TD update operating over a different cortical input. In particular, we study versions of the

successor representation (SR), a predictive state representation which, when combined with a TD reward learning stage mapping these states to values, can produce behavior analogous to model-based learning in some reward revaluation tasks. We show that because the original SR prelearns and caches the future state occupancy matrix, it fails to capture rodent and human model-based behavior in tasks that require subjects to adjust to changes in the transition matrix. However, we introduce a variant that minimizes such precomputation, so that it caches only the agent's action selection policy and not the resulting expected state occupancies. Simulations demonstrate that this more flexible SR strategy also has behavioral limitations that should in principle be detectable; however such limitations have not been exploited by previous experiments. Overall, this approach provides a neurally plausible approximation to model-based learning.

III-8. Back-propagation of prediction error signals in the basal forebrain underlies new learning

Hachi Manzur[1] HEMANZUR@NIH.GOV
Ksenia Vlasov[1] KSENIA.VLASOV@NIH.GOV
Shih-Chieh Lin[1,2] SHIH-CHIEH.LIN@NIH.GOV

[1]National Institutes on Aging
[2]NIH

Prediction error signals have long been proposed in reinforcement learning theories as a key component in driving new learning. Although it is generally assumed that, during learning, prediction error signals should back-propagate from the outcome to the cue predicting that outcome, such back-propagation has rarely been observed experimentally. Here we present electrophysiological evidence for a prediction error signal in the basal forebrain (BF) that back-propagates from the moment of reward to the moment of the reward-predictive cue across consecutive days of training. By carefully constructing a task where rats were required to learn a new stimulus outcome association, we were able to dissect the new learning process into at least three distinct phases involving different exploration and exploitation strategies. The behavioral responses in these distinct phases of new learning were tightly coupled with the activity of a population of noncholinergic BF neurons, respectively corresponding to different stages of prediction error back-propagation. Together, our observations provide strong experimental support for a key role of prediction error signals in the BF during new learning as agents refine their internal models of the environment.

III-9. Learning associations with a neurally-computed global novelty signal

Mohammadjavad Faraji[1] MOHAMMADJAVAD.FARAJI@EPFL.CH
Kerstin Preuschoff[2] KERSTIN.PREUSCHOFF@UNIGE.CH
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH

[1]Ecole Polytechnique Federale de Lausanne
[2]University of Geneva

Being able to detect novel stimuli is necessary for efficiently learning new memories without altering past useful memories (the challenge of the stability-plasticity dilemma). Furthermore, it is critical to consider novelty information (related to the structure of the environment) in addition to reward information in models of reinforcement-based learning. As such, novelty is a crucial factor in both learning and memory formation. However, how it is computed in a neurally plausible way and how it affects (synaptic) learning rules are both questions of debate. Here we propose a model to measure novelty based on the activity of decision units in a decision making process in neural networks. We also propose multiple ways by which novelty can implicitly modulate (gate) Hebbian plasticity to control learning at underlying synapses. We apply our model to a clustering task, in which the total number of clusters is initially unknown to the network. The proposed model is able to add more clusters whenever it judges

an input pattern to be novel, i.e., a pattern which may belong to none of the existing clusters. This model represents an agent that is able to generate (trigger) new states, an essential feature for learning new environments. We argue that the novelty signal in this framework can be interpreted as a (global) modulatory signal, corresponding to the diffusion of a non-specific neuromodulator (e.g., norepinephrine (NE) released from locus coeruleus (LC) neurons) and can modulate the local Hebbian factors (i.e., the coactivity of pre- and post-synaptic neurons) in synaptic plasticity rules. As such, it can be considered as a biologically plausible third factor in multi-factor learning rules.

III-10. It's easier to learn with a bad memory: fast inference in a simple latent feature model

Michael A Buice[1] MICHAELBU@ALLENINSTITUTE.ORG
Peter Latham[2] PEL@GATSBY.UCL.AC.UK

[1]Allen Institute for Brain Science
[2]Gatsby Computational Neuroscience Unit, UCL

The brain is faced with a very hard unsupervised learning problem: it has to make sense of the world just by looking at incoming spike trains from sensory receptors. The classic approach to this problem, pioneered by Olshausen and Field (1996), is to assume a generative model, and learn the parameters of that model. But what generative model should one assume? Most work has considered models based on continuous parametrization. For instance, Olshausen and Field assumed that images consisted of a linear sum of patterns, with the coefficients in the sum taking arbitrary values. While this is a good description in some regimes, the world is, in fact, divided into discrete objects: houses, cows, trees, cars, etc. As a first attempt at learning this kind of structure, we consider a model in which the world consists of a linear sum of discrete patterns, where the coefficients in the sum are either zero or one. Based only on noisy images consisting of multiple patterns, the goal is to learn what the underlying patterns are. Exact inference is impossible, so we take a variational approach: we approximate the true distribution over patterns with an approximate one, and minimize the Kullback-Leibler distance to the true distribution. The resulting update rules for the patterns are similar to those found in Olshausen-Field or other sparse coding approaches. They thus map naturally onto a neural network architecture. The resulting network can learn to dissociate multiple patterns with a small number of trials (as few as 50 trials for 4 patterns). However, after even a small number of trials, the posterior narrows, and the patterns become "frozen," making it impossible to learn new patterns should the environment change. An online decay can fix this problem with only a small effect on performance.

III-11. Evidence accumulation in a changing environment

Zachary Kilpatrick ZPKILPAT@MATH.UH.EDU
Alan Veliz-Cuba ALANAVC@MATH.UH.EDU
Kresimir Josic JOSIC@MATH.UH.EDU

University of Houston

The two-alternative forced-choice paradigm is commonly used to study sensory processing and decision making in humans and other animals. In the classic case, the correct choice does not change throughout a trial. However, natural environments change continuously. We therefore examined how an optimal observer makes an inference when the correct option changes in time. To do so, we derive an optimal procedure for accumulating evidence and deciding between two choices in a changing environment. Following the standard approach, we assume the observer makes sequential measurements, and uses them to update the likelihood ratio between the two choices. When the correct choice remains constant and observations are independent, this likelihood ratio (LR) can be computed as the product of the LRs from each observation. This is no longer true if the correct choice varies.

We provide a modified recursive expression for the LR that shows how an ideal observer discounts (or forgets) previous evidence. Taking the log LR and passing to the continuum limit, we find the evidence accumulation process is described by a nonlinear stochastic differential equation with a well defined equilibrium distribution. In contrast, when the environment does not change, in the continuum limit we obtain a linear drift-diffusion equation and no equilibrium distribution. Furthermore, non-dimensionalization allows us to describe evidence accumulation with a single parameter 'm', the information gained over the expected time between switches in ground truth. As 'm' increases, the asymptotic error rate of the model decreases. Notably, our model generalizes previously derived optimal procedures for change detection.

III-12. A model of perceptual learning: from neuromodulation to improved performance

Raphael Holca-Lamarre[1] RAPHAEL@BCCN-BERLIN.DE
Klaus Obermayer[2] OBY@NI.TU-BERLIN.DE
Joerg Luecke[3] JOERG.LUECKE@UNI-OLDENBURG.DE

[1]Bernstein Center for Computational Neuroscience, Berlin
[2]Technical University of Berlin
[3]University of Oldenburg

Experimental findings suggest that improvements in perceptual sensitivity are due to specific changes in sensory representations, for instance to refinements of individual receptive fields and to increases in the cortical area devoted to task-relevant stimuli. Further evidence indicates that neuromodulators, such as acetylcholine (ACh), orchestrate these changes. So far, scientific studies focused on the qualitative description and interpretation of experimental results. In this work, we address these results using a computational model, allowing us to draw a quantitative relationship between neuromodulation, changes in sensory representations, and measures of perceptual performance. For the model, we find that stimulus-specific disinhibition, such as attributed to ACh, results in an increase in the number of neurons encoding the specific stimulus. Such changes in neural representations lead to improvements in performance that are akin to improved detection: the hit rate increases at the expense of increases in false alarms. Furthermore, if disinhibition is coupled with top-down feedback encoding classification decision, neural representations become more discriminative for the task at hand. In this case, the improvements in performance are akin to improved discrimination: the hit rate increases while the false alarm rate drops. Interestingly, in this case, improvements in performance do not require a change in the number of neurons encoding a specific stimulus, in line with recently published observations. Our model points to a possible mechanism by which neuromodulation, either by itself or paired with top-down feedback, may improve perceptual sensitivity. Additionally, our findings suggest that these two different scenarios lead to distinct changes in both neural representations and perceptual sensitivity, thereby making predictions for future animal experiments.

III-13. Neural mechanisms of rule-based behavior

Flora Bouchacourt FLORA.BOUCHACOURT@GMAIL.COM
Etienne Koechlin ETIENNE.KOECHLIN@UPMC.FR
Srdjan Ostojic SRDJAN.OSTOJIC@ENS.FR

Ecole Normale Superieure

The human brain is able to adjust and exploit multiple strategies for a same task, depending on behavioral demands. The representations of such stimuli-response mapping rules are called task sets. Most of the theoretical research on rule-based behavior is based on computational models at the level of behavior. Little is known however about its neural implementation and mechanisms. We examine a candidate mechanism for neural implementation of task sets by means of synaptic plasticity. Our model is composed of two interacting neural circuits.

The associative network learns one to one associations between visual stimuli and motor responses, but cannot learn more than one stimuli-response mapping. The task rule network learns the representations of those mappings through hebbian and temporal sequence learning mechanisms. Task sets are encoded in its pattern of synaptic connectivity and a feedback to the associative network enables their retrieval. We first implement a rule-independent associative network. Fitting the model to behavioral data, we find that it can account for behavior in the session in which 24 different task sets are presented one after the other. In contrast, it poorly describes the behavior when only 3 task sets are presented repeatedly across the whole session. Introducing the task rule network permits to account for the data. Hence we show the importance of its activity and of its feedback to the associative network for the retrieval of a previously seen rule. Then we describe the effects of progressive learning through synaptic plasticity in the task rule network. Our model explores a mechanism for neural implementation of learning, acquisition and activation of rules towards action, at the boundary between functional and neuronal levels.

III-14. Single trial electrophysiological dynamics predict behaviour during sedation

Laura Lewis[1] LAURALEWIS@FAS.HARVARD.EDU
Robert Peterfreund[1] RPETERFREUND@MGH.HARVARD.EDU
Linda Aglio[2] LAGLIO@PARTNERS.ORG
Emad Eskandar[1] EESKANDAR@PARTNERS.ORG
Lisa Barrett[3] L.BARRETT@NEU.EDU
Sydney Cash[1] SCASH@MGH.HARVARD.EDU
Emery Brown[1] ENB@NEUROSTAT.MIT.EDU
Patrick Purdon[1] PATRICKP@NMR.MGH.HARVARD.EDU

[1]Harvard University
[2]Brigham and Women's Hospital
[3]Northeastern University

Low doses of anesthetic drugs can induce a state of sedation, during which patients drift in and out of consciousness. These fluctuations in behavioural state occur on second-to-second timescales, as patients may respond to sensory stimuli at one moment and not the next. We sought to identify the neural dynamics that differ between the awake and unresponsive states during sedation. We recorded intracranial electrocorticography (ECoG) in 11 patients with intractable epilepsy undergoing propofol anesthesia for clinical reasons. Patients received an infusion of propofol over 14 minutes, resulting in a gradual transition between the awake, sedated, and anesthetized states, while responding to auditory stimuli with button presses. We found that patients spent a median of 4.2 minutes in a state of sedation, reflected by increased reaction times and decreased probability of responding to stimuli. We analyzed evoked responses to the auditory stimuli during sedation and found that the amplitude of slow ($>$80 ms) components was reduced when patients failed to respond, possibly indicating a breakdown of high level cortical processing. To determine which dynamics caused this breakdown, we tested whether single trial dynamics prior to stimulus onset could predict whether patients would respond to sensory stimuli. We found that L1-regularized generalized linear models could successfully predict behavioural state using the spectral content of the ECoG recordings. In particular, the appearance of isolated slow ($<$1 Hz) waves in frontal channels correlated with the failure to respond to a stimulus. Previous reports have shown that slow waves are associated with hundreds of milliseconds of silence in cortical neurons, suggesting that this pattern corresponds to periods of neuronal suppression in frontal cortex. We conclude that single trial fluctuations in cortical dynamics can predict arousal state, and suggest that suppression of frontal cortex may disrupt top-down coordination of behavioural responses to incoming stimuli.

III-15. Humans maintain probabilistic belief states when predicting object trajectories

Matjaz Jogan[1] MJOGAN@SAS.UPENN.EDU
Alan He[1] AHE@SEAS.UPENN.EDU
Alexander Tank[2] ALEXTANK@UW.EDU
Alan A Stocker[1] ASTOCKER@SAS.UPENN.EDU

[1]University of Pennsylvania
[2]University of Washington

The ability to accurately predict the trajectories of moving objects is crucial for an autonomous system that interacts with a dynamic environment. Humans are thought of having a sense of so-called “intuitive physics” that allows them to be quite efficient in making such predictions. Recent experimental results suggest that these predictions reflect the outcome of a probabilistic inference process based on noisy observations and an accurate physics model of the world (Smith et al., 2013; Battaglia et al., 2013). However, it remains unknown whether humans mentally track and update i) an estimate or ii) a full probabilistic description of the object state (belief state) (Lee et al., 2014). We designed a set of psychophysical experiments to specifically distinguish the two hypotheses. Subjects were first asked to predict the collision location of a moving object with a hidden wall. The trajectory of the object was occluded and subjects were only given the object’s initial motion and an acoustic signal at the precise time of collision. Subjects exhibited clear biases in their location estimates that indicated that they were performing probabilistic inference using prior expectations over speed and location. Subjects then repeated the experiment receiving, however, an additional spatial cue about the hidden wall location. By introducing different levels of uncertainty associated with this cue we expected subjects to assign different relative weights in combining the cues if they were maintaining full belief states while tracking. More specifically, by measuring subjects performance for each cue alone we were able to individually predict optimal behavior and verify whether it matched subjects’ actual behavior. We found that subjects’ behavior was indeed well predicted by a Bayesian belief state model that optimally combined cues across space, time, and object motion. Our results suggest that humans maintain and update full belief states when predicting object trajectories.

III-16. Time course of attentional modulations in primary visual cortex

Xiaobing Li[1] XLI@SUNYOPT.EDU
Michael Jansen[1] MJANSEN@SUNYOPT.EDU
Reza Lashgari[2] REZALASHGARI@GMAIL.COM
Yulia Bereshpolova[3] YULBER@GMAIL.COM
Harvey Swadlow[3] HARVEY.SWADLOW@UCONN.EDU
Jose-Manuel Alonso[1] JALONSO@SUNYOPT.EDU

[1]SUNY College of Optometry
[2]Iran University of Science and Technology
[3]University of Connecticut

Spatial attention is known to modulate the activity of single neurons in primary visual cortex (area V1) but it remains unclear how these modulations change over time. To measure the time-course of V1 attentional modulations, we trained two rhesus monkeys to pay attention for 1-3 seconds to one of multiple drifting gratings cued at the beginning of each trial. The monkeys were rewarded to maintain their attention for the entire trial and release a lever as fast as possible after detecting a brief change in either grating color or orientation. The spike trains of each neuron were converted into a continuous spike density function to calculate mean rate (F0) and response power within a low frequency range (LF, 1-7.5 Hz) that contained the stimulus frequency (2-4 Hz). Attentional modulations were measured as a percentage of response change when attention was cued inside versus outside of the receptive field and the attentional time course was calculated by sliding a temporal window of 500 msec (10 msec steps) over the peri-stimulus time histogram (PSTH, n= 61 neurons, average 891 trials / neuron) and

fit with a Weibull function (mean R2 for 2-sec trials: 0.939 for F0 vs. 0.761 for LF). Our results demonstrate that spatial attention increases V1 responses at a rate of 20% per 0.1 seconds and the increase starts decaying 1.5-2 seconds after the stimulus onset, even if both monkeys respond significantly faster to long than short trials (3s trials vs 1s trials: Monkey R, 0.267s vs 0.287s, $p < 0.001$; Monkey S, 0.291s vs 0.315s, p=0.006, t-test). We conclude that the effect of spatial attention on V1 neuronal responses peaks at the median duration of the trial set, however, the benefits of attention on behavioral performance keep increasing with time and are greatest for the longest trials.

III-17. Prediction predicated on prediction error: Boosting the anticipation of delayed rewards

Kiyohito Iigaya[1] KIIGAYA@GATSBY.UCL.AC.UK
Giles W Story[2] G.STORY@UCL.AC.UK
Zeb Kurth-Nelson[2] ZEBKURTHNELSON@GMAIL.COM
Ray Dolan[2] R.DOLAN@UCL.AC.UK
Peter Dayan[2] DAYAN@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University College London

From the gender of foetuses to the forthcoming weather, people often prefer to know about uncertain future outcomes in advance. Experiments suggest that stochastic rewards are more attractive when their fate is revealed (by a predictive cue) sooner rather than later — even, suboptimally, at the expense of mean return [Spetch et al. (1990), Gipson et al. (2009)]. Attempts have been made to understand the neural and computational bases of this so-called 'observing' behaviour. Bromberg-Martin and Hikosaka (2009) showed that the activities of midbrain dopamine neurons directly reflect the preference for observing; later (2011) they showed that the activ- ities of lateral habenula neurons preclude Beierholm and Dayan (2010)'s account based on Pavlovian misbehaviour. Instead, they advocated a direct attraction for 'information' inherent in observing, a the- ory that unfortunately founders on the observation of increased preference for decreased (and thus less entropic) reward probabilities [Roper and Zentall (1999)]. We reconcile these issues via [Loewenstein (1987)]'s proposal of a direct benefit for the anticipation of future reward (along with that of the reward itself). Crucially, we hypothesise that reward prediction errors inspired by the predictive cue can boost anticipation (consistent with the dramatically enhanced excitement that follows the cue; Spetch et al. (1990)). As a consequence, the values of predictive cues can be enhanced by the delay until the reward delivery. Our model provides a ready explanation for both behavioural and neurophysiological findings in observing, including the effects of changing the delay preceding the reward [Spetch et al. (1990)]. We are currently testing it using human psychophysics experiments. Our study suggests a new in- terpretation of reward prediction error signals and reward seeking behaviours, of special relevance to gambling.

III-18. Metaplasticity and choice under uncertainty

Shiva Farashahi[1] SHIVA.GHAANIFARASHAHI.GR@DARTMOUTH.EDU
Hyojung Seo[2] HYOJUNG.SEO@YALE.EDU
Daeyeol Lee[2] DAEYEOL.LEE@YALE.EDU
Alireza Soltani[1] ALIREZA.SOLTANI@DARTMOUTH.EDU

[1]Dartmouth College
[2]Yale University

Much effort in system neuroscience has been focused on studying the effects of reward on decision making. However, we still do not fully understand how reward is integrated and influences choice in real life situations where reward-predicting cues and reward contingencies can change frequently and unpredictably. Here we use

a combination of experimental and modeling approaches to study how reward is integrated over time depending on the uncertainty of reward contingencies in the environment. The experiment is a variation of probabilistic reversal learning task in which we altered both reward probability and the block lengths in order to create two environments that require very different time constants of reward integration. Firstly, our experimental results indicated that subjects can adjust the time constant of their reward integration according to the uncertainty of reward contingencies in each environment. Secondly, we developed a decision-making model endowed with different reward-dependent synaptic plasticity rules to study how optimal reward integration can be achieved at the synaptic level. In our model, synapses could have multiple meta-levels associated with each of the two levels of synaptic strength (weak and strong). We found that whereas a single set of synapses with no metaplasticity enabled the model to perform optimally in a given environment, contributions from synapses with non-optimal learning rates strongly reduced the performance. In contrast, metaplasticity resulted in close to optimal behavior as long as the slowest or fastest transition probabilities include the optimal learning rates in each environment, and model's choice behavior was improved with larger numbers of meta-levels. Finally, we tested the robustness of metaplasticity as a neural mechanism for choice under uncertainty and found that an ensemble of synapses with a wide of range of transition probabilities could provide inputs for optimal choice behavior in very different environments.

III-19. Dynamic correlations between visual and decision areas during perceptual decision-making

Il Memming Park[1,2] MEMMING@AUSTIN.UTEXAS.EDU
Jacob Yates[2] JLYATES@UTEXAS.EDU
Alex Huk[2] HUK@UTEXAS.EDU
Jonathan W Pillow[3] PILLOW@PRINCETON.EDU

[1]Stony Brook University
[2]University of Texas at Austin
[3]Princeton University

The lateral intraparietal area (LIP) reflects the accumulation of motion-related evidence from the middle temporal area (MT) during decisions about motion. However, little is known about the functional connections between MT and LIP. We address this question by analyzing the local field potential (LFP) and spikes recorded simultaneously in MT and LIP during a motion-discrimination task. Monkeys viewed a motion stimulus (presented within the receptive fields of MT neurons), made a decision about which direction it moved, and communicated their choice with a saccade to one of two choice targets (one of which was inside the response fields of LIP neurons). We observed two prominent frequency bands in the LFP of both areas: delta, 2-5 Hz; and alpha, 8-15 Hz. In LIP, the alpha-band power exhibited choice-selective ramping, a signature of evidence integration, and was suppressed prior to saccades, similar to movement-related desynchronization found in motor cortex. Although we expected to observe strong correlations between MT and LIP during the motion-viewing period—when LIP presumably integrates spiking activity from MT—we observed the opposite pattern: the spike-field and field-field coherences between MT and LIP were negligible during decision formation, and were strong only during the inter-trial interval and portions of the trial outside of the motion-viewing period. We obtained similar results in cross-area spike correlation analysis. We hypothesize that both MT and LIP fall into an oscillatory low-dimensional dynamical network during idle states or simple motor planning and execution, but desynchronize during key portions of specific sensory processing and perceptual decision formation. Such within- and between-area desynchronization is consistent with recent results in sensory cortices, but challenge notions that frequency-specific signalling is a critical element of cross-area communication.

III-20. Feedback-related information representation in human medial and lateral prefrontal cortex

Elliot Smith[1] EHS2149@COLUMBIA.EDU
Garrett Banks[1] GPB2111@CUMC.COLUMBIA.EDU
Charles MIkell[1] CBM2104@COLUMBIA.EDU
Shaun Patel[2] SHAUN.PATEL@ME.COM
Emad Eskandar[2] EESKANDAR@PARTNERS.ORG
Sameer Sheth[1] SS4451@COLUMBIA.EDU

[1]Columbia University
[2]Harvard University

Feedback related negativity (FRN) is commonly observed in electroencephalography studies as the difference between potentials evoked by positive and negative feedback about a subject's behavioral performance. The dorsal anterior cingulate cortex (dACC) is thought to be the cortical generator of this signal, as well as the generator of a larger class of cognitive control signals known as frontal midline theta ($FM\Theta$). There is still debate, however, about the origin and significance of the FRN and its relationship to $FM\Theta$ signals. Here we examine direct intracranial recordings from the dACC and lateral prefrontal cortex in humans to better understand the anatomical localization of the FRN and provide insight into its cortical computation at the level of local neuronal populations. Seven human subjects undergoing epilepsy monitoring with intracranial electrodes spanning the mediolateral extent of prefrontal cortex received feedback about their behavioral performance in a Stroop-like task. We show that the FRN is evident in both low and high frequency local field potentials recorded on electrocorticography. However, FRN is larger on medial than lateral contacts for both frequency ranges, and coupling between theta phase of the low frequency LFPs and high frequency LFP amplitude is also greater in medial contacts. We also provide evidence that medial and lateral contacts are functionally connected using Granger causality and conditional mutual information analyses, and that information transfer increases during the feedback period from medial to lateral contacts. Furthermore, the medial to lateral information transfer oscillates with theta-range periodicity. These results provide evidence for the dACC as the cortical source of the FRN, and support the idea that FRN and $FM\Theta$ signals are generated by an architecture in which information from dACC is transferred to lateral PFC using theta modulation. These results have implications for the neurophysiology of reinforcement learning in cognitive control and decision-making.

III-21. A homeostatic reinforcement learning theory of cocaine addiction

Mehdi Keramati[1] M.MAHDI.KERAMATI@GMAIL.COM
Boris Gutkin BORIS.GUTKIN@GMAIL.COM
Serge H Ahmed[2,3] SERGE.AHMED@U-BORDEAUX.FR

[1]Gatsby Computational Neuroscience Unit, UCL
[2]Universite de Bordeaux
[3]IMN UMR CNRS 5293

Several animal models have demonstrated the important role of learning processes in addiction-like behaviours (e.g., instrumental drug-seeking, cue induced relapse, etc.). Thus, a fruitful theoretical approach to understand addiction has been to see it as the brain reward-learning system having gone awry. Other experimental facts (e.g., dose-response and dose-intake curves, loading effect at the beginning of drug self-administration sessions, etc.), on the other hand, have proven the importance of homeostatic regulatory processes in addiction. Following this approach, a second class of theoretical models has depicted addiction as the needs structure being altered by chronic drug-use. Built upon recently proposed 'homeostatic reinforcement learning' theory, we propose a unified theory of cocaine addiction showing how both homeostatic and learning process are critical in the development of addiction. The theory explains key experimental results that none of the two classical approaches can account for separately. We argue that the relative excitability of direct- vs. indirect cortico-basal ganglia pathways by glutamatergic signals is a homeostatically-regulated internal variable. This variable is modulated by striatal dopamine concentration, which in turn is modulated by brain cocaine level. We also assume that chronic cocaine use changes the ideal setpoint of this variable through down-regulating the availability of dopamine D2 receptors. We simulate the model and show that based on these two assumptions, our model explains a wide range of experimental signatures of addiction, notably escalation of cocaine self-administration under long, but not short daily access, loading effect, pre- and post-escalation dose-response curves, as well as dose-intake curves, the effect of post-escalation reduced availability of cocaine on response rate, the effect of session-duration of escalation pattern, post-escalation insensitivity to punishment, priming-induced relapse of escalated cocaine seeking, and the effect of D2 receptor availability on motivation of cocaine, as well as on vulnerability to develop addiction.

III-22. Amygdala activity induced by the inequity predicts long-term change of depression index

Toshiko Tanaka[1] TOSHIKO64@GMAIL.COM
Takao Yamamoto[2] YAMAMOTO.T-AY@NHK.OR.JP
Masahiko Haruno[3] MHARUNO@NICT.GO.JP

[1]Osaka University
[2]NHK
[3]NICT Cinet

Depression is a big problem for the society by which one's personal and economic life deteriorates seriously. Although previous cohort-based studies have implicated a causal relationship between social stresses such as inequity and depression (Wilkinson and Pickett, 2006), little is known about neural substrates for the link, partially due to substantial individual differences in sensitivity. Here, we address whether fMRI brain response induced by inequity can predict the long-term (one-year) change of a depression index (i.e., BDI-II test). We first found that prosocial (n=174) and individualistic (n=59) participants (c.f., Haruno and Frith, 2010) exhibited different distributions of BDI score (p=0.0045) and hypothesized that their differential fMRI activity in correlation with the inequity (i.e., absolute payoff difference) may foresee their depressive tendency. Fifty prosocials and twenty-seven individualists were scanned as a responder in the ultimatum game. We found significant differential in the bilateral amygdala, inferior frontal gyrus (IFG), and medial frontal gyrus (MFG) ($p<0.001$; uncorrected). We also asked the participants to answer the BDI-II test again one year later and computed the BDI change for each participant. To predict the BDI change based on brain activity, we applied a kernel-based (spline) Bayesian regression (Tipping 2001) to beta values (for inequity) in the amygdala, IFG and MFG. Each ROI was defined as a 5mm-radius sphere around the peak voxel. Our one-leave-out test demonstrated a significant positive relationship between the BDI change and predicted values for amygdala (R=0.42, p=0.00015) but for IFG (R=0.10, p=0.39) and MFG (R=0.16, p=0.15). Furthermore, amygdala did not show such relationship for the presentation of a proposer's faces (R=0.028, p=0.25) and offers (R=0.047, p=0.25). These results revealed the first biological link between inequity and depressive tendency and a key role of the amygdala, and may suggest a crucial effect of social comparison on our mental state.

III-23. Voluntary behavioral engagement controls cortical spike timing for enhanced auditory perception

Ioana Carcea IOANA.CARCEA@MED.NYU.EDU
Michele Insanally MNI1@NYU.EDU
Robert Froemke ROBERT.FROEMKE@MED.NYU.EDU

New York University

Voluntary engagement in behavioral tasks enhances performance. Brain states during active behavior can improve perceptual abilities by both increasing detection of sensory input in a background of ongoing activity, and by facilitating recognition of behaviorally relevant inputs over less relevant or distracting inputs (Jaramillo and Zador, 2011; Raposo et al., 2012; Brunton et al., 2013). However, the neuronal mechanisms and computational principles by which voluntary engagement improves stimulus detection and recognition remain poorly understood. We made single-unit recordings from primary auditory cortex of adult rats performing an appetitive auditory go/no-go task (Froemke et al., 2013). Animals self-initiated some of trials ('Self' trials) while other trials were externally triggered ('External' trials). Stimulus detection and stimulus recognition were superior during self-initiated trials, but surprisingly, this improvement was apparently not supported by changes of cortical spike rates or receptive fields. Instead, stimulus detection and recognition were enhanced by changes in spike timing precision over two different epochs. For detection, ongoing activity between self-initiation and stimulus onset showed a progressive decrease in spike timing variability and enhanced information about the upcoming stimulus prior to stimulus onset. In contrast, stimulus recognition was enhanced by changes in spike timing only during stimulus presentation. Intriguingly, if the stimulus occurred earlier than expected, detection was selectively impaired, but if the stimulus occurred later than expected, recognition was specifically disrupted instead. These results demonstrate that voluntary behavioral engagement prepares cortical circuits for sensory processing by two distinct processes during different temporal domains.

III-24. A model-based learning rule that predicts optimal control policy updates in complex motor tasks

Anastasia Sylaidi A.SYLAIDI10@IMPERIAL.AC.UK
Aldo Faisal A.FAISAL@IMPERIAL.AC.UK

Imperial College London

Optimal feedback control (OFC) constitutes one of the dominant computational paradigms that bridges perception and action on account of optimality criteria (Todorov, E., Jordan, M., 2002, Scott, S., 2004, Wolpert, D., Landy, M., 2012). However, the mechanisms that underlie the selection and update of optimal controllers remain unclear, since motor neuroscience has primarily focused to date on capturing learning as an adaptation of task-related parameters. Here, we address this challenge with an OFC framework that explains motor learning in complex motor tasks as a simultaneous update of task representations and action policies. Our hypothesis is that the brain uses a locally linear system identification process to continuously update internal composite representations of body and world dynamics, which in turn determine motor commands, that drive movement and optimize a set of task objectives. The difference between predicted and actually produced movement is observed via perceptual pathways and constitutes a prediction error that propels the learning process through gradient descent steps in the space of internal forward model parameters. We tested this model in an experimental paradigm that instructed subjects to move a virtual object of unknown unintuitive dynamics from start to target locations. After each trial completion subjects received performance feedback estimated by an arbitrary quadratic cost function that captured the instructed task goals. Our model consistently predicted the trial-by-trial progression of motor learning on an individual subject basis. Crucially, it also performed better in capturing end-performance of subjects than an ideal-actor model, which assumes complete knowledge of task dynamics. Our results suggest that the brain can rely on simple strategies of local system identification to learn the near optimal control of complex object manipulation tasks. The proposed framework provides thereby an algorithmic formalization, which can guide further experimental investigations on the neural foundation of cortical action selection and motor learning.

III-25. Context and action inputs in a reinforcement learning model of bird-song: modeling and ultrastructure

Joergen Kornfeld[1] JOERGEN.KORNFELD@GMAIL.COM
Michael Stetner[2] MICHAEL.STETNER@GMAIL.COM
Winfried Denk[1] WINFRIED.DENK@MPIMF-HEIDELBERG.MPG.DE
Michale Fee[2] FEE@MIT.EDU

[1]Max-Planck-Institute for Medical Research
[2]Massachusetts Institute of Technology

Reinforcement learning (RL) associates three signals to optimize behavior: context, action, and reward. These signals are thought to converge in the basal ganglia, often associated with reinforcement learning and highly conserved across vertebrates (Reiner, 2009). Songbirds could follow this algorithm to learn which vocalization (action) to make at each moment in time (context) to maximize the similarity to their tutor's song (reward). Previous research suggests that the songbird basal ganglia receives input from premotor nucleus HVC, which carries timing signals (context), and premotor nucleus LMAN, which generates exploratory vocal variability during learning (action). This led to a new conceptual model (Fee and Goldberg 2011) for RL with distinct functional roles for these two inputs: (1) HVC timing inputs are selectively strengthened during learning, and (2) LMAN action inputs gate plasticity at the HVC inputs, but are not plastic themselves. We demonstrated with numerical simulations that this model can indeed learn to imitate a template song. We next examined whether the functional distinction between HVC and LMAN inputs are reflected in the nature of their synaptic contacts onto striatal medium spiny neurons (MSNs). We hypothesized that the plastic HVC inputs would synapse onto MSN dendritic spines, thought to be a key site of striatal plasticity (Kreitzer and Malenka, 2008). HVC and LMAN axons were reconstructed from a serial block-face electron microscope image stack (Denk & Horstmann 2004) of Area X and were classified based on known morphological features. We identified all synapses formed by putative LMAN and HVC axons and classified them as dendritic spine synapses or dendritic shaft synapses. As predicted, we found that HVC axons form the presynaptic partner of the vast majority of MSN spines ($>$90%). Furthermore, we found that, while HVC axons preferentially terminate on MSN spines, LMAN axons more often terminate on MSN shafts.

III-26. Neural responses in macaque caudate nucleus during visual perceptual learning

Takahiro Doi TAKADOI@MAIL.MED.UPENN.EDU
Joshua Gold JIGOLD@MAIL.MED.UPENN.EDU

University of Pennsylvania

Our ability to detect, discriminate, and identify visual stimuli can improve greatly with training. We study neural mechanisms responsible for such improvements in monkeys learning to perform a visual motion direction-discrimination task. Our previous work showed that training shapes how sensory information is read out to form the decision, possibly via a reinforcement-driven learning process. Here we tested the hypothesis that key features of such a reinforcement-driven process, including signals related to prediction errors that scale with stimulus strength, are represented in neural activity in the basal ganglia during perceptual learning on the task. We recorded single- and multi-unit activity from caudate nucleus, the input stage of basal ganglia, while a monkey was trained to perform a random-dot direction discrimination task. We presented a patch of noisy moving dots for a fixed duration at the center of the visual field. After a variable delay, the monkey indicated the perceived direction by making a saccade to one of the two choice targets. After the monkey established the correct motion-saccade association with an easy stimulus, we gradually introduced weaker motion stimuli (stimuli with a larger proportion of randomly moving dots). As expected, the monkey's sensitivity to weak motion stimuli improved over the course of training. We found that responses of caudate neurons also changed during training, in particular reflecting an increasing sensitivity to motion strength. The selectivity tended to have the same sign for the saccadic choices contralateral and ipsilateral to the recorded hemisphere. These responses are consistent with a representation of choice-independent value, which could, in principle, be used to guide a reinforcement-driven learning process. Thus, the basal ganglia may contribute to a long-term learning process that improves visual perception.

III-27. A deep learning theory of perceptual learning dynamics

Andrew Saxe ASAXE@STANFORD.EDU
Stanford University

With practice, humans and other organisms dramatically improve their accuracy in simple perceptual tasks. This perceptual learning has been the focus of extensive experimentation aimed at identifying its neural basis: what synaptic changes at the neural level underlie the behavioral improvements observed at the psychological level? Part of the difficulty stems from the brain's depth—its many-layered anatomical structure—which substantially complicates the learning process. Determining the locus of neural changes across the cortical hierarchy has been a key focus of empirical investigations, yet while a number of computational models have addressed changes within a single layer, none have explained the distribution and timing of changes across layers. This work develops a quantitative theory of perceptual learning based on gradient descent learning in deep linear neural networks. Despite their linearity, the learning problem in these networks remains nonconvex and exhibits rich nonlinear learning dynamics. The theory gives precise answers to fundamental theoretical questions such as the size and timing of changes across layers. Within a single layer, the theory's predictions coincide with earlier shallow models of perceptual learning: tuning changes target the 'most informative' neurons. Across layers, the theory predicts that changes follow a reverse hierarchy, with higher layers changing earlier, and ultimately more, than lower layers. Furthermore, these changes interact with task precision: coarse discriminations change only higher layers, while fine discriminations produce changes across the entire cortical hierarchy. Finally, the theory addresses the crucial issue of whether learning will transfer to new contexts, predicting that coarse discriminations transfer better than precision discriminations; and that early learning transfers better than late learning. These predictions accord with a diverse set of experimental findings, suggesting that the brain's depth is a key factor influencing the size and timing of receptive field changes in perceptual learning.

III-28. Inferring learning rules from distributions of firing rates in inferior temporal cortex

Sukbin Lim[1] SUKBIN@UCHICAGO.EDU
Jillian L. McKee[1] JMCKEE@UCHICAGO.EDU
Luke Woloszyn[2] LUKE.WOLOSZYN@GMAIL.COM
Yali Amit[1] AMIT@GALTON.UCHICAGO.EDU
David J Freedman[1] DFREEDMAN@UCHICAGO.EDU
David L. Sheinberg[3] DAVID_SHEINBERG@BROWN.EDU
Nicolas Brunel[1] NBRUNEL@GALTON.UCHICAGO.EDU

[1]University of Chicago
[2]Columbia University
[3]Brown University

Information about external stimuli is thought to be stored in cortical circuits through experience-dependent modifications of synaptic connectivity. These modifications of network connectivity should lead in turn to changes in the dynamics of the relevant circuits, as a particular stimulus is presented repeatedly. Here, we ask what plasticity rules are consistent with the differences in the statistics of the visual response to novel and familiar stimuli in inferior temporal cortex, an area underlying visual object recognition. Experimentally, it was shown that the average visual responses decrease with familiarity, while sparseness of the activity and selectivity to stimuli increase (Freedman et al., 2006; Woloszyn and Sheinberg, 2012). We introduce a method to infer the dependence of

the ‘learning rule’ on the post-synaptic firing rate, from comparing the distributions of visual responses to novel and familiar stimuli. This method provides both single neuron static transfer functions (f-I curves), and the simplest firing rate-based synaptic plasticity rule that is consistent with the observed differences between these two distributions. We applied this method to experimental data obtained in ITC neurons in monkeys performing two different tasks, a passive-fixation task (Woloszyn and Sheinberg, 2012), and a delayed match-to-category task. In putative excitatory neurons, the inferred learning rule exhibits depression for low post-synaptic rates, and potentiation for high rates. The threshold separating depression from potentiation is strongly correlated with both mean and standard deviation of the firing rate distribution. Changes in responses of putative inhibitory neurons between novel and familiar stimuli are consistent with an absence of plasticity in those neurons. Finally, we show that a network model implementing a rule extracted from data shows stable learning dynamics, and leads to a sparser representation of external inputs.

III-29. The effect of pooling in a deep learning model of perceptual learning

Rachel Lee RACHELSTEPHLEE@GMAIL.COM
Andrew Saxe ASAXE@STANFORD.EDU
Stanford University

Practice makes perfect: with extensive training, humans and monkeys reach much lower discrimination thresholds in fine visual orientation discrimination tasks. Many empirical studies have sought the neural basis of this perceptual learning, but widely different results have been reported even for paradigms that seem effectively identical. Here we consider one such example: Schoups et al. (2001) had macaques discriminate between oriented sine wave gratings, and found increases in orientation tuning curve slope among the most informative V1 neurons. By contrast, Ghose et al. (2002) found no such change despite also employing a fine orientation discrimination paradigm. The tasks used in these studies differ slightly: in the former, only the phase of the sine grating was randomized between trials; while in the latter, both phase and spatial frequency were varied. A vital challenge for theory is to explain why these apparently minor variations in paradigm could yield such divergent results in electrophysiology. In this work we show that this small task difference is in fact significant: we construct a four layer deep neural network model and show that it exhibits the same pattern of results, namely robust V1 changes when inputs have only random phases, and no noticeable V1 changes when inputs have both random phase and spatial frequency. We trace this behavior to the presence of complex cells that pool across phase, thereby also pooling error signals across examples when only phase is randomized. Our result provides a new theoretical interpretation of these experimental discrepancies as a natural consequence of gradient-based learning in a deep network with pooling-based invariance structure, and gestures to a possibly generalizable principle in perceptual learning: variation along a non-pooled stimulus dimension blocks learning in a given cortical area, even if this stimulus dimension is not relevant to the discrimination.

III-30. Mapping eye movement control circuits across the entire hindbrain

Alexandro Ramirez ADR2110@GMAIL.COM
Emre Aksay EMA2004@MED.CORNELL.EDU
Weill Cornell Medical College

The control of spontaneous eye movements requires the coordinated activity of at least two groups of premotor neurons: burst neurons that fire a brief volley of action potentials resulting in a saccade and Velocity-to-Position Neural Integrator (VPNI) cells that convert burst neuron activity into persistent firing required to hold the eyes fixed at their new location. The mechanisms underlying these firing patterns are thought to rely heavily on network interactions, however the comprehensive locations of VPNI and burst neurons has never been determined in a single species. We address this in the larval zebrafish (7-8 dpf) by coupling focal laser ablations with calcium imaging throughout the entire hindbrain and cerebellum while simultaneously tracking spontaneous eye movements using two-photon microscopy from populations of neurons expressing GCaMP6f. Eye movements were made in the dark to remove confounding effects on cell activity from visual input and resultant feedback. We find the majority of cells correlated with eye movements could be classified as position or velocity related. Velocity neurons were distributed in the rostral-caudal axis across most of the hindbrain (rhombomeres 2-8). Position cells were more clustered: in addition to the previously observed cells in the caudal hindbrain (rhombomeres 7-8), cells were also clustered in more rostral locations (rhombomeres 5-6 and 2). Interestingly, we find a buildup of activity multiple seconds before a saccade in some neurons suggesting a novel cell type involved in eye movement control. We find few neurons in the cerebellum correlated with spontaneous eye movements. Preliminary analysis of laser ablations suggest differential roles for the groups found and possibly inter-group coupling. These results provide the first comprehensive map of the distribution of signals important for the generation of spontaneous eye movements and begin to address the causal roles of different nuclei.

III-31. Extracting grid characteristics from spatially distributed place-cell inputs using non-negative PCA

Yedidyah Dordek YDORDEK@GMAIL.COM
Ron Meir RMEIR@EE.TECHNION.AC.IL
Dori Derdikman DERDIK@TECHNION.AC.IL

Technion - Israel Institute of Technology

The precise interaction between different spatially-dependent cells in the hippocampal formation is under debate. While many models study the downstream projection from grid cells to place cells, recent data has pointed out the importance of the feedback projection from place cells back to grid cells. Few models (e.g., Kropff & Treves, 2008) have studied the nature of this feedback. Here we continue this line of modeling and ask how grid cells are affected by the nature of the input from the place cells. We consider the following questions: How are grid cells formed? And what causes the emergence of their hexagonal shape? We propose a two-layered neural network with feedforward weights connecting place-like input cells to reciprocally connected grid cell outputs. The network is trained by moving a virtual agent randomly in a given space. Place-to-grid weights are learned via a general Hebbian rule while inter-grid weights develop by an anti-Hebbian rule. The architecture of this network highly resembles neural networks used to perform PCA (Principal Component Analysis). In accordance with the known excitatory nature of place-to-grid interactions, our results indicate that if the components of the feed-forward neural network (or the coefficients of the PCA) are enforced to be positive, the final output converges to a hexagonal lattice (Fig. 1). However, without the constraint of positivity we get mostly square lattice results. The emergence of the specific hexagonal shape is investigated in terms of stability and maximum variance by combining mathematical analysis and numerical simulations. For example, when the network is initialized with various spatially dependent inputs (e.g. square, stripe-like), it almost always converges to a hexagonal lattice (Fig. 2). Our results express a possible linkage between place-cell-grid-cell interactions to PCA, suggesting that grid cells represent a process of dimensionality reduction of the information in place cells.

III-32. Minimal ensemble coding of combined stimulus properties in the leech CNS

Jutta Kretzberg JUTTA.KRETZBERG@UNI-OLDENBURG.DE
Friederice Pirschel FRIEDERICE.PIRSCHEL@UNI-OLDENBURG.DE
Elham Fathiazar ELHAM.FATHIAZAR@UNI-OLDENBURG.DE

University of Oldenburg

Leeches are able to discriminate touch locations as precisely as the human fingertip, despite their very low number

of mechanosensory cells. Each patch of skin is innervated by only two P-cells (“pressure”) and two T-cells (“touch”) with overlapping receptive fields. In both cell types, touch stimuli located closer to the center of the receptive field trigger more spikes with shorter latency. However, increasing touch intensity also induces a higher spike count and shorter latency, leading to potential confusion of both stimulus properties. How can such a low number of neurons nevertheless trigger the precise responses observed behaviorally? We performed intracellular double recordings of sensory cells during tactile skin stimulation and analyzed the responses with stimulus estimation techniques. When touch location and touch intensity were varied separately, both stimulus properties could be estimated above chance level based on different response features of both cell types. Optimal estimation of touch location required a temporal response feature — the relative latency of two cells of the same type (T-T or P-P). In contrast, touch intensity could be estimated best based on the summed spike counts of cell pairs, with T-T, P-P and T-P pairs yielding similar results. Different roles of the cell types for stimulus encoding became evident when both stimulus properties were varied in combination. Optimal estimation results were obtained for touch location based on the relative latency of T cells combined with the summed P-cell spike counts for touch intensity. Hence, the leech seems to use a minimal ensemble code for multiplexed encoding of combined stimulus properties. To test whether this encoding strategy is actually decoded in the leech nervous system, we currently perform intracellular and voltage sensitive dye recordings of interneurons receiving synaptic inputs from mechanosensory cells during skin stimulation.

III-33. Bayesian inference of spinal interneuron classes reveals diverse and clustered cell types.

Mariano Gabitto MIG2118@COLUMBIA.EDU
Ari Pakman ARI@STAT.COLUMBIA.EDU
Jay Bikoff JB3098@COLUMBIA.EDU
LF Abbott LFABBOTT@COLUMBIA.EDU
Thomas Jessell TMJ1@COLUMBIA.EDU
Liam Paninski LIAM@STAT.COLUMBIA.EDU

Columbia University

A systematic characterization of the diversity of neuronal cell types has been identified as a key component for analyzing and understanding neural circuits. We have developed CellTypeHMC, a platform for studying cell type diversity in transcriptionally profiled cell populations. A novel Hamiltonian Monte Carlo algorithm is used to estimate genetically defined cell types efficiently. New clustering representations are introduced that organize cell-type variability on the basis of transcription profiles. Furthermore, spatial distributions of gene expression are used to constrain statistical estimates, providing better defined genetic cell types and identifying their spatial locations. Applying this approach to a class of spinal cord interneurons, we identified genetically heterogeneous populations localized in compact spatial domains. Predictions were validated using triple antibody staining for particular subpopulations. By combining genetic and anatomical features, we have begun to assemble a comprehensive catalog of interneuron cell types in the spinal cord. Our classification approach is sufficiently flexible to incorporate molecular, anatomical and physiological definitions of cell type. Altogether, CellTypeHMC can help elucidate the underlying fundamental principles of cell type characterization in the nervous system.

III-34. A flexible code: the retina dynamically reallocates its resources to code for complex motion

Stephane Deny[1] STEPHANE.DENY.PRO@GMAIL.COM
Vicente Botella-Soler[2] VBSOLER@IST.AC.AT
Serge Picaud[3] SERGE.PICAUD@INSERM.FR
Gasper Tkacik[2] GASPER.TKACIK@IST.AC.AT
Olivier Marre[1] OLIVIER.MARRE@GMAIL.COM

[1]Institut de la Vision
[2]IST Austria
[3]Institut de la Vision INSERM UMRS 968

Ganglion cells, the retinal output, perform non-linear computations on moving objects, but it is unclear if they make motion coding more efficient. We recorded large populations of more than 100 ganglion cells using large multi-electrode arrays in the rat retina. We separated the population into different subtypes and displayed either one or two randomly moving bars as a stimulus. When only one bar was displayed, even cells whose receptive field was far from the bar responded reliably to motion. These distant cells coded mostly for the velocity of the bar. They constituted a channel that provided additional information about the bar's trajectory. However, this response to motion outside the receptive field was strongly suppressed as soon as another bar, with a different trajectory, was displayed inside their receptive field. When the second bar was added, most of these distant cells actively suppressed their response to the first bar, and coded the position of the second bar instead. This effect was strongly non-linear. We found a subtype of ganglion cells whose response to a moving bar inside their receptive field could be well predicted by a LNP model. However, neither the response to the most distant bar nor the suppression could be predicted. Our results suggest that the retina is able to reallocate its neuronal resources to code for multiple moving objects: in the absence of motion in their receptive field centers, the ganglion cells can actively respond to distant stimuli, even outside of their surrounds, but switch to encoding dynamic stimulus features in their centers when they are present. We are currently investing the mechanisms behind this non-linear processing, to model how the ganglion cells can adjust their selectivity depending on the amount of information that needs to be encoded.

III-35. Interneurons in a two population grid cell network

Ziwei Huang[1] ZIWEIH@STUD.NTNU.NO
Trygve Solstad[2] TRYGVE.SOLSTAD@GMAIL.COM
Benjamin Dunn[3] BENJAMIN.DUNN@NTNU.NO

[1]Norwegian University of Science and Technology
[2]Sor-Trondelag University College
[3]Kavli Institute for Systems Neuroscience, NTN

The hexagonal firing pattern of entorhinal grid cells could arise from a competitive mechanism mediated by interneurons. Here we asked if a two-population continuous attractor model, consistent with the proposed inhibitory connectivity pattern, could maintain grid cell firing even if interneurons (a) comprise less than 20% of the neural population and (b) lack spatial periodicity, as was recently observed in a sub-population of entorhinal interneurons. First, using non-negative matrix factorization (NMF), we constructed two-population models with varying numbers of interneurons while maintaining the same effective connectivity between grid cells. Surprisingly, network drift decreased exponentially with the number of assumed interneurons and networks having less than 10% interneurons were able to accurately path integrate. The resulting connectivity was patterned with each interneuron receiving projections from either many grid cells with similar spatial selectivity or cells that together formed an inverted grid pattern. In both cases, grid cells with inhomogeneous peak firing rates had lower grid scores than the corresponding interneurons. Interestingly, thought to be outliers, a small number of interneurons with both high grid scores and spatial sparsity have also been observed experimentally. Second, we considered a network

where the connections from grid cells to interneurons were fixed to sparse random values, while back projections were found using NMF. In this case, the spatial selectivity of interneurons decreased dramatically as the variance in grid field firing rates was increased. Although this network produced aperiodic interneurons similar to recordings, a considerably larger proportion of interneurons was required to reach the same level of stability which did not decrease exponentially as in the fully factorized case. Further experiments should be able to determine if reality falls somewhere on the spectrum between these two simple cases.

III-36. Environmental boundaries as an error correction mechanism for grid cells

Kiah Hardcastle KHARDCAS@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Lisa Giocomo GIOCOMO@STANFORD.EDU

Stanford University

Accurate spatial navigation is critical for survival: it enables predator avoidance and piloting to remembered food locations. Navigational ability likely relies on internal neural maps of external space (Tolman 1948). Grid cells, found in medial entorhinal cortex (MEC), may serve this function, as they fire in periodic, hexagonally arranged locations reminiscent of a longitude and latitude coordinate system (Hafting et al., 2005). Many have proposed that this firing pattern arises through path-integration, a landmark-independent calculation of position computed though integration of an inertia-based velocity signal. However, the additive nature of path-integration may lead to accumulating error and, without a corrective mechanism, inaccuracy in position estimates. We hypothesized that environmental boundaries could contribute to such a corrective mechanism. To test this, we examined grid cells in behaving rodents as they traveled across an open arena. First, we find that error in grid cell spiking accumulates at a rate of ~0.015 cm/sec relative to time since the animal last encountered a boundary. Second, the spatial pattern of accumulating error spikes was consistent with coherent drift of grid-like neural activity patterns in attractor network models for grid cells (e.g. Burak & Fiete, 2009). Third, encounters with a boundary correct errors in grid cell spiking perpendicular, but not parallel, to the boundary, consistent with an error correction mechanism driven by border cells that fire along the entire boundary length. Furthermore, we reproduced all of these experimental observations in an augmented attractor network model that combined grid cells and border cells. Connectivity between border cells and grid cells consistent with Hebbian plasticity was sufficient to account for all observations. Our results propose a fundamental role for the use of landmarks in grid-cell driven spatial navigation, and suggest a specific neurobiological mechanism by which neurons that code for environmental boundaries support noise-robust representations of external space.

III-37. Formation of dorso-ventral grid cell modules: The role of learning

Irmak Aykin AYKIN@EMAIL.ARIZONA.EDU
O. Ozan Koyluoglu OZAN@EMAIL.ARIZONA.EDU
Jean-Marc Fellous FELLOUS@EMAIL.ARIZONA.EDU

University of Arizona

Hippocampal place cells are active at specific locations, and grid cells in the medial entorhinal cortex show periodic grid-like firing as a function of location. Both types of cells are organized from dorsal to ventral levels in increasing spatial field sizes. The size gradient seems smooth for place cells, but modular for grid cells. Together, grid and place cells form a topographical map for navigational tasks. This map can be characterized by a weight matrix, where the entries correspond to synaptic connections between grid and place cells. Using a Hebbian plasticity rule with decay, these connections can be strengthened or weakened in accordance with the fields visited during behavior. Thus, a rat learning a path through specific locations will form a weight matrix that could act as a signature for the learned path. Starting with a homogeneous, smooth distribution of firing field gradients of place and grid cells, we study the structure of the weight matrix when following actual paths recorded from rodents. Two conditions were considered: random foraging and learning of paths between specific rewarded locations. The resulting weight matrices show significant differences: a) Weight matrix corresponding to foraging shows a smooth connectivity pattern throughout the synaptic population; b) The weight matrix of the path with reward locations shows a clear pattern consisting of sub-modules organized along the dorso-ventral axis. This result suggests that grid cells synaptically group themselves in modules on a learned path. c) This modularity does not appear in place-place or grid-place connections. d) These results are compatible with, and may partially explain, the electrophysiological results obtained experimentally. Overall, our plasticity-based framework is a novel computational model that suggests a mechanism for the formation of dorso-ventral modules in grid cells.

III-38. Concentration-invariant identification of odors in a model of piriform cortex

Merav Stern[1,2] MERAV.STERN@MAIL.HUJI.AC.IL
Kevin A Bolding[3] BOLDING@NEURO.DUKE.EDU
Larry Abbott[1] LFA2103@COLUMBIA.EDU
Kevin M. Franks[3] FRANKS@NEURO.DUKE.EDU

[1]Columbia University
[2]The Hebrew University of Jerusalem
[3]Duke University

Odor identity is encoded by ensembles of piriform cortex neurons that appear relatively invariant to odorant concentration. In contrast, the olfactory bulb mitral cells that provide sensory input to piriform cortex represent odors through patterns of activity with temporal and spatial structure that scales with concentration. We have constructed a model that shows how concentration-dependent temporal bulb output is transformed into a concentration-invariant cell-identity code in piriform cortex. Our piriform model consists of a population of pyramidal cells that receive mitral-cell input, feedforward inhibition, feedback inhibition and recurrent excitation from other cortical pyramidal cells. Odor identities are defined by specific sequences of glomerular activations, where latency to respond decreases after inhalation onset as concentration increases. Mitral cells belonging to an active glomerulus fire with Poisson statistics, providing input to the model of piriform cortex. The earliest active glomeruli provide suprathreshold input to a small number of the model's cortical pyramidal cells. These cells provide diffuse recurrent excitation to all other pyramidal cells, but this is only sufficient to drive spiking in cells that have already received some subthreshold bulb input. This produces a steep ramping of cortical activity, which then activates strong feedback inhibition to suppress subsequent cortical activity. Cortical odor ensembles are therefore defined by the first few active glomeruli and, because changes in odor concentration alter the latency but not the sequence of glomerular activations, cortical ensembles are largely concentration invariant. Using total spiking over a sniff as input, a simple perceptron accurately discriminates different odors and generalizes across concentrations. Odorant concentration is encoded by the latency to fire and the coherence of spiking in the cortex. Newly acquired experimental data support these modeling results.

III-39. Cortical feedback decorrelates olfactory bulb output in awake mice

Gonzalo Otazu OTAZU@CSHL.EDU
Hong Goo Chael CHAE@CSHL.EDU
Martin Davis DAVISM@CSHL.EDU
Dinu Albeanu ALBEANU@CSHL.EDU

Cold Spring Harbor Laboratory

Mammalian early sensory systems are characterized by large numbers of descending fibers from the cortex that usually outnumber feedforward inputs to the cortex. Despite their anatomical prominence, the dynamics and roles played by feedback signals remain largely unknown. Computational models in the visual and auditory systems proposed that such signals might provide specific information to the periphery helping in the discrimination of the incoming sensory input. We used multiphoton calcium imaging to report the activity of piriform feedback axons in the olfactory bulb of awake head-fixed mice. We found that cortical feedback axons transmit specific odor information to selected targets in the bulb. Feedback responses often outlast odor presentation by several seconds suggesting that cortical feedback reflects olfactory input received in the recent past. Cortical feedback axons represented individual odors through a combination of enhancement and suppression of bouton activity. However, any given bouton responded with stereotypic polarity across multiple odors, preferring either enhancement or suppression, suggesting that odor information is encoded through two distinct channels. To understand the role of cortical feedback in the bulb we silenced pharmacologically the piriform cortex, while imaging the activity of the bulb output neurons (mitral and tufted cells). We found that cortical inactivation increased odor responsiveness of mitral cells, but had little impact on tufted cells. Cortical feedback specifically decorrelated mitral, but not tufted cells responses increasing the discriminability of different odor representations. Our findings support the idea that odor evoked cortical feedback acts differentially on the two output channels of the bulb (mitral vs. tufted cells) facilitating odor identification.

III-40. Coding of concentration in olfactory bulb: Origins of olfactory “where”

Ana Parabucki NAARYA@GMAIL.COM
Torr Polakow TORRPOLAKOW@GMAIL.COM
Alexander Bizer ABIZER@UNIV.HAIFA.AC.IL
Genela Morris GMORRIS@SCI.HAIFA.AC.IL
Roman Shusterman RSHUSTERMAN@UNIV.HAIFA.AC.IL

University of Haifa

Navigation using olfactory cues entails processing of three distinct types of information: odor identity, odor concentration, and time of cue sampling. The last two components govern the odor source localization process. Localization of the odor source relies on sampling and comparing odor concentration. This may be done in the spatial domain, by bilateral sampling and comparing between two nostrils, or temporal domain, by sequential sampling and comparison of concentrations sampled at consecutive times. While the role of the spatial strategy in olfactory navigation has been studied in various species, mechanism of temporal sampling and its contribution to odor localization remain obscure. To study the neural basis of temporal sampling in mouse olfactory bulb, we designed a unique odor delivery setup. We monitored activity of mitral/tufted (M/T) cells in response to an olfactory stimulus, which flickered between two odor concentrations. M/T responses to the flickering task revealed two distinct M/T populations: a) concentration tracking cells — providing information on actual odor concentration, and b) gradient detection cells — providing information on the concentration change. To address the mechanism behind these responses, we used a transgenic mice line in which Channelrhodopsin2 is expressed selectively in olfactory sensory neurons that contain M72 receptors and converge into M72 glomerulus. Light stimulation of this specific glomerulus revealed that the associated M/T cells' responses to identical light stimuli depend on previous stimuli intensity. These responses resembled those seen in gradient detection cells, suggesting that the underlying mechanism might be intrinsic. In summary, our results indicate that M/T cells respond to specific concentration or to change of odor concentration. The coexistence of this two different M/T responses may reveal two different streams of information in the olfactory system, resembling “what” and “where” streams of the visual system.

III-41. Random connections from the olfactory bulb to cortex support generalization across odors and animals

Evan Schaffer[1] ESS2129@COLUMBIA.EDU
Dan Stettler[1] DDS2110@COLUMBIA.EDU
Daniel Kato[1] DK2643@COLUMBIA.EDU
Gloria Choi[2] GBCHOI@MIT.EDU
Richard Axel[1] RA27@COLUMBIA.EDU
LF Abbott[1] LFABBOTT@COLUMBIA.EDU

[1]Columbia University
[2]Massachusetts Institute of Technology

Learning to associate sensory stimuli with appropriate behavioral responses requires a compromise between specificity and generalization. Stimulus representations that vary discontinuously between similar stimuli allow those stimuli to be discriminated and associated with different behaviors. Conversely, representations that vary smoothly between similar stimuli permit generalization of the same behavior across stimuli. In the olfactory system, projections from the olfactory bulb to the piriform cortex are unstructured. Their apparent randomness might suggest that the olfactory system has opted for specificity over generalization because random connections tend to decorrelate responses. To determine how well piriform odor representations support generalization, we compared calcium-imaging data from piriform cortex to a bulb-piriform model with random feedforward connectivity. In our imaging data, although representations of most odorant pairs are uncorrelated, odorants of similar structure have correlated representations. The model can generate correlated output matching our imaging data purely from correlated input because random connectivity does not completely decorrelate. Moreover, many model piriform neurons respond similarly across a class of related odorants despite the randomness of their input. Thus, random connectivity can support generalization, and smooth response tuning does not imply structured connectivity. To test for generalization in learning, we trained a perceptron-like readout receiving piriform input to respond to a single odorant. With no further training, this readout responds to similar odorants and to mixes containing the trained odorant but not to odorants unrelated to the trained odorant. Finally, we studied generalization across animals by generating multiple random models with the same statistics. These models exhibit similar generalization after learning. Thus, behavioral consistency across animals does not require similarity in the wiring of their olfactory systems. This leads to the counterintuitive prediction that different mice with different random odorant representations will tend to make the same mistakes in an olfactory task.

III-42. Neural circuits for the cortical control of the optokinetic reflex

Baohua Liu[1] LBAOHUA@UCSD.EDU
Andrew Huberman[1] AHUBERMAN@UCSD.EDU
Massimo Scanziani[1,2] MASSIMO@UCSD.EDU

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

The cerebral cortex of mammals has the ability to control and modulate innate, reflexive behaviors mediated by subcortical structures. By adjusting these behaviors to the prevailing conditions or according to past experiences the cortex greatly expands the behavioral repertoire of mammals. One example of such cortical modulation is the impact of visual cortex on the optokinetic reflex (OKR), a compensatory eye movement that stabilizes retinal images during self-motion. The initial stages of the OKR are mediated by phylogenetically old subcortical brainstem nuclei of the accessory optic system (AOS). Here we study the cortical control of the OKR as a model system to understand the mechanisms enabling cortex to modulate innate behaviors. We optogenetically silenced the visual cortex of mice to evaluate the cortical contribution to the OKR and to the activity of AOS nuclei. We observed that cortical silencing moderately but significantly reduced the OKR gain by 10-30%, depending on visual stimulus parameters. Furthermore the reduction in OKR gain upon cortical silencing could be largely

accounted for by the concomitant decrease in AOS activity without affecting the sensorimotor transformation function. We further demonstrate that layer V neurons in the visual cortex project to and form functional excitatory synapses onto AOS neurons, implying that the cortex can modulate AOS activity through this direct projection. Finally we discover that surgical disruption of the vestibulo-ocular reflex, another gaze stabilization mechanism, leads to a compensatory enhancement of OKR gain and that this enhancement depends on visual cortex. Our results thus indicate that visual cortex, via its direct corticofugal projection, amplifies the activity in the AOS to modulate OKR behavior. This circuit enables visual cortex to compensate for disruptions in reflexive gaze stabilization mechanisms.

III-43. Sleep selectively resets V1 responses during perceptual learning

Malte Rasch MALTE.RASCH@BNU.EDU.CN
Yin Yan YAK@MAIL.BNU.EDU.CN
Minggui Chen MINGGUI.CHEN@GMAIL.COM
Wu Li LIWU@BNU.EDU.CN

Beijing Normal University

It is well known that sleep has important contributions to the consolidation of memories. The neural basis of the beneficial effect of sleep is, however, relatively unknown. One hypothesis is that synaptic modifications acquired during the day are “down-selected” by the nights sleep, so that a selected few “important” synaptic modifications are more likely to persist, whereas the bulk of “spurious” synaptic modifications are reset to base level. To test this theory, we monitored the population activity of neurons in macaque V1 by chronically implanted electrode arrays during a visual perceptual learning task for two weeks. While it has been reported previously that population responses of macaque V1 progressively change over the days of learning, we here investigate how a night’s rest affects the population responses the next day. By using de-mixing PCA to separate effects of trials and days, we show that the overall population responses are consistently modified within a single day of learning. However, after a night’s rest the acquired overall response dynamics is almost completely reset the next morning, only to be regained by the next day’s efforts, suggesting that the bulk of the acquired modifications are lost during the night. The monkeys’ behavioral performance, on the other hand, tends to be nevertheless improved in early trials in comparison to late trials within a single day. This indicates that some important modifications might be selectively retained. Indeed, the task-relevant aspects of the population code (the ratio between responses of neurons nearby the figure with those on the background) were not affected by a night’s rest and instead improved progressively with training without being reset overnight. Our results thus support the idea that sleep down-selects activity dependent modifications.

III-44. 7T laminar fMRI indicates a preference for stereopsis in the deep layers of human V1

Nuno Goncalves[1] NRG30@CAM.AC.UK
Jan Zimmermann[2] JAN.ZIMMERMANN@NYU.EDU
Hiroshi Ban[3] BAN.HIROSHI@NICT.GO.JP
Rainer Goebel[4] R.GOEBEL@MAASTRICHTUNIVERSITY.NL
Andrew Welchman[1] AEW69@CAM.AC.UK

[1]University of Cambridge
[2]New York University
[3]Osaka University
[4]Maastricht University

Stereopsis, the impression of solid appearance from binocular disparity, depends on a sequence of computations that match corresponding features between the two eyes. The first step in this process is understood as image correlation, captured by the binocular energy responses of V1 neurons. This correlation can be positive (images match between the eyes) or negative (e.g., white features match black features), evoking responses from V1 neurons in both cases. Yet, stereoscopic perception rejects negative correlations. How, and where, is this rejection of anticorrelation achieved? Here we use 7T functional magnetic resonance imaging at sub-millimeter resolution to examine laminar responses to stereoscopic stimuli in V1. Human participants (N=4) viewed random dot stereograms in correlated and anti-correlated form. We measured blood-oxygenation level dependent (BOLD) signals along the calcarine sulcus using a gradient and spin-echo sequence with 0.8 mm isotropic resolution. A general linear model showed increased activity for correlated versus anti-correlated stereograms in deep layers of V1. Conversely, BOLD activity in superficial layers did not differ between these conditions. A series of control experiments ruled out bias in BOLD responses across cortical layers. Using multivariate pattern classification, we found that voxels in deep layers of V1 are weighted most strongly when a classifier is trained to discriminate fMRI activity evoked by correlated vs anti-correlated stereograms. These results indicate a preference for disparities that support perception in the deep layers of V1. This is compatible with an influence of feedback from higher areas, where the stereo correspondence problem has been solved. Such feedback may account for the reduced amplitude of V1 responses to anticorreleated stimuli, which represents a puzzling deviation of these neurons from the strict implementation of a binocular energy computation.

III-45. ON - OFF maps in ferret visual cortex

Gordon Smith[1] GORDON.SMITH@MPFI.ORG
David Whitney[1] DAVID.WHITNEY@MPFI.ORG
Matthias Kaschube[2] KASCHUBE@FIAS.UNI-FRANKFURT.DE
David Fitzpatrick[1] DAVID.FITZPATRICK@MPFI.ORG

[1]Max Planck Florida Institute for Neuroscience
[2]Frankfurt Institute for Advanced Studies

The primary visual cortex of carnivores and primates contains multiple overlapping functional maps, including those of retinotopy and ocular dominance, as well as maps of orientation and direction preference. Orientation and direction selectivity arise through the pooling of inputs originating in ON- and OFF- center retinal ganglion cells. These ON and OFF pathways remain segregated in the LGN and the spatial distribution of thalamocortical axons in layer 4 of carnivores exhibits clustering based on center type. However, previous in vivo imaging studies have been unable to visualize a columnar representation of ON and OFF responses in layer 2/3, suggesting that these pathways converge shortly after entering the cortex. Here we show, using both in vivo wide-field epifluorescence and two-photon imaging of virally-expressed GCaMP6s, strong evidence for the presence of ON - OFF maps in layer 2/3 of ferret visual cortex. Full-field luminance changes drive transient but reliable responses, which are highly patchy across the cortical surface, with largely complementary patterns dominated by responses to either luminance increments (ON) or decrements (OFF). ON - OFF responses can be evoked in nearly 40% of labeled neurons, a subset of which failed to respond to grating stimuli. Cells responsive to luminance steps appear to exhibit significant spatial clustering across the cortical surface, and highly ON - OFF selective domains exhibit reduced orientation selectivity. In addition to full-field luminance changes, ON - OFF selective responses can be evoked by drifting oriented edges, suggesting that this pathway is engaged during normal viewing. Taken together, our results demonstrate the presence of a novel ON - OFF map in the superficial layers of ferret cortex, which may represent a parallel channel for the processing of luminance information in the cortex.

III-46. Rapid optimal integration of tactile input streams in somatosensory cortex

Hannes Saal HSAAL@UCHICAGO.EDU
Michael Harvey MAHARVEY555@GMAIL.COM
Sliman Bensmaia SLIMAN@UCHICAGO.EDU

University of Chicago

Perception reflects the seamless integration of signals from a variety of sensory receptors that each respond to different aspects of the environment. In touch, two afferent classes exhibit different but overlapping sensitivities to skin vibrations between 50 and 800 Hz: rapidly-adapting (RA) afferents are most responsive in the low range, while Pacinian (PC) afferents are most sensitive at high frequencies. Here, we investigate how signals from these two afferent populations are transformed and combined in cortex. To this end, we reconstruct the responses of afferent populations to a wide variety of simple and complex vibrotactile stimuli and record the responses of neurons in primary somatosensory cortex (S1) to these same stimuli. We then develop simple models that allow us to examine how signals from the different afferent populations combine to shape responses in cortex. We find that, while most cortical neurons receive input from both afferent classes, signals from the two classes drive S1 responses differently: RA input is primarily excitatory and thus determines the strength of the cortical response, whereas PC input seems to sharpen the timing of cortical responses without having a major impact on their strength. Importantly, we show that this integration process maximizes information transmission across the range of stimuli tested and does so rapidly. Our findings thus highlight that sensory cortex is tuned to the statistics of its input streams and able to integrate information across them in a near-optimal way.

III-47. Oxytocin enables maternal behavior by balancing cortical inhibition

Bianca Marlin BIANCA.JONES@MED.NYU.EDU
Mariela Mitre MARIELA.MITRE@MED.NYU.EDU
James D'Amour JAMES.DAMOUR@MED.NYU.EDU
Moses Chao MOSES.CHAO@MED.NYU.EDU
Robert Froemke ROBERT.FROEMKE@MED.NYU.EDU

New York University

Oxytocin is important for social interactions and maternal behavior (Insel and Young, 2001; Churchland and Winkielman, 2012; Dulac et al., 2014). However, little is known about when, where, and how oxytocin modulates neural circuits to improve social cognition. Here we describe how oxytocin enables pup retrieval behavior in female mice by enhancing auditory cortical pup call responses. Naive virgin females initially do not retrieve pups. We tested when virgin females first expressed retrieval behavior after interacting with dams and pups, and found that oxytocin (pharmacologically-applied or optogenetically-released) accelerated the time to first retrieval. Expression of retrieval behavior required left but not right auditory cortex, and expression of retrieval behavior was accelerated by oxytocin in left auditory cortex. We made new antibodies specific to the mouse oxytocin receptor, and found that oxytocin receptors were preferentially expressed in left auditory cortex. Electron microscopy revealed oxytocin receptors at synapses, including inhibitory terminals on excitatory neurons. Pup calls are known to evoke more activity in auditory cortical neurons of mothers compared to naive virgin females (Ehret, 2005; Liu et al., 2006; Cohen et al., 2011). We next performed in vivo whole-cell recordings to examine this transformation of responses from the virgin state to the maternal state. We made current- and voltage-clamp recordings to measure spiking and synaptic responses to pup calls. Neural responses to pup calls were lateralized, with co-tuned and temporally-precise responses to pup calls in left primary auditory cortex (AI) of maternal but not pup-naive adults and not right AI. Importantly, pairing calls with oxytocin (pharmacologically or optogentically) enhanced call-evoked responses by balancing the magnitude and timing of inhibition with excitation in virgins. Our results describe fundamental synaptic mechanisms by which oxytocin increases the salience of acoustic social stimuli. Furthermore, oxytocin-induced plasticity provides a biological basis for lateralization of auditory cortical processing.

III-48. Discriminating partially occluded shapes: insights from visual and frontal cortex

Anitha Pasupathy PASUPAT@UW.EDU
Amber Fyall FYALLA@UW.EDU

University of Washington

The primate brain successfully recognizes objects even when they are partially occluded. Visual area V4 and the prefrontal cortex likely play important roles in this perceptual capacity but their contributions are unknown. We measured the responses of neurons in V4 and the ventrolateral prefrontal cortex (vlPFC) while monkeys discriminated pairs of shapes under varying degrees of occlusion. For most V4 neurons, partial occlusion caused a weakening of early (40-70 ms latency) shape selective responses, but over time selectivity gradually increased and peaked around 200 ms after stimulus onset. This delayed emergence of selectivity under occlusion was also recently observed in human inferotemporal cortex (ITC) [1]. In striking contrast to these visual areas, neurons in vlPFC, which receive visual form information primarily from V4 and ITC, showed the opposite trend—responses increased with increasing occlusion. Across the vlPFC population, this response pattern had the effect of amplifying responses and selectivity to occluded stimuli; because signals in vlPFC peaked ~150 ms after stimulus onset, they are appropriately timed to serve as the feedback modulation that facilitates the gradual increase in shape selectivity in V4. We implement a V4-PFC network model wherein vlPFC responses are derived by gain-modulating the occluded shape signals from V4 by the level of occlusion also derived from V4 responses. Responses in vlPFC are then fed back to V4, facilitating an increase in response amplitude and selectivity for the most occluded stimuli. Our experimental results and model provide the first elaboration of how PFC feedback can selectively augment impoverished shape information in feedforward signals and contribute to enhanced shape selectivity in visual cortex and to the successful discrimination of partially occluded shapes.

III-49. Mosaic representation of odors in the output layer of the mouse olfactory bulb

Hong Goo Chae[1] HGCHAE@CSHL.EDU
Daniel Kepple[1] DKEPPLE@CSHL.EDU
Alexei Koulakov[1] KOULAKOV@CSHL.EDU
Venkatesh N Murthy[2] VNMURTHY@FAS.HARVARD.EDU
Dinu Albeanu[1] ALBEANU@CSHL.EDU

[1]Cold Spring Harbor Laboratory
[2]Harvard University

Characterizing the neural representation of chemical space is a formidable challenge in olfaction research. Unlike many other sensory systems, low-dimensional metrics for characterizing odors have remained elusive and it is unclear what features of chemical stimuli are represented by neurons. Here, we have endeavored to relate neural activity in the early olfactory system of mice to the physico-chemical properties of odorants. We imaged odor-evoked responses in identified tufted and mitral cells in awake mice using multiphoton microscopy. Although both mitral and tufted cells responded with diverse amplitudes and dynamics, mitral cells responses were on average sparser and less sensitive to changes in concentration of odorants compared to tufted cells. We characterized odorant features using a comprehensive set of 1,666 physico-chemical properties that has been extensively used previously. Similarity of physico-chemical features of odor pairs was a poor predictor of similarity of the corresponding neuronal representation by mitral or tufted cells. Dimension reduction revealed that ~22 dimensions could explain more than 90% of the variance in neural responses across the population, but fewer dimensions (~12) were necessary if neural activity was projected on to the space of physico-chemical properties. This suggests that factors other than the physico-chemical properties we considered, including non-sensory signals, are required to fully explain the neural responses. Responsive mitral and tufted cells were spatially dispersed, and cells within a local region were functionally heterogeneous with respect to odor identity and concentration. We

used dimension reduction strategies to determine whether any odorant property is laid out in an orderly manner spatially and found only limited and variable dependence of mitral/tufted cell position on odorant characteristics. Our data indicate that novel descriptors are needed to link chemical space to neuronal representations and that odor information leaves the bulb in a mosaic pattern, with substantial local diversity.

III-50. Independent behavioral primitives underlie behavioral response to odors in Drosophila.

Seung-Hye Jung SEUNGHYE.JUNG@DUKE.EDU
Sam Gerber SGERBER@MATH.DUKE.EDU
Jeff Beck BAYESIAN.EMPIRIMANCER@GMAIL.COM
Vikas Bhandawat VB37@DUKE.EDU

Duke University

Animals parse their sensory inputs to orchestrate diverse motor programs. Most prevailing approaches to the study of sensory-motor transformations divide the process into two steps: a decision step in which a sensory cues are analyzed to select a behavior, and an execution step in which the selected behavior is implemented. Support for this serial model of sensory-motor processing derives primarily from studies of trained animals performing learned behavioral associations. How sensory-motor processes unfold during the performance of innate behaviors remains poorly understood. In this study we use Drosophila olfactory system to delineate the logic that underlies innate sensorimotor transformation. We established a novel behavioral paradigm which allows precise stimulus control and a fine-grained analysis of locomotion. A detailed analysis of how different attractive odors modulated a fly's locomotion led us to a model of fly olfaction which had two salient features: First, odors independently modulate a surprising number of locomotor parameters. Second, each olfactory receptor neuron (ORN) affects a subset of ORN classes. Overall, our experiments support the idea that independent behavioral primitives underlie behavioral response to odor. To better understand how olfactory behaviors are decomposed into behavioral primitives, we fit a Hierarchical Hidden Markov Model (HHMM) to the data. A 2-level HHMM with 8 high-level states (behavioral primitives) underlies modulation of behavior by apple cider vinegar (ACV), a food odor. Odors modulate the probability a fly is in a given state. Remarkably, null mutation in four of the seven ORNs (called Orco-ORNs) abolishes modulation in some behavioral primitives without affecting others. Mutation in the other three ORNs (called Ir8a-ORNs) strongly affects the complementary set of primitives. Orco-ORNs are critical for global search to find the odor source, while Ir8a-ORNs affect the switch from global to a local search once the odor is found.

III-51. Rapid linear decoding of olfactory perception during flight

Laurent Badel LAURENT@BRAIN.RIKEN.JP
Kazumi Ohta KAZUMI30@BRAIN.RIKEN.JP
Yoshiko Tsuchimoto TSUCHIMOTO@BRAIN.RIKEN.JP
Hokto Kazama HOKTO_KAZAMA@BRAIN.RIKEN.JP

RIKEN Brain Science Institute

Innate attraction and aversion to odors depend on internal hedonic percepts formed in the brain downstream of olfactory receptors. In the fruit fly Drosophila melanogaster, activating single glomeruli in the antennal lobe (AL) can evoke attraction or aversion, suggesting that hedonic value is encoded at the level of individual glomeruli. However, it is unclear whether this is a property of a few privileged glomeruli, or constitutes a general principle of odor valuation in the fly brain. Moreover, little is known about how well and quickly flies can assess the hedonic valence of discrete odor plumes found in natural environments. Here, we monitor behavioral responses in a flight-simulator setup, and in parallel measure odor-evoked Ca 2+ signals in most AL output neurons. We observe that

flies can make behavioral decisions within a few 100s of ms, and that gradual changes in AL response evoke gradual changes in behavior, suggesting that valence may be decoded by a linear readout of the onset of AL response. Consistent with this hypothesis, we find that a linear model can recapitulate the observed behavior, and predict responses to novel odors not used in the fitting procedure. Individual glomeruli are assigned weights consistent with previous findings, but small in magnitude, indicating that odor valence is not dominated by a few privileged glomeruli, but depends on pooling small contributions over a large number of glomeruli. This conclusion is further supported by genetic silencing of individual olfactory receptors, which is found to have little impact on behavior. We also image Ca 2+ signals in AL output neuron terminals in the lateral horn, and identify an anatomical region that is differentially activated by attractive and aversive odors. Activity of this region alone was sufficient to predict the observed behavior, suggesting a neuronal substrate for coding valence-specific information.

III-52. Functional connectivity shapes the dynamics of visual object recognition

Emin Orhan EORHAN@CNS.NYU.EDU

New York University

Multivariate pattern analysis of neural responses provides a detailed picture of the time-course of object recognition at multiple levels of abstraction. Previous work in this field has shown that representations that support coarse-level classification of visual images (e.g. images of animate vs. inanimate objects) both emerge and decay more slowly than representations that support finer-level classification (e.g. images of cars vs. bikes). However, a mechanistic explanation of this finding is lacking. Are the slower emergence and slower decay of coarse-level representations caused by the same mechanism? What does this finding imply about the structure of the coarse- and fine-level neural representations that support object recognition? Here, I address these questions in the context of simple dynamical neural network models with recurrent connectivity. I argue that the temporal dynamics of object decoding can be informative about the underlying functional connectivity of the neural units. In particular, I show that the slower decay of coarse-level visual representations can only be explained if the functional connectivity matrix has a hierarchical, or ultrametric, structure. However, the slower emergence of these same representations can be explained by the structure of the stimuli alone without assuming any special functional connectivity, thus suggesting that the slower emergence and slower decay of coarse-level visual representations might be caused by different mechanisms. I further show that properties of the connectivity matrix, such as the sparsity of connections, have natural signatures in the temporal dynamics of decoding accuracy. An analytical theory based on linear dynamics, Gaussian noise and linear discriminant analysis qualitatively explains these main results. Numerical simulations with a more realistic noise model and nonlinear dynamics also validate the results. In summary, the results reported here illustrate how we can relate structure to function in simple dynamical neural networks.

III-53. A theoretical framework for optogenetic perturbations in the oculomotor system

Nuno Calaim[1] NUNO.CALAIM@NEURO.FCHAMPALIMAUD.ORG
David Barrett[2] DAVID.BARRETT@ENG.CAM.AC.UK
Pedro Goncalves[3] PEDROG@GATSBY.UCL.AC.UK
Sophie Deneve SOPHIE.DENEVE@ENS.FR
Christian Machens[1,4] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Neuroscience Programme
[2]University of Cambridge
[3]Gatsby Computational Neuroscience Unit, UCL
[4]Ecole Normale Superieure

How do neural networks respond to instantaneous perturbations of their activity? This question has been the subject of intense investigation ever since the advent of optogenetic perturbation techniques, which allow us to instantaneously perturb neural activity and record the response. We do not yet have a theoretical framework to adequately describe the neural response to such optogenetic perturbations, nor do we understand how neural networks can perform computations amid a background of ongoing natural perturbations. In this work, we develop a framework to describe the impact of optogenetic perturbations on the oculomotor integrator (OI). The OI is a neural structure in the hindbrain which is responsible for controlling eye position by integrating eye movement signals to produce eye position signals. We build a spiking network model of the OI from first principles, following the approach of Boerlin et al. 2013. Specifically, we postulate that the connectivity and dynamics of neurons in OI are optimized to represent eye movement signals using a linear decoder (analogous to a dendritic summation). The resulting spiking network replicates key properties of the OI, such as the typical distribution of tuning curves and accurate eye position representation (Aksay et al. 2000, 2004). We can now do simulated optogenetics in our model: we artificially perturb membrane voltages and record the impact of these perturbations. We find that changes in eye position in our model are consistent with recent optogenetic experiments in which the OI was perturbed with Halorhodopsin and Channelrhodopsin (Goncalves et al. 2014). This indicates that the OI acts to instantaneously adjust the activities of the unperturbed neurons in order to compensate for any error in the computation performed by the OI. More generally, these results suggest that our framework may provide a useful and timely tool for characterizing the impact of optogenetic manipulations.

III-54. Efficient probabilistic inference with oscillatory, excitatory-inhibitory neural circuit dynamics

Laurence Aitchison[1] LAURENCE.AITCHISON@GMAIL.COM
Máté Lengyel[2] M.LENGYEL@ENG.CAM.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University of Cambridge

Probabilistic inference is a powerful framework for understanding both behaviour and cortical computation. However, two basic and ubiquitous properties of cortical responses, that neural activity displays prominent oscillations in response to constant input, and large transient changes in response to stimulus onset, are in seeming contrast with models of cortex based on probabilistic inference that typically either completely lack dynamics, or have simplistic dynamics that give neither oscillations nor transients. We asked whether these dynamical behaviors can in fact be understood as hallmarks of the specific representation and algorithm that the cortex uses to perform probabilistic inference. Based on the observation that oscillations are particularly useful for rapidly spanning a large volume of state space, we developed a neural network model that was inspired by a statistical sampling algorithm devised to speed up inference by introducing auxiliary variables, Hamiltonian Monte Carlo (HMC). Our insight was that inhibitory interneurons in a cortical circuit are ideally suited to play the role of auxiliary variables, and their interaction with principal excitatory neurons naturally leads to oscillations. Thus, we constructed our neural network model to have an E-I structure, which indeed resulted in oscillations. Importantly, not only oscillations and transients emerged in this network, but they were beneficial for probabilistic inference which was an order of magnitude more efficient than in a non-oscillatory variant. The network matched two further properties of neural dynamics that would not otherwise appear to be compatible with probabilistic inference. First, the frequency of oscillations as well as the magnitude of transients increases as the contrast of an image stimulus increases. Second, excitation and inhibition are balanced, and inhibition lags excitation. These results suggest a new functional role for inhibitory neurons and neural oscillations in enhancing the efficiency of cortical computations.

III-55. Emergence of identity-independent object properties in ventral visual cortex

Daniel Yamins YAMINS@MIT.EDU
Ha Hong HAHONG@MIT.EDU
James DiCarlo DICARLO@MIT.EDU

Massachusetts Institute of Technology

Inferior Temporal (IT) cortex, the highest area of the ventral visual stream, is well-studied as the neural correlate of viewpoint invariant object recognition. It is commonly hypothesized that the ventral stream builds invariance by generalizing the simple/complex cells of V1, hierarchically pooling units sensitive to identity-preserving transformations (eg. translations, dilations) in successive areas. On this view, IT should be less sensitive than lower and intermediate visual areas (eg. V1, V4) to identity-independent variables like object position and size. Here, we present evidence that this mechanistic picture must be revised. We first optimized a deep neural network for object categorization accuracy. Consistent with previous work (Yamins et. al, 2014), the top layer of this network is predictive of neural spiking responses in IT cortex. Surprisingly, this same network achieves high performance in estimating many non-categorical visual properties, including object position, size, and pose – even though it was only explicitly optimized for categorization. Moreover, we find that all tested identity-independent variables are better estimated at each successive model layer, even as tolerance to these variables simultaneously emerges. These observations make a counterintuitive neural prediction: IT cortex should encode a spectrum of non-categorical visual properties. Neural recordings in V4 and IT and a high-fidelity model of V1 show that this prediction is borne out: the actual IT neural population significantly outperforms lower visual areas such as V1 and V4 in estimating both all tested properties, including those (eg. position) that are normally thought to be supported by low-level ventral areas. Comparing neural populations to human psychophysical measurements, we find that IT cortex is significantly more predictive of human performance patterns than V1 or V4. Our results suggest that the ventral stream should be thought of as representing not merely object category but instead a rich bundle of behaviorally-relevant object-centric properties.

III-56. Identifying proxies for behavior in full-network C. elegans neural simulations

James Kunert KUNERT@UW.EDU
Eli Shlizerman SHLIZEE@UW.EDU
J. Nathan Kutz KUTZ@UW.EDU

University of Washington

The nematode C. elegans has had the connectivity between all 302 of its neurons (its connectome) entirely resolved, and its behavioral responses to stimuli have been shown by PCA to be fundamentally low-dimensional. In our previous work we constructed a full-connectome model for neural voltage dynamics and showed that such a model was capable of generating neural patterns which serve as proxies for these low-dimensional responses. Our study develops a technique for relating simulated neural patterns to experimental observations through simulated ablation. Simulated ablations are performed by disconnecting the "ablated" neurons from the connectome. Neural proxies for behavioral responses should, under ablation, cause changes consistent with observed behavioral changes in experimental ablations. We quantify the impact of simulated ablations by calculating the SVD of the motorneuron voltage dynamics, finding the low-dimensional modes excited within the system by given inputs and comparing these in the ablated versus non-ablated cases. This enables direct comparison with numerous experimental ablation studies. As an example, we consider the stimulation of the PLM mechanosensory neuron pair. These neurons excite forward motion, and our previous study showed that their stimulation within our model generates a consistent two-mode oscillatory response. Experimental studies have shown that the ablation of the AVB interneurons destroys the worm's ability to move forward. Our simulated ablation study replicates this, destroying the system's ability to generate the corresponding neural pattern when AVB is disconnected (but preserving this ability under other ablations). This demonstrates the technique's ability to relate simulated neural activity to experimental observations without making assumptions specific to the response or the model. Thus it may be extended to identify proxies for more complex behaviors within neural simulations,

III-57. Maintaining stable perception during active exploration

Yuwei Cui[1,2] YWCUI@UMD.EDU
Subutai Ahmad[2] SAHMAD@NUMENTA.COM
Chetan Surpur[2] CSURPUR@NUMENTA.COM
Jeff Hawkins[2] JHAWKINS@NUMENTA.COM

[1]University of Maryland, College Park
[2]Numenta Inc.

Our sensory input changes dramatically as the result of our own behavior, including eye movements, head turns, and body movements. Despite these rapid sensory changes, our perception of the world is amazingly stable, and we can reliably discriminate between different patterns. This suggests that we learn stable but distinct representations through active exploration. There is reason to believe that efference copy, an internal copy of the motor signal, is critical for such sensorimotor learning. However the exact brain mechanisms underlying these computations remain unknown. In this study, we propose a computational model of sensorimotor learning and prediction. Sparse distributed representations of visual scenes are built up incrementally by pooling together predictable temporal transitions during exploration. To enable accurate predictions during active exploration, we modified the Hierarchical Temporal Memory sequence-learning algorithm to use both sensory inputs and efference copy signals. To enable forming stable representations of scenes, we implemented a novel temporal pooling learning rule that allows downstream neurons to form connections with upstream neurons that are predicting correctly. The overall model is unsupervised and the architecture is consistent with several important aspects of thalamocortical circuits. We tested the algorithm on a set of simulated environments, as well as a robotics test bed. In both cases the model achieves two desired properties: 1) prediction of future sensory inputs during behavior, and 2) emergence of stable and distinct representations for learned patterns. After learning, the sparse activity of cells in downstream regions is stable despite sensor movements, while different images lead to distinct representations. These results demonstrate how efference copy can be used in sensory cortex to make predictions during behavior. We propose temporal pooling as a novel computational principle for forming invariant representations during unsupervised learning and active exploration.

III-58. Dynamic interaction between Pavlovian and instrumental valuation systems

Rong Guo[1,2] RONG@NI.TU-BERLIN.DE
Jan Glaescher[2] GLAESCHE@UKE.DE
Klaus Obermayer[1] KLAUS.OBERMAYER@MAILBOX.TU-BERLIN.DE

[1]Technische Universität Berlin
[2]University Medical Center Hamburg-Eppendorf

Pavlovian valuation influence on instrumental behavior has recently been studied in humans with the Pavlovian-to-instrumental transfer (PIT) paradigm. While these studies provide evidence for the neural correlates of PIT effect when subjects are tested in extinction, they do not address how the two systems cooperate and compete during learning process. Here we approach this question computationally using reinforcement learning theory in conjunction with model-based fMRI, utilizing a task that requires both types of leanings in parallel. 27 participants were instructed to predict the occurrence of a visual stimulus on either the left or the right side of the screen with the goal of maximizing their rewards. Stimuli appeared probabilistically on the left and right screen side in two

conditions: (a) equal stimulus probability of 0.5 or (b) biased stimulus probabilities of 0.7 and 0.3, counterbalanced between runs. Conditional reward occurred with probabilities of 0.2 and 0.8 for correct predictions on either side. Crucially, the higher reward location was the opposite of where the stimulus appeared most frequently, thus permitting us to disentangle the stimulus-driven and the reward-driven learning processes. We found that subjects tended to trade off the reward against the predictability of stimulus location. We modeled choice behavior with a dynamic hybrid of two Rescorla-Wagner models. We found a co-existence of Pavlovian and instrumental prediction errors in the ventral striatum, suggesting that this region response to surprising perceptual events as well as unexpected reward delivery or omission. These data raise the possibility that Pavlovian learning can be achieved with the exact same neural computations responsible for instrumental learning. Furthermore, amygdala activity correlated with a dynamic tradeoff between the two learning processes, which reflects its prominent role in PIT effect. In summary, our study highlighted the neural computations underlying a dynamic interaction between Pavlovian and instrumental systems for guiding behavior.

III-59. Unexpected functional diversity among retinal ganglion cell types of the mouse retina

Philipp Berens[1,2] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Tom Baden THOMAS.BADEN@UNI-TUEBINGEN.DE
Katrin Franke KATRIN.FRANKE@UNI-TUEBINGEN.DE
Miroslav Roman Roson MIROSLAV.ROMAN.ROSON@GMAIL.COM
Matthias Bethge[3] MATTHIAS@BETHGELAB.ORG
Thomas Euler THOMAS.EULER@CIN.UNI-TUEBINGEN.DE

[1]BCCN Tübingen
[2]BCM Houston
[3]University of Tübingen

Visual processing begins in the retina, where retinal circuits compute many visual features in parallel from the pattern of light falling into the eye. The results of these computations are sent to the brain by the retinal ganglion cells (RGCs). Anatomical studies have identified 15-22 morphologically different RGC types, suggesting that as many functional channels may form the output of the retina. To determine how many such ‘feature channels’ exist in the mammalian retina, we used two-photon calcium imaging to record almost 10,000 cells in the ganglion cell layer using a standardized set of light stimuli. We sorted the cells into functional types using a probabilistic clustering framework. We found that RGCs can be divided into at least 30 functional types based on the visual responses. The responses and morphologies of many of these matched known types, e.g. the ON- and OFF alpha-like RGCs or the local-edge-detector (‘W3’). In addition, we identified new RGC types, such as an OFF DS RGC that does not co-stratify with starburst amacrine cells and a contrast-suppressed type. To test whether these functionally defined RGC groups indeed correspond to single RGC types, we measured how well the dendritic fields of each type covered the retinal surface. Most functional groups had a coverage factor ~1, suggesting that none of them has been spuriously split. Some groups had a coverage factor $>>$ 1, suggesting that they likely consist of multiple subtypes (e.g. the ON-OFF DS RGC had a coverage factor ~4, presumably corresponding to the known subtypes pointing in the four cardinal directions). The summed coverage factors from all groups indicate that there may be around 50 types of RGCs in the mouse. Taken together, our data indicates that information channels from the eye to the brain may be much more diverse than previously thought.

III-60. Fast effects of the serotonin system on odor coding in the olfactory bulb

Vikrant Kapoor VKAPOOR@MCB.HARVARD.EDU
Allison C Provost APROVOST@FAS.HARVARD.EDU
Venkatesh N Murthy VNMURTHY@FAS.HARVARD.EDU
Harvard University

Serotonin is a neuromodulator whose actions are involved in the regulation of brain states over time scales of minutes and hours. Here we show that the serotonergic system plays a role in odor representation in the olfactory bulb on behavioral timescales (hundreds of milliseconds). We used in vivo multiphoton microscopy of mitral and tufted cells expressing calcium indicators, as well as single unit recordings to examine the effects of activation of the serotonergic system in mice. Brief stimulation of the dorsal raphe nucleus, a key serotonergic center, led to excitation of tufted cells at rest and potentiation of their odor responses. By contrast, while mitral cells at rest were excited by DRN activation, their odor responses could be suppressed or potentiated depending on odor and cell identity. This bidirectional modulation led to decorrelation of mitral cell outputs, which could increase discriminability of odors. We then used in vitro slice electrophysiology to uncover the biophysical mechanisms underlying the distinct actions of the serotonergic system on different populations of principal neurons in the olfactory bulb. Using selective optogenetic activation of serotonergic axons, we found that the serotonergic modulation of mitral and tufted cell activity was caused, surprisingly, by both glutamate and serotonin. Our data indicate that the raphe nuclei, in addition to their role in neuromodulation of brain states, are also involved in fast millisecond time scale top-down modulation, similar to cortical feedback. Notably, they can play differential roles in sensitizing or decorrelating different populations of output cells.

III-61. Changes in inhibition explain cortical activity patterns and laminar variability

Carsen Stringer[1] CARSEN.STRINGER@GMAIL.COM
Marius Pachitariu[2] MARIUS@GATSBY.UCL.AC.UK
Dara Sosulski[1] DARA.SOSULSKI@GMAIL.COM
Sanjeevan Ahilan[1] SANJEEVAN.AHILAN.13@UCL.AC.UK
Peter Bartho[3] BARTHO.PETER@KOKI.MTA.HU
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
Michael Hausser[1] M.HAUSSER@UCL.AC.UK
Peter Latham[2] PEL@GATSBY.UCL.AC.UK
Nicholas Lesica[1] N.LESICA@UCL.AC.UK
Maneesh Sahani[2] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]Gatsby Computational Neuroscience Unit, UCL
[3]Hungarian Academy of Sciences

Activity in the brain varies greatly, from desynchronized firing to up-down states. Activity also varies across layers, with superficial and deep layers exhibiting correlated variability while granular layers show near-zero correlations. To understand these behaviors, we combined computational methods with electrophysiological and anatomical studies. Our main result is that inhibition both controls network states and produces diversity across layers. We first analyzed multi-neuron stimulus-driven activity from the deep layers of auditory cortex. We found that intrinsic large-scale fluctuations on multiple timescales can corrupt sensory encoding in awake and anesthetized animals, but are abolished in desynchronized states which encode stimuli with high-fidelity. We constructed a biologically-plausible spiking network model with chaotic population-wide fast (50-200 ms) and slow (500-1000 ms) fluctuations that quantitatively reproduced this behavior. Such chaotic network states are a closer approximation to awake neural dynamics than the classical asynchronous state, because in our data awake states displayed

fast population-wide fluctuations. In simulations, enhanced non-selective inhibition was essential to stabilize fluctuations and improve coding properties. This is consistent with our data: putative fast-spiking inhibitory activity correlated with increases in decoding accuracy and large decreases in tuning widths and noise correlations. Could laminar differences in noise correlations thus be due to inhibitory-stabilization? We obtained laminar distributions of interneurons in auditory cortex from confocal imaging of PV-conditional mouse lines and counted cells using a novel algorithm. The deep layers had the lowest PV interneuron density while the granular layers had the highest. Incorporating this distribution into our model we obtained zero noise correlations in the input layer but positive correlations in the other layers. We further noticed that synchronized states can amplify perturbations of single spikes, as recently observed. This work suggests that a thorough understanding of inhibitory neurons, including subtypes, is critical for understanding the behavior of networks in the brain.

III-62. Modulating response dynamics with the zebrafish habenula

Seetha Krishnan[1] SEETHAKRIS@GMAIL.COM
Ruey-Kuang Cheng[2] RKCHENG@IMCB.A-STAR.EDU.SG
Suresh Jesuthasan[2] SURESHJJ@IMCB.A-STAR.EDU.SG

[1]National University of Singapore
[2]Institute of Molecular and Cell Biology

The behavior of an animal is dependent on context and value of stimuli. This plasticity is enabled by the release of neuromodulators, resulting in the selection of different functional networks under different conditions. Here, we investigate how neuromodulator release is controlled in response to various sensory stimuli with innate value. The evolutionarily conserved habenula complex, located in the diencephalon receives input from sensory systems and basal ganglia, and projects to the raphe, VTA and locus coeruleus. It is thus well placed to dynamically control neuromodulator release. We used both high-speed wide-field and two-photon calcium imaging of transgenic zebrafish to create four dimensional maps of activity in the habenula and outputs, in response to defined stimuli. PCA using Thunder (Freeman et.al.) was used to identify spatial patterns from temporal activity. The habenula responds to light, odor and electric shock in multiple distinct subnuclei. There are differences in the encoding of the type and concentration of odor, wavelength of light, and onset and offset of stimuli. Based on femto-second laser ablation, we propose that the dorsal left subnucleus enables inhibition of serotonergic neurons in the raphe by light, a rewarding stimulus. The ventral habenula, in contrast, mediates excitation of serotonergic neurons in response to light. We are currently testing how multi-modal stimulation is represented, and examining how intra-habenula computation shapes output, given existence of slow waves of evoked activity across the habenula, gap junctions as well as long-range connection between habenula neurons. Thus, the zebrafish habenula, which is easily accessible to in vivo calcium imaging, robustly activated by sensory stimuli and effortlessly manipulated genetically and by laser ablation, is an attractive system for dynamical modeling and the study of stimulus-evoked neuromodulator release.

III-63. Relating functional connectivity in V1 neural circuits and 3D natural scenes using Boltzmann machine

Yimeng Zhang[1] YIMENGZH@CS.CMU.EDU
Xiong Li[2] FLIT.LEE@GMAIL.COM
Jason Samonds[1] SAMONDJM@CNBC.CMU.EDU
Benjamin Poole[3] POOLE@CS.STANFORD.EDU
Tai Sing Lee[1] TAI@CNBC.CMU.EDU

[1]Carnegie Mellon University
[2]Shanghai Jiao Tong University
[3]Stanford University

Bayesian theory has provided a compelling conceptualization for perceptual inference in the brain. To understand the neural mechanisms of Bayesian inference, we need to understand the neural representation of statistical regularities in the natural environment. Here, we investigated empirically how the second order statistical regularities in natural 3D scenes are represented in the functional connectivity of a population of disparity-tuned neurons in the primary visual cortex of primates. We applied the Boltzmann machine to learn from 3D natural scenes and found that the functional connectivity between nodes exhibited patterns of cooperative and competitive interactions that are consistent with the observed functional connectivity between disparity-tuned neurons in the macaque primary visual cortex. The positive interactions encode statistical priors about spatial correlations in depth and implement a smoothness constraint. The negative interactions within a hypercolumn and across hypercolumns emerge automatically to reflect the uniqueness constraint found in computational models for stereopsis. These findings demonstrate that there is a relationship between the functional connectivity observed in the visual cortex and the statistics of natural scenes. This relationship suggests that the functional connectivity between disparity-tuned neurons can be considered as a disparity association field. They also suggest that the Boltzmann machine, or a Markov random field in general, can be a viable model for conceptualizing computations in the visual cortex, and as such, can be used to leverage the natural scene statistics to understand neural circuits in the visual cortex.

III-64. Testing a disinhibitory circuit in mouse visual cortex

Mario Dipoppa MARIO.DIPOPPA@GMAIL.COM
Adam Ranson RANSON.AD@GOOGLEMAIL.COM
Matteo Carandini MATTEO@CORTEXLAB.NET
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK

University College London

The dynamics of cortical networks are shaped by the functional connectivity between different classes of interneurons. In the primary visual cortex (V1) of the mouse, it has been proposed that two of these classes of interneurons are arranged in a disinhibitory circuit[1]. In this circuit, interneurons expressing Vasoactive Intestinal Peptide (VIP+) would inhibit interneurons expressing Somatostatin (SOM+), and thus disinhibit pyramidal cells[2,3]. This disinhibitory circuit seems to capture the effects of locomotion on activity in the absence of visual stimuli[1,4], and to agree with data indicating that SOM+ cells provide surround suppression to pyramidal cells[5]. However, it seems at odds with observations made in the presence of large visual stimuli, where running has been reported to increase visual responses in both SOM+ and VIP+ neurons[2,6]. Perhaps the action of the disinhibitory circuit depends on the presence and size of visual stimuli? To address this question, we used two-photon calcium imaging in mouse V1 to study how locomotion modulates activity in VIP+ and SOM+ interneurons and their properties of spatial integration. Contrary to the predictions of the disinhibitory circuit, locomotion strongly increased visual responses of all cell types considered, and also increased spontaneous activity of SOM+ and VIP+ neurons. Locomotion, moreover, decreased surround suppression[7] in both putative excitatory neurons and in SOM+ interneurons. This observation puts in question the notion that SOM+ interneurons are responsible for surround suppression of pyramidal cells[5]. Hence our data do not appear to be consistent with the simple disinhibitory and the surround suppression circuits that have been proposed for mouse V1. 1 Pfeffer et al., Nat. Neurosci. (2013). 2 Fu et al., Cell (2014). 3 Zhang et al. Science (2014). 4 Niell & Stryker, Neuron (2010). 5 Adesnik et al., Nature (2012). 6 Polack et al., Nat. Neurosci. (2013). 7 Ayaz et al., Current Biol. (2013).

III-65. Hippocampal phase precession is disrupted after medial entorhinal cortex lesions

Christian Leibold[1] LEIBOLD@BIO.LMU.DE
Magdalene Schlesiger[2] SCHLESIGER.M@GOOGLEMAIL.COM
Christopher Cannova[2] CCCANNOV@UCSD.EDU
Brittney Boublil[2] BBOUBLIL@UCSD.EDU
Jena Hales[2] J9DAVIS@UCSD.EDU
Jill Leutgeb[2] JLEUTGEB@UCSD.EDU
Stefan Leutgeb[2] SLEUTGEB@UCSD.EDU

[1]Ludwig-Maximilians-Universität München
[2]University of California, San Diego

The formation of episodic memories requires the compression of sequences of behavioral events to a time scale at which synaptic learning rules, such as spike-timing dependent plasticity, can take effect. In the hippocampus and entorhinal cortex, time compression is thought to occur by the theta-phase precession of place cell assemblies, whereby individual neurons fire at progressively earlier phases of the theta oscillation as the animal traverses a cell's place field. Along with the sequential activation of place cells during locomotion along a linear track, phase precession results in sequences of spike phases within each theta cycle that correspond to the sequences in which place field are arranged, and the sequences that occur over seconds on the spatial scale are expressed over tens of milliseconds on the neuronal scale and therefore on the timescale required for synaptic plasticity. The mechanisms underlying hippocampal phase precession are still unclear. Competing models propose that hippocampal phase precession may either arise exclusively from local network dynamics or, alternatively, may require inputs from additional cortical areas, in particular from the theta-modulated and/or phase-precessing cells in medial entorhinal cortex (MEC). Here we show that complete MEC lesions resulted in a substantial disruption of hippocampal phase precession such that it was no longer detectable in either single runs or when pooling data from all runs through a place field. Despite the substantial disruption of spike timing, theta oscillations were preserved and spatial firing was partially retained after the MEC lesion. However, spatial firing only persisted over short periods of running on the track rather than over the entire experimental session. Our results thus strongly support models in which intrinsic network dynamics within the hippocampus are not sufficient to generate phase precession and in which theta-modulated and/or phase-precessing inputs from the MEC are necessary for hippocampal phase precession.

III-66. The hexagonal grid code as a multi-dimensional clock for space

Martin Stemmler[1,2] STEMMLER@BIO.LMU.DE
Alexander Mathis[3] AMATHSI@FAS.HARVARD.EDU
Andreas Herz[2] HERZ@BCCN-MUNICH.DE

[1]Ludwig-Maximilians-Universität München
[2]Bernstein Center for Computational Neuroscience Munich
[3]Harvard University

Grid cells are neurons that fire whenever the animal's location coincides with a point of an imaginary, hexagonal grid that tessellates space. These neurons, found in the medial entorhinal cortex, pre- and para-subiculum across different mammalian species, are thought to play an important role in spatial navigation. A metric for space, however, cannot be established at the single-grid-cell level; instead, a metric requires an ensemble of neurons organized into separate modules, with each module representing a different spatial scale. We show here how animals can navigate by using simple grid cell population vector averages, much in the same way movement directions can be read out from motor cortex. This decoding strategy enables dead reckoning, is optimal in the face of neuronal noise, and is consistent with many grid cell features: their organization into modules, the alignment of grid cell lattices, and a ratio of successive spatial scales close to 3/2. The periodic nature of grid

cells permits a linear read-out across multiple spatial scales that mimics and generalizes the process of reading a traditional analog clock. Our theory predicts that the nervous system decodes grid cell ensembles using cosine-like tuning curves subject to cosine-like gain fields, akin to the linear gain fields found in parietal cortex.

III-67. Multiple noise sources shape optimal encoding strategies in fundamentally different ways

Braden Brinkman BRADENB@UW.EDU
Alison Weber AIWEBER@UW.EDU
Eric Shea-Brown ETSB@UW.EDU
Fred Rieke RIEKE@UW.EDU

University of Washington

Neural circuits must be able to encode inputs using a limited range of responses in the presence of noise. Influential work by Laughlin showed that neurons should use all responses with equal frequency, devoting a greater range of outputs to more common inputs. This strategy is no longer optimal, however, when the amount or type of noise varies at different processing stages throughout the circuit, or when noise is correlated in parallel circuit pathways encoding the same input. Yet little has been done to systematically study how tradeoffs between different noise sources might impact neural coding. We fill this gap by developing a simple neural circuit model that incorporates multiple noise sources, each arising at a different location in the circuit. Using variational methods, we analytically predict the shapes of the nonlinear input-output relations (“nonlinearities”) that give the best linear decoding of the stimulus. We show that the optimal nonlinearities do not depend on all model parameters, but only on a set of effective parameters, one for each noise source. We numerically verify that optimizing mutual information between the stimulus and responses gives qualitatively similar results. The behaviors predicted by our noise model provide a broader framework to interpret results of several previous studies. The striking qualitative changes we predict for the optimal nonlinearities under different noise conditions indicate that the location at which noise arises in a circuit and the strength of noise correlations are crucial for determining the best encoding strategies.

III-68. Sleep restores variability in synchronization and neuronal avalanches after sustained wakefulness

Christian Meisel[1] CHRISTIAN@MEISEL.DE
Dietmar Plenz[2] PLENZD@MAIL.NIH.GOV

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

Sleep is crucial for daytime functioning and well being. Without sleep optimal brain functioning such as responsiveness to stimuli, information processing, or learning is impaired. Such observations suggest that sleep plays an important role in organizing cortical networks toward states where information processing is optimized. The general idea that computational capabilities are maximized at or nearby critical states related to phase transitions or bifurcations [1] led to the hypothesis that brain networks operate at or close to a critical state. Near phase transitions, a system is expected to recover more slowly from small perturbations, a phenomenon called critical slowing, and observables typically exhibit power-law scaling relationships. Growing experimental evidence on neuronal avalanches [2], i.e. spatiotemporal clusters of synchronous activity in cortex, suggests that the brain under normal conditions resides near a critical state. Here, we hypothesize that sleep deprivation shifts brain dynamics further away from criticality. To address this question, we implanted microelectrode arrays into superficial layers of prefrontal cortex for chronic recordings in the awake, behaving rat. Rats were sleep deprived for 6 h during which MEA activity was recorded. The LFP was z-normalized and thresholded to identify neuronal

avalanches and large synchronization events of increased neuronal spiking. Multiunit activity was identified by offline spike-sorting. During sleep deprivation, spontaneous activity deviated systematically from avalanche dynamics resulting in a growing distortion of the underlying power-law in size distribution. Concomittantly, as a marker for critical slowing down, the recovery from large synchronization events became progressively faster with time awake. All observed changes were reversed by sleep. Our results are in line with previous work suggesting a growing deviation from criticality during wakefulness [3] and could provide a network-level framework for the role of sleep: to reorganize cortical dynamics toward a state where information processing is optimized.

III-69. Feed-forward circuit abnormalities in the mouse model of Fragile X syndrome

Vitaly Klyachko KLYACHKO@WUSTL.EDU
Washington University

Fragile X Syndrome (FXS) is the most common single-gene cause of intellectual disability, often associated with autism and seizures. The common seizure phenotype in FXS patients suggests a dysfunction in circuit excitability (Bassell and Warren, 2008), and many recent studies have implicated excitation/inhibition (E/I) imbalance in the pathophysiology of FXS (Paluszkiewicz et al., 2011). Yet, the mechanisms of circuit hyperexcitability in FXS and the role of inhibitory synapse dysfunction in these defects remain poorly understood. Here, we examined this question in the context of the canonical feed-forward inhibitory circuit formed by the perforant path, the major cortical input to the hippocampus. Perforant path feed-forward circuits exhibited a markedly increased E/I ratio in Fmr1 KO mice, a mouse model of FXS. This E/I imbalance lead to major functional defects in spike discrimination and coincidence detection tasks performed by feed-forward circuits in Fmr1 KO mice. Changes in feed-forward circuits were associated with altered inhibitory, but not excitatory synapse function in Fmr1 KO mice. Perforant path-associated inhibitory synapses exhibited abnormalities in paired-pulse ratio and short-term dynamics during high-frequency trains, consistent with decreased GABA release probability. Inhibitory synaptic transmission in Fmr1 KO mice was also more sensitive to inhibition of GABAB receptors, suggesting an increase in presynaptic GABAB receptor signaling. Indeed, all differences in inhibitory synaptic transmission between Fmr1 KO and WT mice were eliminated by a GABAB receptor antagonist. Inhibition of GABAB receptors or selective activation of presynaptic GABAB receptors also removed all major differences in feed-forward inhibitory circuit properties between Fmr1 KO and WT mice. Our results suggest that the inhibitory synapse dysfunction in the perforant path of Fmr1 KO mice causes E/I imbalance and feed-forward circuit defects, which are predominately mediated by a presynaptic GABAB receptor-dependent reduction in GABA release.

III-70. Computational aspects of visual motion detection in the blowfly

Suva Roy SUVAROY@INDIANA.EDU
Rob de Ruyter van Steveninck DERUYTER@INDIANA.EDU
Indiana University Bloomington

Understanding the computational scheme for motion detection is key to unraveling the functional correlates of the associated circuitry. The Reichardt Correlator model outlines one such scheme, which involves correlation of time delayed inputs from spatially separated locations. While studies in insects and vertebrates provide evidence for a correlation type motion detection, the role of input spatial structure, spatial contrast, adaptation and input noise in the implementation of motion computation is not clear. Using blowfly as the model system and pseudorandom light intensities with specific spatiotemporal correlation as the diagnostic tool, we test both implementation and saliency of operations believed to underlie motion detection. The visual stimulus is presented in a 2-dimensional hexagonal pixel array such that the angular projection of pixels matches the lattice geometry of the blowfly compound eye. Measurements from the horizontal and vertical wide field motion sensitive neurons H1 and V1 reveal

that response depends primarily on correlation between signals from spatial locations separated by not more than 2 units of ‘receptor spacing’. To test robustness of the model’s bilinearity, we recorded neuronal response at different variances of pseudorandom stimuli. The response indeed is an approximately bilinear function of the two correlator inputs, as long as the total flicker energy remains unchanged. However when the flicker energy changes, the response varies inversely with the flicker energy, which the model fails to predict. This indicates that motion estimation is determined not solely by the signal (motion) but rather by the overall signal to noise ratio. We further investigated how motion adaptation shapes time constant of the delay filter in the model. Preliminary results suggest change in time constant of a filter that is perhaps located after the correlation stage.

III-71. Entorhinal inter-laminar coupling during the encoding and recognition of visual stimuli

Nathan Killian[1,2] NKILLIAN@MGH.HARVARD.EDU
Elizabeth Buffalo ELIZABETH.BUFFALO@UW.EDU

[1]Massachusetts General Hospital
[2]Harvard University

Studies in rodents have suggested that slow and fast sub-bands of the gamma rhythm (30-120 Hz) may route separate streams of information through the hippocampal formation to subserve stimulus encoding and recognition. We directly tested this hypothesis with laminar recordings from the EC in monkeys performing a free-viewing recognition memory task in which each image has both an encoding (novel) phase and a recognition (repeated) phase. Superficial layers (I-III) of the entorhinal cortex (EC) provide the majority of the input to the hippocampus and the EC deep layers (V-VI) receive hippocampal output, with only sparse direct superficial to deep connections. Owing to this arrangement, superficial-deep communication reflects on the general input-output relationship of the hippocampal formation comprising the hippocampus, subicular structures, and the EC. Local field potential recordings were localized to individual laminae through registration of current source density (CSD) profiles with histology. Superficial-deep communication in the gamma band was analyzed with Granger causality (GC) to assess the directed influence of hippocampal input signals on output signals. Compared to a reference non-viewing time period, GC increased in the high gamma band (70-120 Hz) during image encoding and increased in the low gamma band (30-70 Hz) during recognition. These results demonstrate, for the first time, that information is routed through the hippocampal formation (from the initial input stage to the output stage) in the gamma band. Furthermore, the gamma rhythm dynamically switches between two distinct sub-bands for visual stimulus encoding and recognition. This scheme could reduce functional crosstalk, i.e. the unwanted mixing of encoding and recognition signals. It follows that differential communication in gamma sub-bands may be utilized by the hippocampal formation to improve the efficiency of recognition memory processes.

III-72. Opponent channel code of auditory space is an efficient representation of natural stereo sounds

Wiktor Mlynarski[1,2] MLYNAR@MIS.MPG.DE

[1]Max-Planck Institute for Mathematics in the Sciences
[2]Massachusetts Institute of Technology

Contrary to early expectations that in the auditory cortex sound position is represented by a topographic map of narrowly tuned units, experimental evidence suggests that spatial information is encoded by the joint activity of two broadly tuned populations known as the opposite channels. Neuronal tuning curves span the entire space surrounding the animal, their slopes are steepest close to the interaural midline, and their peaks are concentrated at lateral positions. It has been argued that observed tuning properties reflect high specialization of sensory neurons to accurately encode the sound position at the midline which is a region of a particular behavioral relevance. This work provides an alternative explanation by demonstrating that a panoramic population code of sound location emerges as a direct consequence of learning an efficient representation of natural binaural sounds. A hierarchical sparse-coding model of binaural sounds recorded in natural settings is proposed. In the first layer, monaural sounds are represented using a population of complex-valued basis functions separating phase and amplitude, both of which are known to carry information relevant for spatial hearing. Monaural input converges in the second layer, which forms a joint representation of amplitudes and interaural phase differences. The model encodes spectrotemporal properties of the sound as well as spatial information with a population of high-level units. Spatial tuning curves were obtained by probing the model with natural sounds recorded at different positions. Responses of the second layer of the model match well the tuning characteristics of neurons in the mammalian auditory cortex. In a broader perspective, the obtained results suggest that various sensory circuits of different functions may share the same design principle - efficient adaptation to the sensory niche.

III-73. The effect of state-space complexity on arbitration between model-based and model-free control

Sang Wan Lee SWLEE@CALTECH.EDU
John P O'Doherty JDOHERTY@CALTECH.EDU

California Institute of Technology

The balance between model-based and model-free reinforcement learning (RL) is suggested to be governed by an arbitration process in which the degree of relative control of the systems over behavior is flexibly adjusted. Precisely how such arbitration occurs is poorly understood. One hypothesis is that the amount of uncertainty associated with the predictions of the two systems are used to allocate control over behavior proportional to each system's precision or reliability. However, uncertainty in a predicted outcome within the two systems is likely to be only one out of several variables involved in arbitration. Here we investigate another potentially important variable: the complexity of the state-space. A more complex state-space might impose increasingly arduous demands on the model-based system, whereas the model-free system could be less affected by complexity. Thus, a natural hypothesis is that more complex state-spaces will be associated with a greater tendency to engage in model-free RL. To test this we designed a novel task in which we systematically manipulated state-transition uncertainty (making the model-based system more or less reliable), goal-values (which if changing can induce reward-prediction errors thereby altering model-free reliability), and state-space complexity. Consistent with our hypothesis, we found behavioral evidence in human participants (N=22) for an effect of state-space complexity on arbitration. However, the nature of that effect was surprising. Rather than resulting in increased reliance on model-free RL as expected, high state-space complexity caused participants to resort to a default bias toward being either model-based or model-free, which varied across participants. One intriguing interpretation of these findings is that when state-space complexity is high, the arbitration process itself is switched off, leaving participants to instead rely on a default strategy irrespective of other variables otherwise driving arbitration such as reliability.

III-74. Spatial rate/phase correlations in theta cells can stabilize randomly drifting path integrators

Joseph Monaco[1] JMONACO@JHU.EDU
H Tad Blair[2] BLAIRLAB@GMAIL.COM
Kechen Zhang[1] KZHANG4@JHMI.EDU

[1]Johns Hopkins University
[2]University of California, Los Angeles

The spatial firing of place cells and grid cells is thought to reflect the association of environmental features, such as external sensory cues and local boundaries, with the path integration of idiopathic cues, such as movement direction and speed. Models of path integration as temporal phase interference among neural oscillators, supported by evidence of theta cells with directionally tuned burst frequency (Welday et al., 2011), must address the critical problem that interference patterns randomly drift in space in the presence of intrinsic phase noise (Monaco et al., 2011; Blair et al., 2014). Here, we present a synchronization theory in which a hypothetical population of theta-rhythmic (6–10 Hz) “location-controlled” oscillators (LCOs) mediates the ability of environmental features to stabilize path integration in a downstream layer of velocity-controlled oscillators (VCOs). We suppose that representations of objects/landmarks or boundaries are weakly theta-modulated and combine to form spatially modulated inputs to the LCOs that are fixed to the environment. This environmental drive combines with ongoing theta oscillations to create a robust correlation between higher firing rates and earlier theta phases of bursting. We show that this rate/phase correlation in LCOs is sufficient to selectively entrain VCOs to prevent drifting with noise. This entrainment requires associative learning during early theta phases at LCO-to-VCO synapses to construct spatially antiphase inputs that drive the environmental feedback to VCOs. We provide a mathematical derivation, an abstract rate-phase model, and an implementation of this sensory feedback mechanism in a spiking network model of phasic bursting neurons. Notably, we present preliminary recording data from subcortical regions in rats, including lateral septum, showing theta cells that qualitatively match the spatial rate/phase correlations of our hypothesized LCOs. These results support a hybrid “place-to-grid” framework where temporal and attractor mechanisms may be complementary without depending on the fine tuning of phase or connectivity.

III-75. Temporal integration in sound texture perception

Richard McWalter[1] RMCW@ELEKTRO.DTU.DK
Josh McDermott[2] JHM@MIT.EDU

[1]Technical University of Denmark
[2]Massachusetts Institute of Technology

Many aspects of auditory perception depend critically on the integration of acoustic information over time, but the mechanisms involved remain poorly understood. Sound textures — signals composed of collections of similar acoustic events — provide a novel avenue for investigating integration. Recent evidence that texture perception may be mediated by time-averaged statistics of peripheral auditory representations (McDermott and Simoncelli, 2013) raises the question of the temporal window over which sound information is integrated. We probed the averaging process using ‘texture gradients’ — signals synthesized such that their statistical properties change over time. We hoped to measure how far back in time the stimulus history would bias texture judgments. On each trial, subjects heard a texture gradient followed by a probe texture with constant statistical properties. Subjects were told that the properties of the gradient stimulus would change over time, and were instructed to compare the end of the gradient to the probe, judging which was closer to a standard reference texture. The gradient stimulus either changed continuously in statistical properties, sometimes concluding with segments with constant statistics, or stepped abruptly from one set of statistics to another. Performance differed for gradients with the same end point in texture space but different starting points, suggesting that stimulus history biases texture judgments via an averaging process. The biasing effect persisted when the gradients concluded with a 1-second segment with constant statistics, but were substantially reduced when the constant segment was extended to 2.5 seconds. An ideal observer model operating on statistics measured from the stimuli mirrored the psychophysical results when a window duration of 1 second or greater was used to compute statistics. The results suggest that texture perception is mediated by an averaging window several second in extent. This duration places a strong constraint on the neural locus of integration.

III-76. Towards understanding mechanisms of pain transmission: a systems theoretic approach

Pierre Sacre P.SACRE@JHU.EDU
William Anderson WANDERS5@JHMI.EDU
Sridevi Sarma SREE@JHU.EDU

Johns Hopkins University

Chronic pain affects about 100 million American adults—more than the total affected by heart disease, cancer, and diabetes combined. Despite their great need, neuropharmacology and neurostimulation therapies for chronic pain have been associated with suboptimal efficacy and limited long-term success as their mechanisms of action are unclear. Understanding the mechanisms of pain transmission to predict its modulation by therapies is therefore essential toward pain management, yet current models suffer from several limitations. In particular, they are not amenable to analysis and fail to provide a comprehensive mechanistic understanding of pain transmission. Using mathematical reduction techniques that exploit time-scale separation, we investigated the cellular dynamics in the dorsal horn of the spinal cord—the first central relay of sensory and pain inputs to the brain. This study proposes a low-dimensional reduced model of dorsal horn transmission neurons and discusses the impact of cellular changes on pain transmission to the brain. The reduced model is sufficient to capture the rich dynamics of transmission neurons in the dorsal horn—from tonic to plateau to endogenous bursting. This cellular switch of firing patterns contributes to a functional switch of information transfer—from faithful transmission to enhancement to blocking of nociceptive information, respectively. In addition, a dynamic balance of intrinsic membrane properties drives the cellular switch from one firing mode to another and therefore the functional switch from one information transfer mode to another. This low-dimensional reduced model is amenable to tractable analysis of the mechanisms of pain transmission and open the door to predict outcomes of refined and/or novel neuromodulation pain therapies.

III-77. Different mechanisms underlie direction selectivity in OFF and ON starburst amacrine cell dendrites.

James Fransen JWFRAN04@LOUISVILLE.EDU
Bart Borghuis BART.BORGHUIS@LOUISVILLE.EDU

University of Louisville

Direction selectivity (DS) is a canonical neural computation. In the retina, direction selective ganglion cells respond strongly to visual motion in a preferred direction and weakly to motion in the opposite, null direction. The source of ganglion cell DS are the presynaptic ON- and OFF-type starburst amacrine cells. While ON SACs have been studied intensely and a mechanism for their DS has been proposed, direct measurements of OFF SAC DS have been lacking and the underlying mechanism has not been explored. We hypothesized that OFF SAC DS differs from ON SAC DS, for two reasons. First, the dominant model for ON DS proposes a dendrite-autonomous mechanism including a soma->dendrite voltage gradient caused by tonic excitation. While tonic glutamate release is present at ON bipolar terminals, this model fails for OFF SACs because their presynaptic OFF bipolar cells lack tonic release. Second, recent EM reconstruction of the presynaptic OFF SAC circuit suggested a Reichardt-type correlation mechanism, driven by sluggish central and fast peripheral dendritic excitation. Here, we combined two-photon imaging and targeted whole-cell electrophysiology to explore the mechanisms underlying OFF SAC DS. While fluorescence imaging showed symmetric bipolar cell glutamate release for inward vs. outward radial motion, current and voltage responses in both ON and OFF SACs were distinctly asymmetric, with faster depolarization during outward motion. Circular white noise analysis of synaptic input across the dendritic arbor showed substantial differences in the sign, amplitude, and spatial organization of inhibition in ON versus OFF SACs. Furthermore, OFF but not ON SACs showed a pronounced temporal gradient of excitatory input, with central excitation trailing peripheral excitation, consistent with the connectomic model. Model simulations showed that linear summation of temporally graded excitatory input across the dendritic arbor suffices for generating DS in OFF SACs. Thus, a dendrite-independent, Reichardt-type mechanism contributes to OFF SAC DS.

III-78. Mechanisms for shaping receptive field in monkey anterior inferior temporal cortex

Keitaro Obara[1,2] KEITARO.O-5@BRAIN.RIKEN.JP
Kazunori O'Hashi[1] KAZ_OI@BRAIN.RIKEN.JP
Manabu Tanifuji[1] TANIFUJI@RIKEN.JP

[1]RIKEN Brain Science Institute
[2]Waseda University

Early behaving monkey studies with fixation tasks showed that the receptive field (RF) of neurons in inferior temporal (IT) cortex is fairly large, e.g. 10.3 degree (Op de Beeck and Vogels, 2000). On the other hand, a recent report revealed that RF can be as small as 2.6 degree when monkeys performed a task to recognize a small stimulus presented near the fixation point (DiCarlo and Maunsell, 2003). What are the factors that caused differences in the RF structure? A possibility is that spatial attention dynamically modulated shapes and sizes of RF. To test this possibility, we trained a monkey with a spatial attention task (Posner, et al., 1978) where the monkey had to detect a change in luminance of one of two dots presented on the screen (90 %change in one dot and 10 %in the other). During the time before the change in dot luminance, object stimuli were flashed at various locations on the screen. We recorded and analyzed multi-unit neural responses to the object stimuli in anterior IT cortex. The spatial attention modulated magnitude of object responses that caused shift of RF shapes toward the attended location, indicating that spatial attention was indeed one of the factors that controls RF shapes of IT neurons. Furthermore, we examined the effect of attention on onset latency of object responses. In contrast to the magnitude of responses, we found that onset latency of object responses, shorter at fovea than periphery, was not affected by spatial attention. The difference in latency was not due to attentional facilitation of processing caused by fixation. We suggest that foveal and peripheral regions of RF are constructed through inherent circuit modules characterized by different latencies, and the outputs of these modules are controlled by spatial attention.

III-79. Correlations and signatures of criticality in neural population models

Marcel Nonnenmacher[1] MARCEL.NONNENMACHER@TUEBINGEN.MPG.DE
Christian Behrens[2] CHRISTIAN.BEHRENS@BETHGELAB.ORG
Philipp Berens[3,4] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Matthias Bethge[2] MATTHIAS@BETHGELAB.ORG
Jakob Macke JAKOB@TUEBINGEN.MPG.DE

[1]Max Planck Institute for Biological Cybernetics
[2]University of Tübingen
[3]BCCN Tübingen
[4]BCM Houston

Large-scale recording methods make it possible to measure the statistics of neural population activity, and thereby to gain insights into the principles that govern the collective activity of neural ensembles. One hypothesis that has emerged from this approach is that neural populations are poised at a 'thermo-dynamic critical point', and that this has important functional consequences (Tkacik et al 2014). Support for this hypothesis has come from studies that computed the specific heat, a measure of global population statistics, for groups of neurons subsampled from population recordings. These studies have found two effects which—in physical systems—indicate a critical point: First, specific heat diverges with population size N. Second, when manipulating population statistics by introducing a 'temperature' in analogy to statistical mechanics, the maximum heat moves towards unit-temperature for large populations. What mechanisms can explain these observations? We show that both effects arise in a simple simulation of retinal population activity. They robustly appear across a range of parameters including biologically implausible ones, and can be understood analytically in simple models. The specific heat grows with N whenever the (average) correlation is independent of N, which is always true when uniformly subsampling a large, correlated population. For weakly correlated populations, the rate of divergence of the specific heat is proportional to the correlation strength. Thus, if retinal population codes were optimized to maximize specific heat, then this would predict that they seek to increase correlations. This is incongruent with theories of efficient coding that make the opposite prediction. We find criticality in a simple and parsimonious model of retinal processing, and without the need for fine-tuning or adaptation. This suggests that signatures of criticality might not require an optimized coding strategy, but rather arise as consequence of sub-sampling a stimulus-driven neural population (Aitchison et al 2014).

III-80. Robust nonlinear neural codes

Qianli Yang[1] QY8@RICE.EDU
Xaq Pitkow[1,2] XAQ@RICE.EDU

[1]Rice University
[2]Baylor College of Medicine

Most natural task-relevant variables are encoded in the early sensory cortex in a form that can only be decoded nonlinearly. Yet despite being a core function of the brain, nonlinear population codes are rarely studied and poorly understood. Interestingly, the most relevant existing quantitative model of nonlinear codes is inconsistent with known architectural features of the brain. In particular, for natural population sizes, such a code would contain more information than its sensory inputs, in violation of the data processing inequality. This is because the noise correlation structures assumed by this model provides the population with an information content that scales with the size of the cortical population, and this correlation structure could not arise in cortical populations that are much larger than their sensory input populations. Here we provide a valid theory of nonlinear population codes that obeys the data processing inequality by generalizing recent work on information-limiting correlations in linear population codes. Although these generalized, nonlinear information-limiting correlations bound the performance of any decoder, they also make decoding more robust to suboptimal computation, allowing many suboptimal decoders to achieve nearly the same efficiency as an optimal decoder. Although these correlations are extremely difficult to measure directly, particularly for nonlinear codes, we provide a simple, practical test by which one can use choice-related activity in small populations of neurons to determine whether decoding is limited by correlated noise or by downstream suboptimality. Finally, we discuss simple sensory tasks likely to require approximately quadratic decoding, to which our theory applies.

III-81. Structural plasticity generates efficient network structure for synaptic plasticity

Naoki Hiratani[1] HIRATANI@BRAIN.RIKEN.JP
Tomoki Fukai[2] TFUKAI@RIKEN.JP

[1]The University of Tokyo
[2]RIKEN Brain Science Institute

Structural plasticity of synaptic connection is critically important for learning process in the brain. Even in the brain of adult animals, 5-15% of spines are created and eliminated everyday, and resultant synaptic connection structure is highly non-random in local cortical circuits. On the other hands, EPSP/IPSP (excitatory/inhibitory post synaptic potential) sizes have rich information capacity and are also modifiable by variety of synaptic plasticity. Correspondingly, previous theoretical results suggested that a randomly connected network is already computationally powerful enough under appropriate synaptic plasticity. Therefore, the functional advantage of structural modification of synaptic connections remains unknown in local circuits. It is also unclear how synaptogenesis generates structure suitable for a particular circuit function. To clarify these points, we constructed a computational model that performs a probabilistic inference from noisy stimuli, and studied the division of labor between structural plasticity and weight modifications and the potential mechanism of structural plasticity through theoretical and numerical analyses. We found that structural plasticity is beneficial for learning in both sparse and dense networks and, in particular, increases transfer entropy between input and neuronal outputs in sparsely connected networks. Furthermore, we found that adequate network structure naturally emerges through dual Hebbian-type learning for both synaptic weight and structural plasticity. Our model of structural plasticity well accounts for several experimental observations on spine dynamics.

III-82. Synaptic clustering or scattering? A model of synaptic plasticity in dendrites

Romain Cazé R.CAZE@IMPERIAL.AC.UK
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK
Simon Schultz S.SCHULTZ@IMPERIAL.AC.UK

Imperial College London

A large body of theoretical work has shown how dendrites increase the computational capacity of neurons. This work predicted that synapses active together should be close together in space, a phenomenon called synaptic clustering. Experimental evidence has shown that, in the absence of sensory stimulation, synapses nearby on the same dendrite can coactivate more than would be expected by chance. Synaptic clustering, however, is far from ubiquitous: other groups have reported that nearby synapses can respond to different features of a stimulus during sensory evoked activity so that synapses tend to activate in a scattered fashion across multiple dendrites. To unify these state dependent experimental results, we use a computational framework to study the formation of a synaptic architecture — a set of synaptic weights — and its function. We present the conditions under which a neuron can learn such synaptic architecture: (1) presynaptic inputs are organized into correlated groups of neurons; (2) the postsynaptic neuron is compartmentalized; and (3), the synaptic plasticity rule is local within a compartment. Importantly, we show that given the same synaptic architecture, synaptic clustering is expressed during learning, whereas synaptic scattering is present under evoked activity. This allows sensory neurons' responses to be sparse. Interestingly, reduced dendritic morphology, a hallmark of several diseases, leads to, in our model, a pathological hyperactivity. This work therefore unifies a seemingly contradictory set of experimental observations: we demonstrate that the same synaptic architecture can lead to synaptic clustering and scattering, and therefore both properties can co-exist in the same neuron.

III-83. Direct evidence of altered cell excitability by extracellular electric fields

Belen Lafon BELULAFON@GMAIL.COM
Asif Rahman ASIFTR@GMAIL.COM
Marom Bikson BIKSON@CCNY.CUNY.EDU
Lucas Parra PARRA@CCNY.CUNY.EDU

The City College of the New York

It is well known that extracellular fields increase neuronal firing. The conventional explanation for this is that fields depolarize a pyramidal soma therefore increasing the likelihood for firing. This seemed like an adequate explanation until it was discovered that synaptic efficacy is also altered by extracellular fields. Thus all prior empirical data on firing is confounded by this possible synaptic effect. Here we use a combination of in vitro experiments and computational modeling to determine whether cell excitability -the likelihood of firing for a given and fixed synaptic input — could be altered by extracellular field stimulation. First we experimentally quantified the effect of stimulation on synaptic input (fEPSP in hippocampal slices), which revealed an increase in synaptic efficacy with somatic depolarization (corresponding to dendritic hyperpolarization) as has been previously reported.A two-compartment neuron model replicates the observed effect and shows that this increase is due to an increase in dendritic driving force when it is hyperpolarized. The same model shows that cellular excitability is also increased.

To quantify this experimental finding we measured the neuronal input/output function (I/O) and found that depolarizing fields also increase excitability, seen as a leftward shift of the I/O function. The computational model points towards a synergy between the change in synaptic input and neuronal output resulting from compartment specific polarization. Mechanistically, dendritic hyperpolarization increases synaptic driving force thus facilitating synaptic transmission (EPSP); simultaneously somatic depolarization leads to an increased firing probability. The differential polarization synergistically enhances firing probability which ultimately will affect local computations and alterations in behavior.

III-84. Homeostatic synaptic depression explains receptive field development by implicit whitening

Carlos Brito CARLOS.STEIN@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Synaptic plasticity is believed to underlie cortical receptive field formation from natural input statistics. While nonlinear Hebbian potentiation can explain this development, it assumes artificially decorrelated inputs and artificial stability constraints, typically attributed to depression and homeostatic mechanisms. When combined with inhibitory plasticity, these mechanisms may have canceling effects, culminating in a run away of both excitatory and inhibitory connections. Here we demonstrate how weight dependent synaptic depression resolves both the limitation of decorrelation and stability. The linear anti-Hebbian character of synaptic depression is shown to make the plasticity rule invariant to second-order input statistics and explains receptive field development without the requirement of whitened input. It also provides robustness to heterogeneities in pre-synaptic firing rates, dendritic attenuation and input redundancies. In a spiking network with both excitatory and inhibitory plasticity, the weight dependent homeostasis deters synaptic run away behavior while preserving the implicit whitening capability, allowing for the stable development of simple cells from natural images. Moreover, recurrent excitatory synapses developed preferential connections between neurons with similar orientations preferences as observed experimentally. These findings give a precise functional interpretation for synaptic potentiation, depression and homeostasis in cortical plasticity, which appears optimally designed for robust feature learning, and elucidates how neurons may self-organize into operative sensory networks.

III-85. Discrimination and production of spatiotemporal patterns with a single recurrent neural network

Vishwa Goudar VISHWA.GOUDAR@GMAIL.COM
Dean Buonomano DBUONO@UCLA.EDU

University of California, Los Angeles

The discrimination and production of complex spatiotemporal patterns is fundamental to most of the computations the brain performs, including speech recognition and production. While significant progress has been made towards developing models that process spatial information (e.g., orientation selectivity or character recognition), most neurocomputational models have failed to emulate the brain's inherent ability to process temporal information. A number of models have proposed that temporal and spatiotemporal computations arise from the dynamics of recurrent neural networks (Durstewitz and Deco, 2008; Rabinovich et al., 2008; Buonomano and Maass, 2009). A limitation of these models is that the most computationally interesting dynamical regimes are also chaotic. Recent work, however, has described techniques that tune the recurrent synaptic weights to create "dynamic attractors"—computationally rich yet locally stable neural trajectories (Laje and Buonomano, 2013). Here was ask if this approach is suitable for sensory processing, e.g., in discriminating the complex spatiotemporal patterns of spoken digits? We establish that dynamic attractor regimes provide a computational basis for speech

discrimination. Furthermore, a single network can be tuned to function in a sensorimotor mode: discriminating spoken digits and reporting the classification through nonlinear motor patterns (Fig. 1). Simulations show that the underlying learning rule utilizes the inherent variations across speakers and utterances of each digit (together with injected neural noise) to tune the RNN to produce digit-specific locally stable trajectories. These dynamic attractors yield speaker and utterance invariant digit discrimination via simple linear readouts. Performance is the result of two properties: (1) the formation of dynamic attractors that capture the spatiotemporal structure of stimuli; and (2) networks that are naturally able to support temporal warping (the ability to recognize similar stimuli played at different speeds). These results also provide one of the first demonstrations of a single network performing both a sensory and motor task.

III-86. Synaptic efficacy tunes speed of activity propagation through chains of bistable neural assemblies

Hesam Setareh HESAM.SETAREH@EPFL.CH
Moritz Deger MORITZ.DEGER@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Reliable propagating of spiking activity is crucial for information transmission and processing in the brain. Models of activity propagation, such as synfire chains (SC) and feed-forward (FF) networks that propagate fluctuations in firing rate through sequentially connected neuron groups, gathered considerable attention in the last decades. In these models, however, activity travels unidirectionally with a high speed of propagation. In principle, such models might also explain how neural circuits can generate trajectories of activity to control stereotypic behaviors, like body movements. However, such behaviors take place on a time scale of seconds, while SC or FF models propagate activity on the time scale of synaptic transmission, which is on the order of milliseconds. Therefore, a SC with a large number of steps is needed to bridge these two time scales, which leads to an excessive demand of neurons for a rather simple function. Here we present an alternative model of activity propagation through nearest-neighbor coupled neuronal assemblies, which we call an ‘excitation chain’. This model does not require exclusive feedforward motifs of the neuronal network structure but propagates activity in both directions. Moreover, the speed of activity propagation can be regulated without modification of the synaptic transmission delay, but by manipulating specific synaptic efficacies. We demonstrate in simulations that the difference between the activation time of the first and the tenth group of an excitation chain can be on the order of 500ms while the synaptic delay is 1ms. The model contains several neuronal assemblies connected in sequence, but bidirectionally. Relatively strong synaptic efficacies and connection probabilities within each assembly, together with short-term synaptic depression, make each assembly act as an excitable element. A chain of these assemblies may propagate activity in a cascade-like fashion with tunable speed, which makes the excitation chain a versatile functional component to generate stereotyped neural activity patterns.

III-87. Relating spontaneous dynamics and stimulus coding in competitive networks

Aubrey Thompson[1] AUBREYTH@ANDREW.CMU.EDU
Ashok Litwin-Kumar[2] AK3625@COLUMBIA.EDU
Brent Doiron[3] BDOIRON@PITT.EDU

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

Understanding the relation between spontaneously active and stimulus evoked cortical dynamics is a recent challenge in systems neuroscience. Recordings across several cortices show highly variable spike trains during spontaneous conditions, and that this variability is promptly reduced when a stimulus drives an evoked response. We have shown how networks of spiking neuron models with clustered excitatory architecture capture this key feature of cortical dynamics. In particular, clusters show stochastic transitions between periods of low and high firing rates, providing a mechanism for slow cortical variability that is operative in spontaneous states. We expand on our past work and explore a simple Markov neural model with clustered architecture, where spontaneous and evoked stochastic dynamics can be examined more carefully. We model the activity of each cluster in the network as a birth-death Markov process, with positive self feedback and inhibitory cluster-cluster competition. Our Markov model allows a calculation of the expected transition times between low and high activity states, yielding an estimate of the invariant density of cluster activity. Using our theory, we explore how the strength of inhibitory connections between the clusters sets the maximum likelihood for the number of active clusters in the network during spontaneous conditions. We show that when the number of stimulated clusters matches the most-likely number of spontaneously active clusters then the mutual information between stimulus and response is maximized. This then gives a direct connection between the statistics of spontaneous activity and the coding capacity of evoked responses. Further, our work relates two disparate aspects of cortical computation–lateral inhibition and stimulus coding.

III-88. Accounting for time delays in dimensionality reduction of neural population activity

Karthik Lakshmanan KARTHIKL@ANDREW.CMU.EDU
Patrick Sadtler PSADTLER@PITT.EDU
Elizabeth Tyler-Kabara ELIZABETH.TYLER-KABARA@CHP.EDU
Aaron Batista APB10@PITT.EDU
Byron Yu BYRONYU@CMU.EDU

Carnegie Mellon University

Dimensionality reduction methods have been applied to study neural population activity during decision making, motor control, olfaction, working memory, and in other contexts. These methods seek to extract a small number of latent variables that can be thought of as a common input that drives the activity of the recorded neurons. Although much has already been seen using these techniques, most current methods assume instantaneous relationships between the low-dimensional latents and high-dimensional population activity. However, latents originating in a driving area in the brain may follow different pathways to reach the recorded neurons. Physical conduction delays as well as different information processing times along each pathway may introduce different time delays from the driving area to each neuron. A dimensionality reduction technique that can account for these delays can provide a more compact representation of the latent space, thereby facilitating easier interpretation. In this work, we introduce a novel probabilistic technique: time-delay Gaussian-process factor analysis (TD-GPFA) that performs dimensionality reduction while accounting for a constant time delay between each pair of latent and observed variables. TD-GPFA models the temporal dynamics of each latent variable using a Gaussian process, allowing us to tractably learn these delays over a continuous domain. We verified using simulated data that TD-GPFA is able to recover the time delays and the correct dimensionality of the latent space. We then applied TD-GPFA to population activity recorded from the macaque primary motor cortex during a reaching task. TD-GPFA is able to describe the neural activity using a more parsimonious latent space than GPFA - a method that has been used to interpret motor cortex data, but does not account for time delays. More broadly, TD-GPFA can help unravel the circuit mechanisms underlying population activity by taking into account physical delays in the system.

III-89. Behaviorally relevant information is revealed in synchrony of neuronal ensembles

Neda Shahidi NEDA.SHAHIDI@UTH.TMC.EDU
Ming Hu MING.HU@UTH.TMC.EDU
Ariana Andrei ARIANA.R.ANDREI@UTH.TMC.EDU
Valentin Dragoi VALENTIN.DRAGOI@UTH.TMC.EDU

University of Texas at Houston

Synchrony of neuronal firing is typically associated with sensory coding. For instance, in visual cortex, the precise coincidence of spikes is observed during passive visual stimulation (Kohn & Smith, 2005). However, whether spike synchrony in neuronal populations contains behaviorally relevant information is poorly understood. We calculated synchrony among multiple neurons (up to 14 in each area) recorded simultaneously using laminar probes in V1 and V4 of behaving monkey. Two animals were engaged in a delayed-match-to-sample task in which they performed correctly on approximately 70% of trials. To measure spike synchrony we calculated empirical probabilities of joint spike events (JSEs) for all possible sub-populations of simultaneously recorded neurons (Pipa et al., 2008). The effect of common visual stimulation or co-fluctuations in firing rates were corrected by jittering. We found that higher-order synchrony among cells occurred before a behavioral decision, and was associated with correct responses. The difference in synchrony between correct and incorrect trials was statistically significant (Wilcoxon signed rank, $P<0.05$) in V4 but not in V1. Interestingly, this effect was only observed for JSE patterns that included triplets or quartets, but not pairs, indicating that assemblies of 3 or more cells may contain behaviorally relevant information. We have also examined the synchrony between simultaneously recorded cells in V1 and V4, and found maximum synchrony in assemblies of 3 or more cells when the delay between V1 and V4 was about 30-ms. We conclude that although the behavioral modulation of single cell firing rates is modest in V1 and V4, synchronous firing of three or more neurons is predictive of behavioral outcomes.

III-90. Peripheral versus internal factors to cortical variability and information

Ruben Coen-Cagli[1] RUBEN.COENCAGLI@UNIGE.CH
Ingmar Kanitscheider[2] INGMAR.KANITSCHEIDER@UNIGE.CH
Adam Kohn[3] ADAM.KOHN@EINSTEIN.YU.EDU
Alexandre Pouget ALEXANDRE.POUGET@UNIGE.CH

[1]University of Geneva
[2]University of Texas at Austin
[3]Albert Einstein College of Medicine

Variability in cortical activity arises from at least two major sources: noise from the sensory periphery propagating to the cortex and internal variability. Internal variability reflects ongoing fluctuations in the activity produced by cortical circuits, including but not limited to recently-described “global” fluctuations in activity. A long-standing proposal is that internal fluctuations are primarily responsible for limiting information in cortex. We sought to test the relative contribution of peripheral noise and internal factors to variability and information. We recorded V1 population responses to gratings corrupted by white noise. We systematically varied the noise amplitude (termed, external noise) thereby controlling the amount of input information, and embedded subsets of frames with an identical noise pattern (frozen seed). First, we found that information in population responses was unaffected by low levels of external noise, but decreased markedly with noise exceeding roughly 5% of the pixel range, similar to levels of noise that limit performance in behavioral experiments. Second, information in population responses was several times larger with frozen than random seed, despite previous findings in area MT which indicated that variability intrinsic to random motion stimuli was responsible for only a minor proportion of cortical variability, if at all (Britten et al., 1993; Bair et al., 2001). Third, noise correlations decreased with external noise. Crucially, we found that all of these results were well captured by a model in which information is limited by variability in the sensory periphery and in the image. Therefore, our results suggest that, while a substantial part of the measured cortical response variability is internal and stimulus independent, cortical information is in fact constrained by computation performed by those circuits on sensory inputs with limited information and corrupted by peripheral noise.

III-91. A common topographic and functional organization for normalization across cortical areas

Joshua Alberts JOA39@PITT.EDU
Douglas Ruff RUFFD@PITT.EDU
Marlene Cohen COHENM@PITT.EDU

University of Pittsburgh

Normalization, in which a neuron's response is divisively scaled when multiple stimuli are presented, is thought to underlie many sensory, motor, and cognitive properties. We hypothesized that if normalization is a general computation involving networks of neurons, signatures of normalization should be similar in different cortical areas and might be revealed by recording from multiple neurons simultaneously. We recorded simultaneously from several dozen electrodes in primary visual cortex (V1) and the middle temporal area (MT) while two monkeys viewed superimposed orthogonal drifting gratings. We compared the response of each unit to superimposed gratings to its response to each grating alone. We observed a continuum of normalization, ranging from suppression from nonpreferred stimuli (strong normalization) to additive responses to superimposed stimuli (no normalization). Our data set revealed three new observations about normalization: 1) As long as a unit responded differently to the two gratings, its degree of normalization did not depend on whether the stimuli were optimized for its tuning, suggesting that normalization is a property of the neuron rather than a response to specific stimuli. 2) There was topographic organization for normalization, meaning that units with a similar degree of normalization were located near each other in the brain. 3) Units that showed strong normalization shared less trial-to-trial variability (had lower noise correlations) than units that showed weak normalization, suggesting that normalization reflects how much a neuron's response is affected by the activity of a large group of neurons. Notably, these observations were similar in V1 and MT, implying that characteristics of normalization are not specific to a particular area. Together, our results suggest that normalization may reflect a neuron's role in the local network, and that the mechanisms underlying factors that divisively scale neuronal responses may share the topographic organization typical of sensory or motor tuning properties.

III-92. Sparse random matrices for neural networks, spectrum density and phenomenology.

Hervé Rouault[1,2] ROUAULTH@JANELIA.HHMI.ORG
Shaul Druckmann[1,2] DRUCKMANNS@JANELIA.HHMI.ORG

[1]Janelia Farm Research Campus
[2]Howard Hughes Medical Institute

Large neural networks can exibit diverse dynamical behaviors: chaos, oscillations, etc. A powerful approach to establishing the existence and prevalence of dynamical regimes in neural circuits consists of studying the linear stability of circuit dynamics and weakly non-linear analysis around its fixed points. Central to such analyses, is the eigendecomposition of the coupling matrix Jij of the neural circuit. The eigenvalue spectrum of a network reveals many properties of the dynamics, for instance, the stability of the dynamical system is set by the eigenvalue with highest real value. Accordingly, it is crucial to be able to determine the eigenvalue spectrum density, and associate different circuit structures with the corresponding eigenvalue spectrum densities. An important structural aspect of many neural circuits is that they conforms with the Dale's law which states that neurons typically do not create both excitatory and inhibitory projections. In terms of the connectivity matrix, this states that the coefficients of

one column of the matrix) should all have the same sign. The eigenvalue density has been determined before for dense matrices that conform with Dale's law and has been refined to allow consideration of the case of several neural populations. However, connections in neural circuits are typically quite sparse, and previous analyses do not hold for sparse matrices. Here, we followed mean-field approaches in order to come up with exact self-consistent analytic expressions for the spectrum density in the limit of sparse matrices (for both the symmetric and more realistic Dale's law network matrices). Moreover, we studied the phenomenology of localization properties of the eigenvectors, finding strong differences in important dynamical regimes between sparse and dense matrices, and further confirmed these results by simulating sparse non-linear neural network to observe the development of dynamics beyond the instability threshold of the network.

III-93. A normative model of neural adaptation via local approximate Bayesian inference predicts V1 response to dynamic stimuli

Alessandro Ticchi[1,2] ALESSANDRO.TICCHI@GMAIL.COM
Aldo Faisal[2] A.FAISAL@IMPERIAL.AC.UK

[1]University of Bologna
[2]Imperial College London

Adaptation plays a fundamental role in neural coding and brain dynamics (Ermentrout, 2001; Crook et al., 1998; Clague et al., 1997). By interpreting adaptation as a continuous inference problem, we derive from first principles a normative model of neural firing rate adaptation to the instantaneous input statistics, and we show that model predictions can explain experimental observations of neural response adaptation in the presence of dynamic stimuli. We assume that the aim of adaptation is to efficiently encode stimuli whose statistics vary in time (Wark et al., 2007) and maximise information transfer, which can be achieved through predictive coding (Deneve, 2008; Boerlin et al., 2013). We implement this mechanism in two steps: 1) the neuron infers the instantaneous stimulus statistics and 2) it adapts its internal response properties accordingly, to maximize information transfer. In the inference process (1), the neuron exploits both the information encoded in its own spiking history and a structured temporal prior over possible dynamics of the encoded stimulus (generative model). For a broad class of stimulus dynamics, the optimal Bayesian posterior can be accurately approximated using a gamma distribution. We derive the evolution equations for the parameters of this distribution based on newly observed inter spike intervals. Under the biologically 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 intercellular messenger molecules in presence of spike events (Halnes et al., 2013). We suggest that these chemical signals encode the updated belief about the instantaneous input statistics of the neuron, and are used to regulate the neural response accordingly (Turrigiano, 2008). Our first principles model of adaptation predicts specific spike adaptation dynamics, which we tested on in vivo neural recordings (Benucci et al., 2013).

III-94. Dimensionality, coding and dynamics of single-trial neural data

Peiran Gao PRGAO@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU

Stanford University

The design of classical neuroscience experiments often involves tightly controlled and repeated simple behaviors or sensory stimuli, over many trials, while recording neural activity. This design allows trial averaging to combat variability. However, to understand more ethologically relevant, complex behaviors, where the same behavioral state or stimulus may rarely be encountered twice, trial averaging is infeasible. Moreover, despite significant advances in recording technologies, in mammalian systems controlling complex behaviors, we still record a small subset, 100~1K, of the number of behaviorally relevant neurons, 1M~1B. This raises a fundamental question: what can we learn about the dynamical properties of neural circuits controlling complex behaviors at such overwhelming levels of subsampling, without trial averaging? To answer this, we initiate a new theory of single-trial data analysis that provides analytical principles for guiding experimental design, and algorithmic methods for extracting structure from multineuronal recordings. We consider model neural circuits with N neurons in which neural activity patterns during an experiment explore K dimensions, and we can only record M<<N neurons for P stimuli/behaviors. Moreover, single neurons have a finite signal-to-noise ratio (SNR). We formulate the structure discovery problem as a low-rank matrix de-noising problem, which we analyze by employing non-commutative probability theory. Our analysis reveals sharp phase transitions, in the M-vs-P plane, in our ability to accurately recover neural dimensionality, decode stimuli on single trials, and recover dynamic circuit properties. In relevant limits, a sufficient condition for accuracy is given by SNR*sqrt(MP) $>$ K. Thus, underlying simplicity in either the behavior or circuit dynamics, as exemplified by small dimensionality K, makes accurate single trial analysis in the vastly under-sampled recording regime possible. Moreover, the algorithms involved, like matrix de-noising and subspace identification, do not require computationally expensive internal loops, like expectation-maximization.

III-95. Modulating regularity of synchronized burst dynamics with noise in cultured/model neuronal network

June Hoan Kim LOVEPHY@GMAIL.COM
Ho Jun Lee SNIPER0712@KOREA.AC.KR
Kyoung Lee KYOUNG@KOREA.AC.KR
Korea University

Synchronized bursts (SBs) are one of the most noticeable dynamic features of neural networks, being essential for various phenomena in neuroscience. They arise in various parts of the brain exhibiting different spatiotemporal modes and play differing roles; simple periodic SBs in neuropathological states or complex periodic sequences under information process, which present even almost random dynamics. Here, one of the fundamental questions is: what governs the regularity or the randomness of SB sequences? In this study we found that the strength of neural noise with respect to the overall excitability of a given culture system, or vice versa, was a key factor determining the regularity of SB sequences. An optimal noise level gave the most precise periodic SB sequence — a phenomenon which we can coin as “SB coherence resonance.” We carefully controlled two different types of neural noise: synaptic noise by different amount of bicuculline, and ion-channel noise by different light intensity in opto-genetically modified neuronal networks with channelrhodopsin-2. In experiments, we controlled the strength of neural noise indirectly, by manipulating the overall excitability of the system with different amounts of bicuculline or with different levels of light intensity in the case of optogenetic light-stimulation. Both methods didn’t show any big difference as far as the phenomenon of the coherence resonance in bursting dynamics. Furthermore, we showed that the system had a hidden stable focus having a slow time scale relevant for the bursting dynamics. Its existence was confirmed through the emergence of a slow underdamped oscillation followed by an electrical pulse stimulation and the phenomenon of frequency-locking subject to a periodic stimulation. All of these experimental observations were successfully recapitulated in computer simulations of a neural network model proposed by Izhikevich, suggesting that the same phenomena may occur in many in vivo as well as in vitro neural networks.

III-96. Influence of recurrent synaptic strengths and noise on grid firing and gamma oscillations

Lukas Solanka LUKAS.SOLANKA@ED.AC.UK
Mark van Rossum MVANROSS@INF.ED.AC.UK
Matthew Nolan MATTNOLAN@ED.AC.UK
University of Edinburgh

Cognitive functions are often associated with changes in gamma frequency oscillations and may involve fine-tuning of synaptic strength. However, a mechanistic relationship between synaptic strength, gamma oscillations, and neural computations underlying cognition has not been determined. We explore this relationship by systematically changing recurrent synaptic strengths in a spiking continuous attractor network model that can generate both grid firing fields and theta-nested gamma oscillations. We find that grid firing and gamma oscillations are sensitive to changes in the strength of excitation and inhibition. However, gamma oscillations carry relatively little information about the presence of stable grid firing fields or the ability of the networks to form a stable attractor state, suggesting that it is not possible to predict whether a network successfully encodes position by measuring the amplitude or frequency of gamma oscillation. Instead, robust grid firing fields generated by the attractor networks are compatible with a wide range of gamma amplitudes and frequencies. Unexpectedly, moderate neural noise promotes generation of both gamma oscillations and grid field computation. The range of synaptic strengths supporting gamma oscillations and grid firing fields is greatly increased with moderate noise. This effect is attributable to noise desynchronizing the epileptic-like states present in noise-free networks. Our results suggest independent roles for gamma oscillations and grid firing during cognitive processing and highlight the beneficial effect of noise on neural computation.

III-97. Puzzle imaging the brain: Using large-scale dimensionality reduction algorithms for localization

Joshua Glaser J-GLASER@U.NORTHWESTERN.EDU
Konrad Kording KK@NORTHWESTERN.EDU

Northwestern University

Many neural properties, including activity, connectivity, and genetics, are critical in understanding neural systems. Understanding the spatial relationship between neurons (and their properties) is an important step towards understanding their interplay. Therefore, there is an acute interest in developing neural imaging techniques that scale. In typical imaging applications, one strives to leave the sample in one piece, or if necessary to cut it into a small number of pieces. However, a radically different approach, “Puzzle Imaging,” would be to cut the sample into a huge number of small pieces and later puzzle the pieces back together. Small pieces are attractive because they can often be efficiently analyzed. For example, DNA sequencing works very efficiently on small volumes. Many researchers are currently working on encoding neural properties (connectivity, activity, and genetics) in DNA. Thus, puzzle imaging would allow efficiently localizing these neural properties across the brain. Puzzle imaging would require efficiently solving a very large puzzle. Molecular techniques may also hold the key to making the puzzle problem easier. Molecular barcoding promises to allow a very large number of markers (“colors”), which would greatly simplify the problem of asking which puzzle pieces border one another. Ultimately, puzzle imaging is a dimensionality reduction problem, a mapping of a huge number of pieces and their properties into three-dimensional space. We describe two concrete neuroscience examples in which puzzle imaging could be beneficial: (1) “Voxel Puzzling,” in which a relatively high-resolution 3-dimensional brain map is reproduced by giving DNA barcodes to neurons; and (2) “Connectomics Puzzling,” in which neural connections are used to recover neural locations. We develop two dimensionality reduction algorithms that would allow large-scale puzzle imaging, and use these algorithms to demonstrate the capabilities of puzzle imaging in the above two examples.

III-98. A spiking neuron that learns to tell the future

Johanni Brea[1] JOHANNIBREA@GMAIL.COM
Alexisz Gaal[2] GAAL.ALEXISZ@GOOGLEMAIL.COM
Robert Urbanczik[1] URBANCZIK@PYL.UNIBE.CH
Walter Senn[1] SENN@PYL.UNIBE.CH

[1]University of Bern
[2]New York University

Animals can learn to make predictions, as e.g. in weather prediction, associating the sound of a bell with food delivery in classical conditioning, or predicting the reward of following a certain strategy in a given environment. Here we propose a biologically plausible learning mechanism that allows a neuron to learn to make predictions on a behavioral timescale. To learn predictions – like bell -> food – the neuron receives the following input: sensory signals from the cue (bell) evoke a spatiotemporal activity pattern at dendritic synapses and the sensory input from the teaching signal (food) reaches synapses proximal to the soma. The dendritic synaptic strengths are altered such that sensory input from the cue leads to postsynaptic spiking prior to the arrival of the teaching signal, thereby predicting the teaching signal. In functional terms, plasticity in dendritic synapses is driven by the discrepancy between the local dendritic potential and the future discounted somatic firing, which depends on the dendritic potential and the teaching input. In contrast to more abstract models of supervised learning, we model the dynamics of the neuron and its environment in physical time. The plasticity rule is a form of spike timing dependent plasticity in which a presynaptic spike followed by a postsynaptic spike leads to potentiation. Even if the plasticity window has a width of 30 milliseconds, associations on the time scale of seconds can be learned. We illustrate the model with an example of classical trace conditioning. In reinforcement learning terminology the plasticity rule is a Monte Carlo method for value estimation; its relation to a biological implementation of temporal difference learning is discussed.

III-99. Characterizing memory formation and storage at the mesoscopic scale

Patrick Lawlor[1] PATRICK-LAWLOR@NORTHWESTERN.EDU
Shoai Hattori[2,1] SHATTORI@BU.EDU
Craig Weiss[1] CWEISS@NORTHWESTERN.EDU
John Disterhoft[1] JDISTERHOFT@NORTHWESTERN.EDU
Konrad Kording[1] KK@NORTHWESTERN.EDU

[1]Northwestern University
[2]Boston University

Memory formation and storage are distributed across multiple brain regions. As such, between-region connectivity changes are central in leading theories of memory, including the Standard Consolidation Theory (SCT), and the Multiple Trace Theory (MTT). According to SCT, for example, memories are gradually transferred from temporary hippocampal-cortical connections to more permanent cortico-cortical connections. But despite the importance of circuit-level changes, most previous work has examined either microscopic changes in individual cells, or macroscopic changes in large brain areas using lesions or fMRI. Little work has examined the neural mechanisms of memory at the mesoscale, where these theories' predictions should be readily measurable. This study aims to characterize connectivity changes over the course of memory formation by integrating electrophysiology and machine learning. We recorded from the rabbit during acquisition and retention of trace eyeblink conditioning, a model of associative learning. We simultaneously recorded single units from three brain regions implicated in memory formation and storage: dorsal CA1 hippocampus, anterior thalamus, and medial prefrontal cortex (mPFC). To estimate functional connectivity, we modeled neural activity using the Generalized Linear Model framework. We find that although neural activity predicts behavior and can be accurately modeled, functional connectivity remains relatively static over the course of learning. Hippocampal neurons exhibit strong mutual functional connectivity throughout learning, but most other types of functional connectivity are relatively rare, considerably weaker, and

do not change systematically with learning. Possible interpretations of our results include: 1) connectivity changes are weak and/or highly distributed, making them difficult to detect at the single-neuron level; 2) unmeasured brain regions may mediate information transfer; and 3) learning may be mediated by changes in timing/synchronization rather than simply increased connectivity strength. This study nevertheless demonstrates the feasibility of modeling memory-related circuit activity and highlights important questions for memory at the systems level.

III-100. The dynamics of growth cone morphology

Geoffrey Goodhill G.GOODHILL@UQ.EDU.AU
Richard Faville RICHARD.FAVILLE@GMAIL.COM
Daniel Sutherland DANIEL.SUTHERLAND@UQCONNECT.EDU.AU
Brendan Bicknell BRENDAN.BICKNELL@GMAIL.COM
Andrew Thompson A.THOMPSON4@UQ.EDU.AU
Zac Pujic Z.PUJIC@UQ.EDU.AU
Biao Sun SUNPO32@HOTMAIL.COM
Elizabeth Kita BEKITA@GMAIL.COM
Ethan Scott ETHAN.SCOTT@UQ.EDU.AU

The University of Queensland

Normal brain function depends on the development of appropriate networks of connections between neurons. A critical role in guiding axons to their targets during neural development is played by neuronal growth cones. These have a complex and rapidly changing morphology, but a quantitative understanding of this morphology, its dynamics, and how these are related to growth cone movement, is lacking. Here we use eigenshape analysis (principal components analysis in shape space) in conjunction with the Bayesian Information Criterion to uncover the set of 5-6 basic shape modes that capture the most variance in growth cone form. By analysing how the projections of growth cones onto these principal modes evolve in time we found that growth cone shape oscillates with a mean period of 30 mins. The variability of oscillation periods and strengths between different growth cones was correlated with their forward movement, such that growth cones with strong, fast shape oscillations tended to extend faster. A simple computational model of growth cone shape dynamics based on dynamic mictotubule instability was able to quantitatively reproduce both the mean and variance of oscillation periods seen experimentally, suggesting that the principal driver of growth cone shape oscillations may be intrinsic periodicity in cytoskeletal rearrangements.

III-101. Synchrony between and within populations of Purkinje cells and basket cells

Andrei Khilkevich KHILKEVICH@UTEXAS.EDU
Hunter Halverson HUNTER.HALVERSON@GMAIL.COM
Michael Mauk MIKE@MAIL.CLM.UTEXAS.EDU

University of Texas at Austin

Much of what is known about computation in the cerebellum is based on recordings from individual Purkinje cells (PCs), the sole output of the cerebellar cortex. Sets of PCs that converge onto common follower cells and that receive input from related climbing fibers are organized in parasagittal stripes. We obtained simultaneous recordings from multiple PCs and cerebellar cortex interneurons to determine the degree and type of synchrony present along a parasagittal stripe. These recordings were obtained with chronically implanted tetrode drives during delay eyelid conditioning training with different interstimulus intervals (ISIs). Simultaneously recorded PCs showed highly correlated learned decreases in activity, which were coupled with the onsets of conditioned eyelid responses (CRs). However, we did not observe spike-to-spike synchrony in PCs firing. In addition, we

investigated whether there is synchrony between PCs and other types of neurons in the cerebellar cortex. The identity of each cell was determined based on statistical properties of their baseline spike trains, using a published algorithm. A subset of single-units was identified as basket/stellate cells. Within this category, some putative basket cells (BCs) showed bursts of activity tightly time-locked to the time of CR. Simultaneously recorded BCs showed correlated increases in activity during the expression of CRs, while simultaneously recorded pairs of PCs and BCs showed robust inverse changes in activities during CRs. These observations are consistent with a recent study demonstrating that PCs collateral axons synapse onto neighboring BCs, complementing a known notion of BCs inhibiting PCs along the parasagittal stripe. In addition, our analysis showed that changes in PCs firing in general preceded ones in BCs. These data provide evidence that the connectivity between BCs and PCs makes the activity of each PC in a parasagittal stripe to influence neighboring PCs to fire in a similar way, thus producing synchrony of PCs activity.

III-102. Characterizing asynchronous activity in networks of excitatory and inhibitory spiking neurons

Srdjan Ostojic SRDJAN.OSTOJIC@ENS.FR

Ecole Normale Superieure

Networks of excitatory and inhibitory neurons form the basic computational units in the mammalian cortex. Within the dominant paradigm, neurons in such networks encode and process information by asynchronously emitting action potentials. It has recently been argued that excitatory-inhibitory networks exhibit two qualitatively different types of asynchronous activity depending on synaptic coupling strength, but further characterization of the two types of activity is needed. Using numerical simulations, we show that above a critical value of synaptic coupling, the activity of sparsely connected networks of integrate-and-fire neurons becomes strongly sensitive to parameters such as the refractory period and synaptic delays. This stands in strong contrast with classical asynchronous activity in which the refractory period and delay play a marginal role. More specifically, mean firing rates are independent of the refractory period at low coupling, but highly increase as the refractory period is decreased at strong coupling, although the firing rates remain well below saturation. A similar sensitivity is seen in higher order statistics. This qualitative change in behavior can be understood in terms of the recently proposed heterogeneous instability of the classical asynchronous state. The corresponding extended mean-field theory reproduces the dependence on refractory period, and produces predictions for the location of the transition. These predictions are in agreement with simulations in part of the parameter space, but not all. We discuss the reasons for the observed quantitative deviations, other limitations of the theory and possible alternative explanations.

III-103. Periodic forcing of stabilized E-I networks: Nonlinear resonance curves and dynamics

Romain Veltz[1] ROMAIN.VELTZ@INRIA.FR
Terrence J Sejnowski[2] TERRY@SALK.EDU

[1]Inria Sophia Antipolis Méditerranée
[2]Salk Institute for Biological Studies

Inhibition stabilized networks (ISNs) are neural architectures [Tsodyks et al., 1997] with strong positive feedback among pyramidal neurons balanced by strong negative feedback from inhibitory interneurons, a circuit element found in the hippocampus and the primary visual cortex [Ozeki et al., 2009]. In their working regime, ISNs produce damped oscillations in the gamma-range 30-80Hz in response to inputs to the inhibitory population. In order to understand the properties of interconnected ISNs, we studied periodic forcing of ISNs using a rate model. Our study is also motivated by the driving of inhibitory populations using optogenetics [Cardin et al., 2009]. We first study the periodic responses with the same frequency as the forcing, independently of their

stability, also called phase-locked (PL). The resonance curve is the amplitude of the PL response as function of the forcing frequency. We show that the resonance curve presents several localized peaks and characterize the scaling of the peaks as function of the forcing frequency. We then study the stability of the PL solutions [Gambaudo, J. M., 1985], to assess what is observable, and the properties of the quasiperiodic responses that can be produced by the network. More particular, periodically forced ISNs respond with (possibly multi-stable) PL activity whereas networks with sustained intrinsic oscillations respond more dynamically to periodic inputs with quasiperiodic response. Hence, the dynamics are surprisingly rich and phase effects alone do not adequately describe the network response. For example, the fact that the network responds with its intrinsic frequency when forced at twice its intrinsic frequency is important and often overlooked [Cardin et al., 2009, Ostojic et al., 2009]. This strengthens the importance of phase-amplitude coupling as opposed to phase-phase coupling in providing multiple frequencies for multiplexing and routing information.

III-104. Does topography play a functional role in decoding spatial information?

Lilach Avitan[1] L.AVITAN@UQ.EDU.AU
Zac Pujic[1] Z.PUJIC@UQ.EDU.AU
Philip Dyer[1] PHILIP.DYER1@UQCONNECT.EDU.AU
Nicholas Hughes[2] NICHOLAS.HUGHES1@UQCONNECT.EDU.AU
Ethan Scott[1] ETHAN.SCOTT@UQ.EDU.AU
Geoffrey Goodhill[1] G.GOODHILL@UQ.EDU.AU

[1]The University of Queensland
[2]The Queensland Brain Institute

Topographic maps are a very common organizing principle of neural wiring manifested throughout the nervous system. However it is still unclear whether preservation of spatial information is essential for neural coding. To address this we performed non-invasive functional imaging of visually-evoked activity in the zebrafish optic tectum. A spot stimulus was presented at 3 nearby positions in the visual field of zebrafish larvae, while performing confocal calcium imaging of tectal neurons loaded with fluorescent calcium indicator (OGB-AM1). All stimuli evoked broad population responses in the tectum, and there was strong overlap between the responses for each spot position. We then compared two methods for decoding the stimulus position from the pattern of activity. The first method relied purely on the topographic location of the population response in the tectum, and achieved a decoding accuracy of 50-70%. The second method ignored topographic location in the tectum, and instead used maximum likelihood (ML) decoding with no information about tectal cell locations. This method achieved a decoding accuracy of 90-99%. Since behavioural experiments by others show that this level of spatial discrimination is not a challenging task for zebrafish larvae, this suggests that the topography of the map in zebrafish tectum is unlikely to be the primary means by which fine-scale spatial information is decoded. Thus, the common assumption that topographic representations are necessary for sensory decoding may be incorrect.

III-105. Towards a computational model of dyslexia

Sagi Jaffe-Dax[1] SAGI.JAFFE@MAIL.HUJI.AC.IL
Ofri Raviv[1] OFRI.RAVIV@GMAIL.COM
Nori Jacoby[1,2] NORI.VIOLA@GMAIL.COM
Yonatan Loewenstein[1] YONATAN.LOEWENSTEIN@MAIL.HUJI.AC.IL
Merav Ahissar[1] MSMERAVA@GMAIL.COM

[1]The Hebrew University of Jerusalem
[2]Massachusetts Institute of Technology

Dyslexics are diagnosed for their poor reading skills. Yet, they characteristically also suffer from poor verbal memory, and often from poor auditory skills. This combined profile was previously explained in broad cognitive terms. We now hypothesize that Dyslexia can be understood computationally as a deficit in integrating prior information with noisy observations. To test this hypothesis we analyzed performance in an auditory discrimination task using a two-parameter computational model. One parameter captures the internal noise in representing the current event and the other captures the impact of recently acquired prior information. We found that Dyslexics' perceptual deficit can be accounted for by inadequate adjustment of these components: low weighting of their implicit memory in relation to their internal noise. Using ERP measurements we found evidence for Dyslexics' deficient automatic integration of experiment's statistics. Taken together, these results suggest that Dyslexia can be understood as a well-defined computational deficit.

III-106. Evidence for directionality in orbitofrontal local field potentials

Erin Rich ERIN.RICH@BERKELEY.EDU
Jonathan Wallis WALLIS@BERKELEY.EDU

University of California, Berkeley

Neurons in the orbitofrontal cortex (OFC) are modulated by stimulus value. A standard view holds that OFC computes values by combining highly processed sensory inputs with information such as context and memories, to compute a subjective and temporally specific outcome prediction used to guide behavior. Because sensory afferents enter OFC posteriorly, this framework suggests that the flow of information should be from posterior toward anterior. To assess this, we recorded single neurons and local field potentials (LFPs) from multiple sites spanning 11mm of the macaque OFC. Subjects performed a reward preference task, in which they chose between picture stimuli predicting rewards of different types (primary versus secondary) and different amounts, and consistently chose pictures that predicted the larger reward regardless of reward type. During recording, we assessed neural responses to forced-choice trials, in which subjects were shown only one of these reward-predicting pictures, and received the corresponding outcome. Among single neurons, 33% encoded the predicted reward value, regardless of reward type, and 14% encoded the value of only one type of reward. In contrast, reward value but not reward type could be decoded from LFPs, with the best decoding in theta (4-8 Hz) and high-gamma (HG) frequencies (70-200 Hz). Further analysis revealed that HG amplitudes on 61% of single electrodes encoded value within 500ms following stimulus onset. Of these, 92% had higher HG amplitudes when the subject expected larger rewards. Interestingly, this encoding appeared earliest at the most anterior OFC sites, and progressively later at more posterior electrodes. Therefore value information is represented at the level of the LFP, and appears to progress in an anterior to posterior direction, opposite of the directionality predicted by the standard view of OFC processing. Instead, our results may be evidence of a top-down flow of value information within OFC.

Author Index