COSYNE 2016

cosyne.org

MAIN MEETING

Salt Lake City, UT

Feb 25–28

Program Summary

Thursday, 25 February

4.00p	**Registration opens**
5.00p	**Welcome reception**
5.45p	**Opening remarks**
6.00p	**Session 1: Engineering neural circuits**
	Invited speakers: Xiao-Jing Wang, Blaise Agüera y Arcas
8.00p	**Poster Session I**

Friday, 26 February

7.30a	**Breakfast**
8.30a	**Session 2: Memory and temporal integration**
	Invited speaker: Mark Goldman; 3 accepted talks
10.30a	**Session 3: Network dynamics**
	5 accepted talks
11.45p	**Lunch break**
2.00p	**Session 4: Human computation**
	Invited speakers: Paul Smolensky, Edward Chang; 1 accepted talk
4.15p	**Session 5: Dissecting cortical circuits**
	Invited speaker: Sonja Hofer; 2 accepted talks
5.30p	**Dinner break**
7.30p	**Poster Session II**

Saturday, 27 February

7.30a	**Breakfast**
8.30a	**Session 6: Visual processing**
	Invited speaker: Greg DeAngelis; 3 accepted talks
10.30a	**Session 7: Sensorimotor integration**
	Invited speaker: Mala Murthy; 2 accepted talks
11.45p	**Lunch break**
2.00p	**Session 8: Attention and action selection**
	Invited speaker: Marisa Carrasco; 4 accepted talks
4.15p	**Session 9: Motor control**
	Invited speaker: Reza Shadmehr; 2 accepted talks
5.30p	**Dinner break**
7.30p	**Poster Session III**

Sunday, 28 February

7.30a	**Breakfast**
8.30a	**Session 10: Great expectations**
	Invited speaker: Peggy Series; 3 accepted talks
10.30a	**Session 11: Generating perceptions**
	Invited speaker: Leslie Vosshall; 2 accepted talks
11.45p	**Lunch break**
2.00p	**Session 12: Neural control of behavior**
	Invited speaker: Richard Mooney; 3 accepted talks

Poster Session Topics

Topic Area	**Session I** *Thursday*	**Session II** *Friday*	**Session III** *Saturday*
Bayesian models, optimality	1–3	1–2	1–2
Behavior, psychophysics	4–9	3–9	3–8
Cognition: attention, memory	10–14	10–12	9–11
Cognition: decision making, perception	15–21	13–20	12–20
Computational modeling	22–26	21–26	21–26
Learning: memory, plasticity	27–31	27–31	27–30
Learning: reward, reinforcement	32–38	32–36	31–36
Motor systems	39–46	37–45	37–45
Network dynamics	47–56	46–55	46–55
Oscillations, rhythms	57–59	56–58	56–58
Population coding/decoding	60–69	59–67	59–68
Sensory processing: vision	70–82	68–81	69–81
Sensory processing: audition	83–88	82–88	82–87
Sensory processing: other	89–91	89–92	88–91
Single-neuron computation	92–96	93–96	92–96
Techniques, methods, algorithms	97–105	97–105	97–105


electrophysiology3
STATES OF AMERICA
ONE DIME
ecube
v1.00

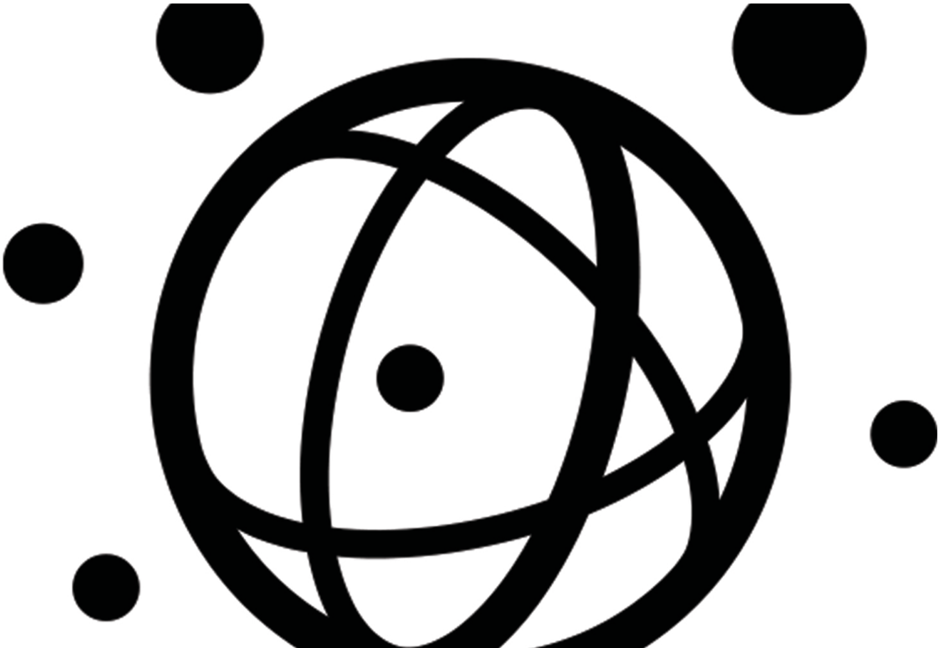

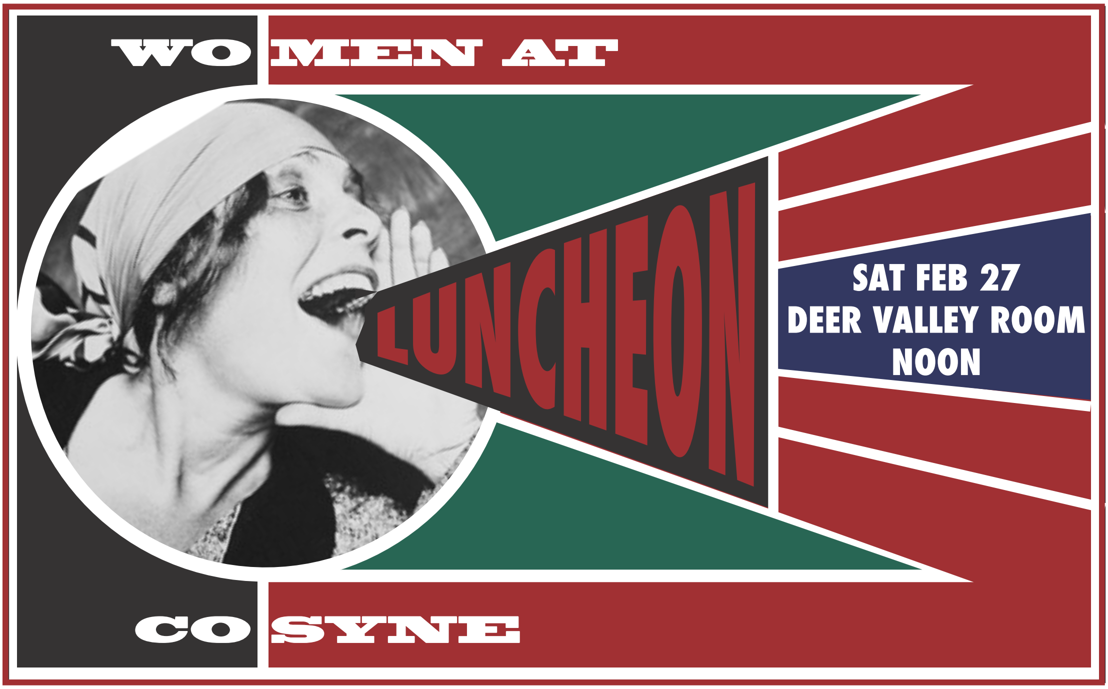

Attendance is limited to the first 100 registrants.

Please go to http://womenatcosyne.eventzilla.net/web/event?eventid=2138836343 to register.

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 2016 Leadership

Organizing Committee

General Chairs

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

Program Chairs

Megan Carey (Champalimaud), Emilio Salinas (Wake Forest University)

Workshop Chairs

Claudia Clopath (Imperial College), Alfonso Renart (Champalimaud)

Undergraduate Travel Chairs

Jill O'Reilly (University of Oxford), Robert Wilson (University of Arizona)

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

Megan Carey (Champalimaud), co-chair
Emilio Salinas (Wake Forest University), co-chair
Alex Huk (University of Texas at Austin)
Andrea Hasenstaub (University of California, San Francisco)
Arianna Maffei (Stony Brook University)
Bing Brunton (University of Washington)
Brent Doiron (University of Pittsburgh)
Bruno Averbeck (NIH)
Damon Clark (Yale University)
Daniel O'Connor (Johns Hopkins University)
Daniel Butts (University of Maryland)
Eric Shea-Brown (University of Washington)
Jessica Cardin (Yale University)
Joseph Paton (Champalimaud)
Johannes Burge (University of Pennsylvania)
Kathy Nagel (New York University)
Kechen Zhang (Johns Hopkins University)
Lindsey Glickfeld (Duke University)
Long Ding (University of Pennsylvania)
Marta Zlatic (Janelia Farm Research Campus)
Mehrdad Jazayeri (Massachusetts Institute of Technology)
Na Ji (Janelia Farm Research Campus)
Paul Miller (Brandeis University)
Sam Sober (Emory University)
Tatjana Tchumatchenko (Max Planck Institute)

Reviewers

Thomas Akam, Athena Akrami, Justin Ales, Ana Amador, Jan Antolik, Behtash Babadi, Omri Barak, Renata Batista-Brito, Rudy Behnia, Daniel Bendor, Gordon Berman, James Bigelow, Bart Borghuis, Ethan Bromberg-Martin, Michael Buice, Jianhua Cang, Rishidev Chaudhuri, Eugenia Chiappe, Yuwei Cui, Rodica Curtu, Thaddeus Czuba, Stephen David, Christopher DiMattina, Anita Disney, Takahiro Doi, Felice Dunn, Mark Eldridge, Tatiana Engel, Sean Escola, Christopher Fetsch, Greg Field, James Fitzgerald, Alfredo Fontanini, Makoto Fukushima, Stefano Fusi, Sunil Gandhi, David Gire, Jesse Goldberg, Mark Goldman, Tim Hanks, Bryan Hansen, Kathryn Hedrick, Ann Hermundstad, Michael Higley, Mark Histed, Elizabeth Hong, Bryan Hooks, Marc Howard, Chengcheng Huang, Wentao Huang, Arvind Iyer, James Jeanne, Kresimir Josic, Ingmar Kanitscheider, Sung Soo Kim, Alex Kwan, Giancarlo La Camera, Phil Larimer, Kenneth Latimer, Brian Lau, Mark Laubach, Arthur Leblois, Andrew Leifer, Ashok Litwin-Kumar, Eran Lottem, Kenway Louie, Pedro Maia, Brian Malone, Jean-Baptiste Masson, James McFarland, Joseph McGuire, Melchi Michel, Jason Middleton, Stefan Mihalas, Joseph Monaco, Masayoshi Murakami, Gabe Murphy, Kae Nakamura, Matthew Nassar, Ian Nauhaus, Gabriel Ocker, Tomoko Ohyama, Srdjan Ostojic, Il Memming Park, Simon Peron, Leopoldo Petreanu, Elizabeth Phillips, Dietmar Plenz, Cindy Poo, Barry Richmond, Dima Rinberg, Sandro Romani, Melissa Runfeldt, Aravi Samuel, Takashi Sato, Casey Schneider-Mizell, Benjamin Scholl, Johannes Seelig, Matt Smear, Spencer Smith, Adam Snyder, Alireza Soltani, Armen Stepanyants, Jordan Taylor, Gelsy Torres-Oviedo, John Tuthill, Marcel van Gerven, Stephen Van Hooser, Clement Vinauger, Martin Vinck, Andrew Welchman, Corette Wierenga, Robert Wilson, Hongdian Yang, Byron Yu, Jianing Yu, Edward Zagha, Joel Zylberberg

Conference Support

Administrative Support, Registration, Hotels

Denise Acton, Cosyne

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 2016.

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

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 2016 recipients are:

Christopher Wilson, Madineh Sedigh-Sarvestani, Carsen Stringer, Anqi Wu, Fanny Cazettes, Chunyu Duan, Katherine Morrison, Morgan Taylor, Julia Veit, Alexandra Constantinescu, Jennifer Blackwell, Laureline Logiaco, Genevieve Yang, Stefania Sarno, Mor Ben-Tov, Ashesh Dhawale, Sakyasingha Dasgupta, Satohiro Tajima, Srinivas Gorur-Shandilya, Kiyohito Iigaya, Jacques Bourg, Emin Orhan, Conor Dempsey, Kaushik J Lakshminarasimhan, David Alex Mely

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 2016 recipients are:

Kathleen Martin, Thomas Roseberry, Gregory Corder, Ulises Pereira, Gregory Telian, Bradley Voytek, Noga Weiss Mosheiff, Sofia Soares, Angela Langdon, Ji Hyun Bak, Wei-Mien Mendy Hsu, Neta Ravid Tannenbaum, Xin Ru (Nancy) Wang, Francesca Mastrogiuseppe, Kathryn Tabor, Corinna Lorenz, Tiberiu Tesileanu, Jordan Guerguiev, Vivek Athalye, James Cooke, SangWook Lee, Richard Lange, Xinping Li, Richard Gao

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 2016 Cosyne Mentors are listed below, each followed by their mentee:

Blake Richards and Annik Yalnizyan-Carson, Ilya Nemenman and Caroline Holmes, Damon Clark and Emilio Salazar Cardozo, Michael DeWeese and Bernal Jimenez

Cosyne Undergraduate Travel Grant Program

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

The 2016 recipients are:

Maricarmen Hernandez, Brenda Vega, Sun Mi Kim, Ana Gomez del Campo, Valerie Zhao, Sara Golidy, Oihane Horno, Mehma Singh, Lucy Lai, Bhadra Chembukave, Toren Wallengren, Ariel Herbert-Voss, Kanupriya Gupta

Program

Note: Printed copies of this document do not contain the abstracts; they can be downloaded at:
http://cosyne.org/c/index.php?title=Cosyne2016_Program
Institutions listed in the program are the primary affiliation of the first author. For the complete list, please consult the abstracts.

Thursday, 25 February

4.00p **Registration opens**

5.00p **Welcome reception**

5.45p **Opening remarks**

Session 1: Engineering neural circuits

(Chair: Emilio Salinas)

Blaise leads a team at Google focusing on Machine Intelligence for mobile devices—including both basic research and new products. His group works extensively with deep neural nets for machine perception, distributed learning, and agents, as well as collaborating with academic institutions on connectomics research. Until 2014 he was a Distinguished Engineer at Microsoft, where he worked in a variety of roles, from inventor to strategist, and led teams with strengths in interaction design, prototyping, computer vision and machine vision, augmented reality, wearable computing and graphics. Blaise has given TED talks on Seadragon and Photosynth (2007, 2012) and Bing Maps (2010). In 2008, he was awarded MIT's prestigious TR35 ("35 under 35").

8.00p **Poster Session I**

Friday, 26 February

7.30a **Continental breakfast**

Session 2: Memory and temporal integration

(Chair: Alireza Soltani)

10.00a **Coffee break**

Session 3: Network dynamics

(Chair: Tatyana Sharpee)

11.45p **Lunch break**

Session 4: Human computation

(Chair: Stephanie Jones)

3.45p **Coffee break**

Session 5: Dissecting cortical circuits

(Chair: Joe Paton)

Saturday, 27 February

Session 6: Visual processing

(Chair: Eero Simoncelli)

Session 7: Sensorimotor integration

(Chair: Jennifer Raymond)

Session 8: Attention and action selection

(Chair: Richard Krauzlis)

Session 9: Motor control

(Chair: Jose Carmena)

Sunday, 28 February

Session 10: Great expectations

(Chair: Leslie Osborne)

9.45a	Temporal expectations in reward prediction: "what" and "when" computations in the basal ganglia A. Langdon, Y. Takahashi, G. Schoenbaum, Y. Niv, Princeton University	46
10.00a	**Coffee break**	

Session 11: Generating perceptions

(Chair: Damon Clark)

10.30a	Solving the stimulus-percept problem for olfaction Leslie Vosshall, Rockefeller University **(invited)**	31
11.15a	A computational role for cortical feedback in odor detection from complex scenes G. Otazu, P. Masset, D. F. Albeanu, Cold Spring Harbor Laboratory	47
11.30a	In vivo characterization of synaptic reliability at the thalamocortical synapse of cat V1 M. Sedigh-Sarvestani, M. M. Taylor, L. A. Palmer, D. Contreras, University of Pennsylvania	47
11.45p	**Lunch break**	

Session 12: Neural control of behavior

(Chair: Long Ding)

2.00p	A synaptic and circuit switch for control of flexible behavior K. Kuchibhotla, J. Gill, E. Papadoyannis, T. Hindmarsh Sten, R. Froemke, New York University	48
2.15p	Ring attractor dynamics in the Drosophila central brain S. S. Kim, H. Rouault, J. Seelig, S. Druckmann, V. Jayaraman, Janelia Farm Research Campus	48
2.30p	Neural ensemble dynamics underlying a long-term associative fear memory B. Grewe, J. Gruendemann, J. Marshall, J. Parker, J. Z. Li, A. Luethi, M. Schnitzer, Stanford University	49
2.45p	Motor circuits for listening and learning Richard Mooney, Duke University **(invited)**	31
3.30p	**Closing remarks**	

Poster Session I — 7:30 pm Thursday 25 February

Poster Session II 7:30 pm Friday 26 February

Poster Session III 7:30 pm Saturday 27 February

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

This PDF can be downloaded at: http://cosyne.org/cosyne16/Cosyne2016_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. Building a large-scale brain model: a dynamics- and function-based approach

Xiao-Jing Wang XJWANG@NYU.EDU
New York University

Decision-making, working memory and other cognitive functions involve many brain regions that are interconnected through feedback loops. With the advance of new experimental tools ranging from connectomics to massive recording and imaging from animals performing controlled behavioral tasks, time is ripe to investigate the dynamical inner working of large-scale brain circuits. One may justifiably wonder “what does it mean to build a whole brain model?” or mischievously ask “so what’s the difference with the European Human Brain Project?” In this talk, I will present our recent work on large-scale brain circuit modeling, centered around formulating new questions that do not arise urgently in studies of local circuits but must be confronted with a large brain system. We have developed large-scale cortex modeling of macaque and mice, using recently published databases of directed and weighted connectivity. We found that, by taking into account quantitative heterogeneity across cortical areas, such a large network naturally gives rise to a hierarchy of timescales: early sensory areas respond rapidly to an external input and the response decays away immediately after stimulus offset (appropriate for sensory processing), whereas association areas higher in the brain hierarchy are capable of integrating inputs over a long time and exhibit persistent activity (suitable for decision-making and working memory). Slower association areas have a disproportionate impact on the global brain dynamics, suggesting a reevaluation of the analysis of functional connectivity by taking into account interareal heterogeneity. This model can be expanded to incorporate a laminar structure of the cortex, and to investigate frequency-dependent feedforward versus feedback neural signaling. Moreover, in such a complex brain system, routing of information between areas must be flexibly gated according to behavioral demands. For instance, when you try to read a book in a noisy café, it is desirable for your brain to “gate in” visual information while “gating out” auditory inputs. We propose such a gating mechanism with a disinhibitory circuit motif implemented by several subtypes of inhibitory neurons. Our results establish a circuit mechanism for a hierarchy of “temporal receptive windows”, and show that distributed processing in the brain is heterogeneous and gated by nonlinear microcircuit motifs that are under the control of executive signals. The model provides a computational platform for investigating dynamics and functions of the large-scale brain.

T-2. Engineering neural-ish systems

Blaise Agüera y Arcas BLAISEA@GOOGLE.COM
Google

Neural nets are finally coming of age. Modern convolutional and recurrent neural networks are sweeping the field as practical solutions to hard machine perception problems. They are also increasingly addressing other tasks that brains can do but have been historically hard for computers: description, translation, chat and natural language, robotics, and even synthesis (hallucination and neural net art). While these artificial systems are not usually designed to be biologically plausible in their implementation details, they are decidedly more "neural" than previous approaches to AI or feature-engineered machine learning. There are, for example, meaningful direct comparisons to be made between electrophysiological recordings from brains and activity from artificial deep neural nets under the same stimuli. These comparisons have begun to allow us to gain greater intuition about both the natural and the artificial systems.

T-3. Microcircuits for memory storage and neural integration

Mark Goldman MSGOLDMAN@UCDAVIS.EDU
University of California, Davis

Memory over time scales of seconds to tens of seconds is thought to be maintained by patterns of neural activity that are triggered by a memorized stimulus and persist long after the stimulus is turned off. This presents a challenge to current models of memory-storing mechanisms, because the typical time scales associated with cellular and synaptic dynamics are two orders of magnitude smaller than this. While such long time scales can easily be achieved by bistable processes that toggle like a flip-flop between a baseline and elevated-activity state, many neuronal systems have been observed experimentally to be capable of maintaining a continuum of stable states. For example, in neural integrator networks involved in the accumulation of evidence for decision making and in motor control, individual neurons have been recorded whose activity reflects the mathematical integral of their inputs; in the absence of input, these neurons sustain activity at a level proportional to the running total of their inputs. In this talk, I will present and compare different biologically motivated circuit motifs for the accumulation and storage of signals in short-term memory.

T-4. Combinatorial representations and symbolic computation with distributed neural activation patterns

Paul Smolensky[1,2] SMOLENSKY@JHU.EDU
[1]Johns Hopkins University
[2]Aix-Marseille University

A fundamental property of higher human cognition is that information processed at any given moment consists of multiple elements arranged in complex combinations. I will spell out the challenge that this poses for brain theory: How can distributed neural activation patterns encode such structured combinatorial information? I will propose a mathematical solution and argue the adequacy of this neural encoding scheme for the complex computations required for human cognition. Extant and potential future evidence that the brain actually deploys such a coding scheme will be discussed.

T-5. Functional organization of human auditory speech cortex

Edward Chang CHANGED@NEUROSURG.UCSF.EDU
University of California, San Francisco

A unique and defining trait of human behavior is our ability to communicate through speech. While much of this processing has been localized to the peri-sylvian cortex, the fundamental organizational principles of the neural circuits within these areas are largely unknown. In this talk, I will present new results from our research to detailed the functional organization of the human higher-order auditory cortex, known as Wernicke's area. I will focus on how neural populations in the superior temporal lobe encode acoustic-phonetic representations of speech, and also how they integrate influences of linguistic context to achieve perceptual robustness.

T-6. Putting vision into context: visual processing during behavior

Sonja Hofer SONJA.HOFER@UNIBAS.CH
University of Basel

Our perception of the environment relies on information flow from the sensory organs to the brain. However, perception is also highly dependent on the context in which a given stimulus occurs, such as expectations, intentions and actions. Such contextual information can strongly modulate visual signals and influence how visual stimuli are perceived and interpreted. The talk will cover recent work in the lab on how such contextual signals emerge during learning, how they are integrated with the visual information and by which pathways they might be conveyed.

T-7. Neural computations underlying perception of depth from motion

Greg DeAngelis GDEANGELIS@MAIL.CVS.ROCHESTER.EDU
University of Rochester

A fundamental computational problem for the visual system involves constructing a three-dimensional (3D) representation of the world from the 2D images formed on the retina of each eye. While the neural mechanisms of depth perception from binocular cues have been well studied, relatively little has been known about how the brain computes depth from the relative motion between objects (motion parallax) that arises during translation of the observer. Importantly, motion parallax by itself can be ambiguous with regard to depth sign (near or far relative to the point being fixated), and theory shows that the brain must combine visual image motion with signals related to rotation of the eye relative to the scene, in order to compute depth from motion parallax. I will describe a series of studies in which we have identified the neural computations by which visual motion signals are combined with both non-visual and visual signals regarding eye rotation to generate a representation of depth based on motion parallax.

T-8. Singing on the fly: Sensorimotor integration and acoustic communication in Drosophila

Mala Murthy MMURTHY@PRINCETON.EDU
Princeton University

Social interactions require continually adjusting behavior in response to sensory feedback. For example, when having a conversation, sensory cues from our partner (e.g., sounds or facial expressions) affect our speech patterns in real time. Our speech signals, in turn, are the sensory cues that modify our partner's actions. What are the underlying computations and neural mechanisms that govern these interactions? To address these questions, my lab studies the acoustic communication system of Drosophila. To our advantage, the fly nervous system is relatively simple, with a wealth of genetic tools to interrogate it. Importantly, Drosophila acoustic behaviors are highly quantifiable and robust. During courtship, males produce time-varying songs via wing vibration, while females arbitrate mating decisions. We discovered that, rather than being a stereotyped fixed action sequence, male song structure and intensity are continually sculpted by interactions with the female, over timescales ranging from tens of milliseconds to minutes—and we are mapping the underlying circuits and computations. We have also developed methods to relate song representations in the female brain to changes in her behavior, across multiple timescales. Our focus on natural acoustic signals, either as the output of the male nervous system or as the input to the female nervous system, provides a powerful, quantitative handle for studying the basic building blocks of communication.

T-9. Attention and early vision

Marisa Carrasco MARISA.CARRASCO@NYU.EDU
New York University

Attention allows us to select relevant sensory information for preferential processing. I will discuss effects of attention on early visual processes. I will present psychophysical and fMRI studies regarding the effects of endogenous (voluntary) and exogenous (involuntary) covert attention—the selective processing of visual information without eye movements—on the perception of basic visual dimensions. Specifically, I will showing how contrast sensitivity increases at the attended location (or for the attended features) at the expense of reduced sensitivity at unattended locations (or for the unattended features), and discuss these results in reference to a normalization model of attention.

T-10. Encoding of action by Purkinje cells of the cerebellum

Reza Shadmehr SHADMEHR@JHU.EDU
Johns Hopkins University

Execution of accurate movements depends critically on the cerebellum, suggesting that Purkinje cells (P-cells) may predict motion of the moving body part. Yet, this encoding has remained a long-standing puzzle. For example, in case of saccadic eye movements, P-cells show little consistent modulation with respect to amplitude or direction, and critically, their discharge lasts longer than duration of a saccade. Here, we analyzed P-cell discharge in the oculomotor vermis of behaving monkeys and found neurons that increased or decreased their activity during saccades. We estimated the combined effect of these two populations via their projections on the output nucleus and uncovered a simple-spike population response that precisely predicted the real-time motion of the eye. When we organized the P-cells according to each cell's complex-spike directional tuning, the simple-spike population response predicted both the real-time speed and direction of saccade multiplicatively via a gain-field.

This suggests that the cerebellum predicts the real-time motion of the eye during saccades via the combined inputs of P-cells onto individual nucleus neurons. A gain-field encoding of simple spikes emerges if the P-cells that project onto a nucleus neuron are not selected at random, but share a common complex-spike property.

T-11. Mental illness as a deficit in probabilistic inference

Peggy Series PSERIES@INF.ED.AC.UK
University of Edinburgh

A growing idea in computational neuroscience is that perception and cognition can be successfully described in terms of Bayesian inference: the nervous system would maintain and update internal probabilistic models that serve to interpret the world and guide our actions. This approach is increasingly recognised to also be of interest to Psychiatry. Mental illness could correspond to the brain trying to interpret the world through distorted internal models, or incorrectly combining such internal models with sensory information.

I will describe work pursued in my lab that aims at uncovering such internal models, using behavioural experiments and computational methods. In health, we are particularly interested in clarifying how prior beliefs affect perception and decision-making, how long they take to build up or be unlearned, how complex they can be, and how they can inform us on the type of computations and learning that the brain performs. In mental illness, we are interested in understanding whether/how the machinery of probabilistic inference could be impaired, and/or relies on the use of distorted priors. I will describe recent results relevant to the study of Schizophrenia and Depression.

T-12. Solving the stimulus-percept problem for olfaction

Leslie Vosshall[1,2] LESLIE@MAIL.ROCKEFELLER.EDU
[1]Rockefeller University
[2]Howard Hughes Medical Institute

Olfaction is the least understood of the senses. Despite many centuries of thought about how smell “works,” we still have no way to predict what a molecule will smell like. Given a chemical structure, the only way to determine what olfactory percept it gives is to smell it. This stimulus-percept problem was solved long ago for color vision and tone hearing. Solving the stimulus-percept problem for olfaction is more complicated because odor stimuli do not vary along a predictable axis. We do not know how many different smells exist, and we do not know how odors are arranged in perceptual space. The perception of an odor stimulus can be described in a number of distinct ways: how intense is it, how pleasant is it, what does it remind me of, how can I describe it, how similar is it to another smell, can I tell these two closely related smells apart? I will discuss the current understanding of the stimulus-percept problem, what data may allow us to solve it, and what implications the problem has for how brain circuits represent smells.

T-13. Motor circuits for listening and learning

Richard Mooney MOONEY@NEURO.DUKE.EDU
Duke University

A classical view is that sensation begets movement, which implies a flow of information from sensory to motor regions of the brain. However, the flow of information from motor to sensory regions also is critical to perception and motor learning. Motor to auditory interactions are especially important to hearing, because they can serve to suppress responsiveness to self-generated sounds while boosting sensitivity to unexpected sounds arising from sources in the environment. Motor to auditory interactions are also thought to facilitate the learning of sound-generating behaviors, such as speech and musicianship, by conveying a motor-related prediction of the auditory consequences of the ensuing movement. The brain can then compare this predictive signal to movement-related auditory feedback to generate an error signal that can guide motor learning. The circuit, cellular and synaptic mechanisms that mediate such motor to auditory interactions in the vertebrate brain are poorly understood. I will discuss research from our group in both songbirds and mice that use a variety of methods, including in vivo cellular imaging, electrophysiology, viral gene transfer and optogenetics, to map, monitor and manipulate motor to auditory interactions important to auditory processing and vocal motor learning. These studies reveal features of central brain organization that are likely to be relevant to human auditory function, especially in the context of speech perception and learning.

T-14. Circuit principles of memory-based behavioral choice

Marta Zlatic[1] ZLATICM@JANELIA.HHMI.ORG
Kathi Eichler[1,2] EICHLERK@JANELIA.HHMI.ORG
Feng Li LIF@JANELIA.HHMI.ORG
Claire Eschbach ESCHBACHC@JANELIA.HHMI.ORG
Akira Fushiki[1,2] FUSHIKIA@JANELIA.HHMI.ORG
Jim Truman TRUMANJ@JANELIA.HHMI.ORG
Bertram Gerber BERTRAM.GERBER@LIN-MAGDEBURG.DE
Aravinthan Samuel[3] ADTSAMUEL@GMAIL.COM
Marc Gershow GERSHOW@GMAIL.COM
Albert Cardona[1,2] CARDONAA@JANELIA.HHMI.ORG
Andreas Thum[4] ANDREAS.THUM@UNI-KONSTANZ.DE

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

A single nervous system can generate many distinct behaviors. Choosing which behavior to generate based on sensory inputs and previous experience is crucial for the survival of any organism. To understand the circuit principles by which experience-driven behavioral choices are made it is essential to determine the architecture of networks that mediate these functions with synaptic resolution, and determine the causal relationships between the structural motifs and function. We use the genetically tractable insect model system, the Drosophila larva, with a 10,000-neuron nervous system and uniquely identifiable neurons to combine three levels of analysis: i) circuit mapping using electron microscopy (EM); ii) physiological measurements of neural activity and iii) neural manipulation in freely behaving animals to dissect the logic of memory-based behavioral choice. In an EM volume that spans the entire nervous system we have reconstructed a complete wiring diagram of the higher order parallel fiber system for associative learning in the insect brain, the Mushroom Body (MB), including the pathways from the conditioned and unconditioned sensory neurons to the MB, and from the MB to distinct types of descending neurons from the brain that mediate distinct aspects of the conditioned response. The revealed a slew of interesting microcircuit motifs in the higher order parallel fiber system, and we are modeling the potential roles in different kinds of learning. We also identified multiple pathways by which the MB could modulate innate odor responses downstream—affecting both distinct types of command-like neurons as well as the sensitivity to sensory cues following conditioning. Using calcium imaging and optogenetic manipulation of individual MB input and output neurons we elucidated the logic of punishment and reward encoding by the ensemble of dopaminergic MB input neurons and the logic by which the MB interacts with pathways for innate responses to odor. Understanding how memories and learned behaviors are encoded throughout the larval nervous system may provide direct insight into these processes in the larger insect and vertebrate nervous systems.

T-15. Midbrain dopamine neurons directly modulate duration judgments

Sofia Soares SOFIA.SOARES@NEURO.FCHAMPALIMAUD.ORG
Bassam Atallah BASSAM.ATALLAH@NEURO.FCHAMPALIMAUD.ORG
Alessandro Braga ALESSANDRO.BRAGA@NEURO.FCHAMPALIMAUD.ORG
Thiago Gouvea THIAGO.GOUVEA@NEURO.FCHAMPALIMAUD.ORG
Tiago Monteiro TIAGO.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG
Joe Paton JOE.PATON@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Centre for the Unknown

Adaptive behavior involves doing the right thing at the right time. Yet time judgments are subjective, and depend on factors including arousal, movement state, and attention. The basal ganglia (BG) are critical for timing behavior on the scale of seconds. Yet, little is known about how distinct components of BG circuitry contribute to time judgments. Here we probe the role of a critical component of the BG, midbrain dopaminergic (DAergic) neurons, that have been previously implicated in interval timing. We measured and manipulated the activity of DAergic neurons in mice judging time in an interval discrimination task. DAergic neural activity was recorded using GCaMP6f together with fiber photometry. We observed reliable responses tied to various task events such as interval onset, interval offset and reward delivery. Strikingly, larger DAergic responses predicted short judgments about near boundary intervals. This was true in both hemispheres and independent of whether judgments were correct. To test whether changes in DAergic activity were sufficient to cause a systematic change in timing judgments, on a subset of trials we photo-activated DA neurons using channelrhodopsin-2 (ChR2). This manipulation resulted in a horizontal shift in the psychometric curve towards short choices on stimulated trials, consistent with the photometry data. Notably, this effect was not accompanied by any significant change in response times, and did not last beyond the stimulated trial, arguing against value and learning-related interpretations, respectively. Taken together, our results are consistent with the notion that activity in midbrain DAergic neurons directly modulates animals' internal representation of elapsed time. These data suggest a novel role for midbrain DAergic neurons, demonstrating one route through which time judgments might be modulated by behavioral state.

T-16. Neural integration underlying a time-compensated sun compass in the Monarch butterfly

Eli Shlizerman SHLIZEE@UW.EDU

University of Washington

Eastern North American monarch butterflies use a time-compensated sun compass during their fall migration. The compass is an active control mechanism that adjusts the migrant's flight to the southerly direction. While the antennal genetic circadian clock and the azimuth of the sun are instrumental for proper function of the compass mechanism, it is unclear how these signals are represented at the neuronal level and how compass neurons are wired to determine necessary flight corrections. To address these questions, we examined the shape and neuroanatomy of the compound eye to construct a receptive field model that encodes the azimuthal position of the sun and then derived conditions for the neural circuit that integrates sensory inputs to provide flight direction control. A particular combination of clock and azimuth encoding, based on two anti-phase signals, satisfies these conditions and provides a model for a time-compensated sun compass. To verify the model we compared activity predicted by our model with recordings from intrinsic neurons interconnecting the two lateral accessory lobes, the presumed location of the sun compass output circuit, and found a striking similarity in responses. Using the model to simulate flight trajectories and comparing these with trajectories of tethered butterflies flying in a flight simulator we find similar distribution characteristics and dynamics, in particular ease-in dynamics, the possibility for rotations and indication of a separatrix. The model of the neural circuit, the conditions that we derive and recordings that we have obtained demonstrate that an integration mechanism based on matching anti-phase neural signals can guarantee robust trajectories reaching the southwest bearing from almost any heading at any time of day.

T-17. Efficient coding of a dynamic trajectory predicts non-uniform allocation of grid cells to modules

Noga Weiss Mosheiff NOGA.WEISS@MAIL.HUJI.AC.IL
Haggai Agmon HAGGAI.AGMON@MAIL.HUJI.AC.IL
Avraham Moriel AVRAHAM.MORIEL@MAIL.HUJI.AC.IL
Yoram Burak YORAM.BURAK@ELSC.HUJI.AC.IL

The Hebrew University of Jerusalem

Recent experiments established that grid cells in the entorhinal cortex are functionally organized in discrete modules with uniform grid spacing. The grid spacings approximately form a geometric series. The experimental data suggests that the number of cells decreases sharply with grid spacing, in marked disagreement with existing theories. Here, we postulate that the entorhinal cortex is adapted to represent a dynamic quantity (the trajectory of the animal in space), while taking into account the temporal statistics of this variable. We first develop a theory for efficient coding of a variable that dynamically follows the statistics of a simple random walk. A central prediction of the theory is that module neuron population sizes should sharply decrease with the increase of grid spacing, in agreement with the trend seen in the experimental data. Another prediction is that the ratio between grid spacings should approach a constant value in the modules with the smallest spacing, which is consistent with experimental data and with previously proposed models. Next, we identify a remarkably simple, near optimal scheme for readout of the grid cell code by neural circuitry, in which model place cells linearly sum inputs from grid cells, using an exponential temporal kernel, whose decay time depends on the spacing of the presynaptic grid cell. The simple readout scheme can be optimized for trajectories that deviate in their temporal statistics from a simple random walk. As an extreme case we consider motion at constant velocity in an unknown direction. Even for such motion, we obtain from the optimization qualitatively similar results as for random walk statistics. Thus, we propose that the sharp decrease in module population sizes, with increase of the grid spacing, is an outcome of the efficient coding hypothesis, if the dynamic nature of motion in space is taken into account.

T-18. Slow adaptation facilitates excitation-inhibition balance in the presence of structural heterogeneity

Itamar Landau[1] ITAMAR.LANDAU@MAIL.HUJI.AC.IL
Robert Egger ROBERT.EGGER@TUEBINGEN.MPG.DE
Vincent J Dercksen[2] DERCKSEN@ZIB.DE
Marcel Oberlaneder[3] MARCEL.OBERLAENDER@TUEBINGEN.MPG.DE
Haim Sompolinsky[1] HAIM@FIZ.HUJI.AC.IL

[1]The Hebrew University of Jerusalem
[2]Zuse-Institut Berlin
[3]Max Planck Institute for Biological Cybernetics

Traditional analysis of cortical network dynamics has most commonly treated simple random graph structure. We present anatomy-based estimates of connectivity statistics within local circuits of the rat barrel cortex. We observe that the in-degree from within a single cell-type is significantly broader than expected from a simple random graph, and that in-degrees from different cell-types are substantially correlated. Simulations of LIF networks with connectivity structure from data reveal unbalanced networks in which a large majority of neurons are totally silent and those cells that fire do so at high rates and with temporal regularity. Analytically, we study a generic model of networks with broad and correlated in-degrees. We show that in general, networks with broad in-degree distributions cannot maintain the dynamic balance of excitation and inhibition, and the dynamics are mean driven. Correlated in-degrees can mitigate this effect and enable the recovery of balance and fluctuation-driven irregular firing. We analytically determine the structural boundary for maintaining balance and find that the connectivity estimates from anatomy fall outside of the balance regime. We present a novel dynamical state in which a slow adaptation current corrects for the structural imbalance locally and facilitates the global emergence of balance.

We find that moderate adaptation currents, of the same order of magnitude as those observed in both excitatory and inhibitory cortical neurons, are sufficient to significantly mitigate the impact of structural imbalance. Finally, we explore the relationship between connectivity and activity that emerges in the adaptation-facilitated balanced state. Population activity is primarily distributed along a single dimension of the underlying connectivity structure, and this dimension is determined by the excitation-inhibition balance rather than by the structure of network connectivity.

T-19. Efficient signal processing in random networks that generate variability

Sakyasingha Dasgupta[1] SAKYASINGHA.DASGUPTA@RIKEN.JP
Isao Nishikawa[2] GOLDBERGVA@GMAIL.COM
Kazuyuki Aihara[2] AIHARA@SAT.T.U-TOKYO.AC.JP
Taro Toyoizumi TARO.TOYOIZUMI@BRAIN.RIKEN.JP

[1]RIKEN Brain Science Institute
[2]The University of Tokyo

The Source of cortical variability and its influence on signal processing remain an open question. We address the latter, by studying two types of randomly connected networks of quadratic integrate-and-fire neurons with balanced excitation-inhibition that produce irregular spontaneous activity patterns: (a) a deterministic network with strong synaptic interactions that actively generates variability by chaotic dynamics (internal noise) and (b) a stochastic network that has weak synaptic interactions but receives noisy input (external noise), e.g. by stochastic vesicle releases. These networks of spiking neurons are analytically tractable in the limit of a large network-size and channel-time-constant. Despite the difference in their sources of variability, spontaneous activity patterns of these two models are indistinguishable unless majority of neurons are simultaneously recorded. We characterize the network behavior with dynamic mean field analysis and reveal a single-parameter family that allows interpolation between the two networks, sharing nearly identical spontaneous activity. Despite the close similarity in the spontaneous activity, the two networks exhibit remarkably different sensitivity to external stimuli. Input to the former network reverberates internally and can be successfully read out over long time. Contrarily, input to the latter network rapidly decays and can be read out only for short time. The difference between the two networks is further enhanced if input synapses undergo activity-dependent plasticity, producing significant difference in the ability to decode external input from neural activity. We show that, this difference naturally leads to distinct performance of the two networks to integrate spatio-temporally distinct signals from multiple sources. Unlike its stochastic counterpart, the deterministic chaotic network activity can serve as a reservoir to perform near optimal Bayesian integration and Monte-Carlo sampling from the posterior distribution. We describe implications of the differences between deterministic and stochastic neural computation on population coding and neural plasticity.

T-20. Long-term stability in behaviorally relevant neural circuit dynamics

Ashesh Dhawale[1] DHAWALE@FAS.HARVARD.EDU
Rajesh Poddar[1] RAJESH.PODDAR@GMAIL.COM
Evi Kopelowitz[1] KOPELOWITZ@FAS.HARVARD.EDU
Valentin Normand[2] VALENTIN.NORMAND@ENS.FR
Steffen Wolff[1] STEFFENWOLFF@FAS.HARVARD.EDU
Bence Olveczky[1] OLVECZKY@FAS.HARVARD.EDU

[1]Harvard University
[2]Ecole Normale Superieure

The goal of systems neuroscience is to understand how neural activity generates behavior. A traditional experimental approach is to record from neural populations at times when subjects perform designated tasks. While

such intermittent recordings provide brief ‘snapshots’ of task-related neural dynamics, they fail to address how neural activity is modulated outside of task context, or how it changes across behavioral states and time. Addressing these questions requires tracking the activity of neuronal populations continuously over weeks and months in behaving animals. Such experiments face significant technical challenges, including processing vast amounts of neural and behavioral data. We present a low-cost, fully automated experimental platform that allows neural activity and behavior to be recorded continuously over several months. The large datasets we generate are analyzed using a novel processing pipeline, where the key step is a spike-sorting algorithm that allows for automatic identification and tracking of single units in terabyte-sized datasets even when units have non-stationary spike-waveforms. We used our system to record activity in large populations of single neurons in motor cortex and striatum, often holding units for several weeks. In conjunction with the neural recordings, high-resolution behavioral data was acquired using high-speed cameras and head-mounted 3-axis accelerometers, which, together with local field potentials, were used to identify epochs of sleep, rest, grooming, feeding, and to track and quantify movement kinematics during execution of a skilled motor task. We found that average firing rates and correlation structure in neuronal populations were stable across many days, even as they varied substantially across different behavioral states in a single day. Additionally, we found the motor representations of skilled behaviors to be remarkably stable at the single unit level, even over month-long timescales. These results demonstrate that neural circuits can maintain distinct task representations with long-term stability at the level of single neurons.

T-21. Stability and drift of motor sequencing in the songbird HVC

William Liberti[1] BLIBERTI@BU.EDU
Jeffery Markowitz[1] JMARKOW@BU.EDU
Daniel Leman[1] DPLEMAN@BU.EDU
Derek Liberti[2] DLIBERTI@BU.EDU
L Nathan Perkins[1] LNP@BU.EDU
Carlos Lois[3] CLOIS@CALTECH.EDU
Darrell Kotton[2] DEREKLIBERTI@GMAIL.COM
Timothy Gardner[1] TIMOTHYG@BU.EDU

[1]Boston University
[2]Boston Medical Center
[3]California Institute of Technology

Motor skills can be maintained for long timescales- for days, years or even decades. However, little is known about the mechanistic basis of this stability. Some propose that while motor skills can remain stable for years, the individual neurons controlling them may significantly change their firing properties over the course of hours. Others contend that the tuning of individual neurons is as stable as the motor skill itself. Merging these two viewpoints, the central hypothesis of this presentation is that the brain encodes learned behaviors on two distinct levels- a mesoscopic level that is highly stable, and a microscopic level in which single neurons can be influenced by a history of reward or punishment. We examine the question of motor coding stability in one of the most stable of all animal behaviors: birdsong. The extreme precision and long-term stability of song structure presents a unique opportunity to observe how motor memories are maintained at the network level. Using genetically encoded calcium indicators and miniature head-mounted microscopes, we observe that firing patterns of excitatory projection neurons in the premotor cortical area HVC drift over a timescale of days. In contrast, electrophysiological recordings reveal both multiunit firing patterns and local field potentials in this region persist for weeks to months. These ensemble patterns persist after peripheral nerve damage, revealing that sensory-motor correspondence is not required to maintain the stability of the underlying neural ensemble. Single neuron recording of inhibitory interneurons reveal stable patterns over the same timescale, suggesting that this cell type largely contributes to the ensemble pattern. These observations suggest a mesoscopic principle of motor stability in HVC: stable behavior is supported by stable inhibitory dynamics, combined with fine scale exploration in the firing patterns of excitatory neurons.

T-22. Optimal synaptic strategies for different timescales of memory

Subhaneil Lahiri SULAHIRI@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU

Stanford University

Systems consolidation suggests that different brain regions are specialized to store memories over different timescales. Similarly, synapses mediating memory have highly complex, diverse molecular signaling pathways, varying across brain regions. This suggests the possibility that synaptic diversity across brain regions may be adapted for different timescales of memory storage. We are left with the fundamental question: what type of molecular synaptic dynamics is suitable for storing memory over any given timescale? To address this, we systematically analyze an extremely broad class of models where synaptic plasticity is implemented by stochastic transitions between internal functional states of a sub-synaptic molecular network. Such models are essential in navigating stringent tradeoffs between learning and remembering, known as the stability-plasticity dilemma. Previous work (e.g. [Fusi, Drew, Abbott, 2005]) analyzed this tradeoff in models with only one topological structure of transitions between sub-synaptic states. This leaves open the nature of this tradeoff over all possible sub-synaptic networks. Rather than analyze one model, we analyze the space of all possible models and elucidate principles that determine how sub-synaptic network structure can be ideally adapted to the time-scale over which memories are stored. We find that as this timescale increases, initially synapses are forced to grow a chain of internal states with rapid transitions, while at even longer timescales, synapses are further forced to exhibit slow stochastic transitions. We also discuss the design of synaptic physiology experiments to test our theoretical predictions. We find conventional methods for probing synaptic plasticity cannot discern the relevant synaptic dynamics. Instead we propose new classes of experiments and data-analysis procedures in which more subtle protocols for probing plasticity can yield systems identification of the synaptic dynamics so crucial for storing memories at a particular time-scale.

T-23. Transient competitive amplification during states of cortical activation

Jacques Bourg[1] JACQUES.BOURG@NEURO.FCHAMPALIMAUD.ORG
Nivaldo Vasconcelos[2] NIVALDO.VASCONCELOS@NEURO.FCHAMPALIMAUD.ORG
Peter Bartho[3] BARTHO@KOKI.HU
Alfonso Renart[1] ALFONSO.RENART@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]University of Minho
[3]Hungarian Academy of Sciences

A number of studies have shown that variability in cortical circuits is low-dimensional and structured around temporal fluctuations in population firing rate (PFR; e.g., Okun et al., Nature, 2015). Fluctuations in PFR, however, are more prominent in periods of cortical inactivation (CI, typical of quiescent periods during wakefulness), being quite small during cortical activation (CA—typical of attentive wakefulness; Renart et al., Science, 2010). We investigated the state-dependence of cortical dynamics using population recordings during urethane anesthesia, which is known to produce periods of both CA and CI. We show that spontaneous-activity fluctuations in the rat auditory and somatosensory cortices during states of CA are also low-dimensional, but that the structure of this variability is not related to changes in PFR. PCA revealed that structured fluctuations are confined to an approximately one-dimensional 'competitive axis', and that they take place while the PFR remains approximately constant. We show that competition is strongest between physically close neurons and that the competitive structure is approximately preserved over time-scales ranging from a few tens to many hundreds of ms. What mechanisms could generate this type of dynamics? Competitive amplification is typically generated by placing a neural circuit close to a pitchfork bifurcation. We identified a novel circuit motif capable of generating competition in the absence of critical slowing-down through non-normal amplification (Murphy & Miller, Neuron, 2009), which we denote as transient competitive amplification (TCA). We show that TCA requires asymmetric connectivity between competing populations, and that it naturally generates idiosyncratic but robust features of the data, such as a systematic time lag in the negative correlations between competing neurons. Our work uncovers a novel dynamical regime of cortical circuits, and suggests that non-normal amplification is important for this type of dynamics.

T-24. Inhibitory control of correlated cortical variability

Carsen Stringer[1] CARSEN.STRINGER@GMAIL.COM
Michael Okun[2] M.OKUN@UCL.AC.UK
Peter Bartho[3] BARTHO@KOKI.HU
Kenneth Harris[2] KENNETH.HARRIS@UCL.AC.UK
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK
Peter Latham[1] PEL@GATSBY.UCL.AC.UK
Nicholas Lesica[2] N.LESICA@UCL.AC.UK
Marius Pachitariu[2] MARIUS10P@GMAIL.COM

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University College London
[3]Hungarian Academy of Sciences

The firing rate of a neuron in the mammalian cortex fluctuates in coordination with the activity of its neighbors. The nature of this relationship varies across behavioral states, and affects the reliability of the neuron's sensory representation. We reproduced the rich range of statistical structures present in multi-neuron recordings using different operating regimes of a single deterministic spiking network model with intrinsic variability. In our model, inhibition controls population-wide variability, a result confirmed in experiments. Further, we show that spiking network models explain aspects of neuronal variability that cannot be accounted for in external noise models. We fit the parameters of a spiking network model to the statistics of spontaneous and driven activity from 60 different electrophysiology datasets of 20-100 neurons recorded in the sensory cortices of mice, rats and gerbils, using novel computational techniques. We used graphics processing units to simulate networks of 512 spiking neurons at 10000x real-time and simulated a million different parameter sets. We computed summary statistics, which capture the timescales and correlations of the activity, for each network and each dataset and chose the fits based on these statistics. The model successfully fit the diversity of timescales and correlations present in the neuronal activity. To investigate the consequences for coding, we drove the networks with external stimulus inputs. Evoked responses were least correlated in networks with the largest inhibitory-to-excitatory firing rate ratios. High inhibition abolished population fluctuations and enhanced coding properties. In agreement with the model, we found that high levels of narrow-spiking inhibition did indeed correlate with reduced noise correlations in sound-evoked recordings from auditory cortex. In the neuronal recordings, noise correlations decreased as mean stimulus-evoked population firing rates increased. This result was reproduced in the spiking network model but not in a model with external noise that was also fit to the datasets.

T-25. Bayesian multisensory integration by dendrites

Joao Sacramento SACRAMENTO@PYL.UNIBE.CH
Walter Senn SENN@PYL.UNIBE.CH

University of Bern

Animals must make sense of an uncertain world to survive. A widespread strategy to improve perceptual accuracy is to integrate multiple independent sensory measurements of a common stimulus, whenever possible. A simple but successful approach is to assume that the measurements are corrupted by Gaussian noise. Cast as a Bayesian inference problem, the optimal solution is then to average the inputs, while weighting each by its own

precision, the reciprocal of the noise variance; this combination rule seems to pervade behaviour. How does the brain carry out such computation? Unlike previous theories, we propose that this might be possible already at the single-neuron level. We study a conductance-based, multi-compartmental model neuron and find that it is naturally equipped to implement Bayesian estimation. Dendritic branches receive clustered excitatory and inhibitory inputs from specific sensory modalities; a strong coupling of the dendritic tree to the soma ensures proper averaging across branches. We develop a set of rules to study the steady-state compartmental voltages that can be applied to arbitrarily shaped dendritic trees. Our analysis reveals that the interplay of branch-specific excitation and inhibition defines the measurement while the sum of both determines its precision. Thus, input stimulation strength enables one modality to suppress others without affecting the encoded stimulus. We argue for such co-modulation of excitation and inhibition as a general and plausible dynamic weighting mechanism, capable of operating on the fast perceptual timescale. Finally, we discuss the relevance of our findings in the light of learning and plasticity in the presence of uncertainty.

T-26. Functional clustering of synaptic inputs in primary visual cortex

Daniel Wilson DAN.WILSON@MPFI.ORG
David Whitney DAVID.WHITNEY@MPFI.ORG
Benjamin Scholl BEN.SCHOLL@MPFI.ORG
David Fitzpatrick DAVID.FITZPATRICK@MPFI.ORG

Max Planck Florida Institute for Neuroscience

The response properties of cortical neurons depend on the proper integration of activity supplied by thousands of axon terminals forming synapses within their dendritic fields. Fundamental to understanding synaptic integration is establishing how functionally defined synaptic inputs are spatially arranged within the dendritic field. It has been suggested that clustering of functionally similar inputs along dendritic branches could contribute to response selectivity, potentially through nonlinear dendritic mechanisms. The spatial organization of these synaptic inputs can play an important role in how single neurons integrate information, but the role of these input patterns in cortical computation is poorly understood. While clustered synaptic activity patterns have been observed in hippocampal neurons in vitro and in cortical neurons in vivo, evidence for a strong clustering of functionally similar synaptic inputs contributing to a neuron's selectivity remains elusive. Here we demonstrate strong functional clustering of synaptic inputs with similar orientation preference for individual neurons. Functional synaptic clusters are correlated with orientation tuning sharpness and input-output nonlinearities evident in comparing summed spine tuning and somatic tuning. These nonlinearities cannot be accounted for by spike threshold alone and, consistent with a dendritic nonlinearity, dendrites with more co-tuned spine clusters show greater rates of local dendritic calcium events. Our results suggest that functional clustering of synaptic inputs plays a significant role in shaping the selective responses of cortical neurons through local dendritic processing.

T-27. The role of target selective descending neurons in dragonfly prey selection during free behavior

Huai-Ti Lin[1,2] LINH10@JANELIA.HHMI.ORG
Anthony Leonardo[1,2] LEONARDOA@JANELIA.HHMI.ORG

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

Dragonflies are excellent aerial predators that capture flying insects on the wing. Their success hinges on both the prey pursuit strategy and the preparatory head tracking that centers the target in the visual fovea to extract prey information [1]. We have recently found that the dragonfly selectively pursues prey satisfying a certain range of angular size-speed ratio without true estimation of prey distance and speed. With a simple behavioral model,

we can show that such a heuristic rule effectively implements prey selection that screens out uncatchable prey. What's the neural substrate for this heuristic? A class of target selective descending neurons (TSDNs) carry prey angular velocity information from the visual system in the head to the body [2] and innervate various motor units such as wings and neck. To understand whether they play a role in prey selection, we first characterized the target joint angular size-speed receptive field (or 2D tuning map) of two specific TSDNs in an immobilized animal. The TSDN target tuning map from the immobilized preparation differs with the prey selection behavioral heuristic in that it favors larger prey sizes that are normally rejected by the animal. To assay whether TSDNs have the same tuning map during behavior, we used an RF powered telemetry backpack [3] to record from the same TSDNs in the dragonfly during foraging. Preliminary data on the TSDN speed-size tuning from the freely behaving animal show good match to the prey selection heuristic. This suggests that 1) the TSDN tuning is modulated by behavioral states, and prior publications on TSDNs recorded under restrained conditions may not directly reflect the functional tuning during behavior. 2) the TSDN tuning could directly implement or be involved in the prey selection heuristic that we observed from the behavior.

T-28. Corollary discharge mediates sensorimotor integration in a C. elegans neural circuit for thermotaxis

Ni Ji[1] SEAURCHINJINNI@GMAIL.COM
Vivek Venkatachalam[1] VIVEK@PHYSICS.HARVARD.EDU
Maria Lim[2] LIM@LUNENFELD.CA
Taizo Kawano[2] T3KAWANO@GMAIL.COM
Christopher Clark[3] CHRISTOPHER.CLARK@UMASSMED.EDU
Hillary Rogers[1] HILLARYRODGERS@GMAIL.COM
Mark Alkema[3] MARK.ALKEMA@UMASSMED.EDU
Mei Zhen[2] ZHEN@LUNENFELLD.CA
Aravinthan Samuel[1] ADTSAMUEL@GMAIL.COM

[1]Harvard University
[2]Samuel Lunenfeld Hospital
[3]University of Massachusetts

For an animal to efficiently navigate its habitat, the nervous system must constantly coordinate sensory input with the animal's current motor state to determine its future behavior. While the phenomenon of sensorimotor integration has been described in a variety of nervous systems, the circuit mechanism and the functional impact of sensorimotor integration on behavior remain elusive. In this study, we investigated the neural circuitry underlying C. elegans thermotaxis behavior to identify circuit mechanisms of sensorimotor integration. To understand how thermosensory input is processed by the neural network, we simultaneously imaged the calcium activity of sensory, inter-, and motor neurons across the thermosensory circuitry in animals that freely navigating a spatiotemporal temperature gradient. By analyzing the correlation of neural activity across the circuit in relation to the sensory input and the motor output, we identified a group of upper layer interneurons that simultaneously encode both the thermosensory stimuli and the motor state of the animal. Through cell ablation and genetic and optogenetic manipulations, we confirmed that the robust motor encoding by the upper layer interneurons depends on a corollary discharge pathway stemming from the motor circuitry. When corollary discharge is genetically perturbed, the activity of the upper layer interneurons no longer couples with that of the motor circuit, but instead encodes solely the sensory input. To assess the contribution of the corollary discharge pathway to thermotaxis behavior, we quantitatively analyzed the thermotaxis behavior in the presence and absence of corollary discharge. We observed that animals with disrupted corollary discharge fail to sustain locomotory states beneficial to the thermotaxis behavior. To provide mechanistic insights into this behavior phenotype, we built a dynamical systems model to demonstrate how sensorimotor integration through corollary discharge can contribute to a robust biased-random-walk strategy, which ultimately ensures robust thermotaxis behavior.

T-29. CA1 firing fields represent an abstract coordinate during non-spatial navigation

Dmitriy Aronov DARONOV@PRINCETON.EDU
Rhino Nevers RNEVERS@PRINCETON.EDU
David W Tank DWTANK@PRINCETON.EDU

Princeton University

Hippocampal neurons fire in discrete place fields within a spatial environment. However, the hippocampus is critical not only for spatial navigation, but for a variety of memory-guided behaviors. Accordingly, previous studies have shown that hippocampal activity can be modulated by non-spatial features of the animal's experience, such as sensory stimuli or the amount of elapsed time. It is therefore possible that place fields are not fundamentally a framework for representing location; rather, discrete firing fields in the hippocampus may be a general framework for representing values of continuous, behaviorally-relevant variables. To test this idea, we trained rats in an acoustic virtual reality apparatus to 'navigate' on an abstract 'linear track', with the non-spatial axis defined by sound frequency. Animals pressed a joystick to change the frequency of a pure tone and were required to release it in a target zone between two frequency values. By analogy with velocity, the rate of change of frequency was proportional to the deflection of the joystick and was therefore controlled by the rat. The scaling factor between joystick deflection and velocity was varied across trials to uncouple sound frequency from elapsed time. Using tetrodes, we recorded units in CA1 of the dorsal hippocampus and found that individual cells formed discrete firing fields that, across the population, spanned the entire behavioral task. This activity sequence expanded and compressed with trial duration, indicating that individual cells were active at particular phases of the task and did not simply represent elapsed time. Our results show that a non-spatial axis can be represented by hippocampal activity in a fashion similar to the previously known representation of location. They suggest a model in which discrete firing fields are flexibly used by the hippocampal neurons to represent arbitrary, abstract spaces that support not only spatial navigation, but cognition in general.

T-30. Exploration flattens prefrontal target selectivity, enhances learning in network states and behavior

Becket Ebitz[1] REBITZ@GMAIL.COM
Eddy Albarran[2] ALBARRAN@STANFORD.EDU
Tirin Moore[2] TIRIN@STANFORD.EDU

[1]Princeton University
[2]Stanford University

In variable, uncertain environments, actors must balance the drive to maximize immediate reward ('exploitation') with periods of discovery ('exploration'), during which they can learn about the value of options. A balance between exploration and exploitation is critical for flexible decision-making, but few electrophysiological studies have examined how the brain implements these states. This is a striking omission in regions responsible for directing attention because attention plays a central, yet paradoxical role in these decision states. Selective attention is required for learning about uncertain values, yet previously-rewarded targets capture attention at the expense of alternatives. How does the brain permit learning about alternatives in the presence of these salient options? In order to address these questions, we identified monkeys' endogenous transitions between exploration and exploitation while we recorded simultaneously from populations of single neurons the frontal eye fields (FEF), a prefrontal region that controls overt and covert attention. In many laboratory tasks (and during the exploit states here), single neurons in the FEF are target-selective: their activity is enhanced when the target selected by attention is in their response field. During exploration, however, single-neuron target-selectivity was substantially reduced. At the network level, target-selectivity became stronger over time within exploitation trials, but we found little of this amplification during exploration. After transitions into exploration, target-selective network states re-developed across trials at a rate proportional to the rewards the monkeys received. However, outside of these

brief plastic periods, network states were insensitive to rewards. Thus, selective attention signals in the prefrontal cortex are reduced during exploration, but reward learning in behavior and network states is enhanced. These results suggest a model in which reduced target-selectivity permits exploratory choices, while long-hypothesized changes in learning rate during exploration may be due to other mechanisms, such as changes in the level of reward-dependent network plasticity.

T-31. Postponement of evidence accumulation in area LIP until action-selection is possible

Shushruth Shushruth[1,2] FS2478@CUMC.COLUMBIA.EDU
Michael Shadlen[1,2] SHADLEN@COLUMBIA.EDU

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

In perceptual decision-making tasks where the animal knows in advance the motor actions associated with the different decision outcomes, the choosing of the appropriate motor action is concomitant with the evaluation of the sensory evidence bearing on the choice. Animals can also make perceptual decisions without knowing the decision-outcome to motor-action mapping. In such cases, it is commonly assumed that a decision is made about the abstract property of the sensory stimulus (e.g., rightward), and the decision-outcome subsequently guides action selection. We trained a monkey to perform an abstract version of a direction discrimination task. The monkey had to decide whether the net direction of stochastic random-dot motion (RDM) was rightward or leftward and indicate its decision by making a saccade to a cyan or an yellow choice target, respectively. Crucially, the targets were not present during motion viewing but appeared 1/3 sec after, anywhere in the visual field. The animal was allowed to report its decision as soon as the targets appeared. Surprisingly, the saccadic latencies (go-RTs) were long and depended on the stimulus strength. A bounded drift-diffusion model provided a satisfactory account of both the animal's accuracy and go-RTs suggesting a coupling between the go-RTs and the sensory evidence presented beforehand. During this action-selection epoch, neural activity in the area LIP built up at a rate that depended on the strength of RDM presented earlier in the trial. These psychophysical and neurophysiological findings suggest that a bounded evidence accumulation process undergirds action selection in this task. Since the samples of evidence accumulated during action selection were informed by sensory information acquired 100s of ms in the past, the accumulation, evident in LIP, might represent some form of recapitulation of the sensory evidence stream, allowing decision making to occur in the framework of the revealed decision-outcome to motor-action mapping.

T-32. History-dependent variability in population dynamics during evidence accumulation in cortex

Ari Morcos ARIMORCOS@FAS.HARVARD.EDU
Christopher D Harvey CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

Harvard Medical School

A fundamental feature of neural processing is combining ongoing activity dynamics with stimuli from the outside world. The posterior parietal cortex (PPC) performs this combination for a variety of computations, including evidence accumulation during decision-making. However, due to previous technical limitations, it remains poorly understood how the combination of external stimuli and internal activity, and thus evidence accumulation, are represented in a population of neurons. Here we developed an evidence accumulation task for mice navigating in a virtual environment and used two-photon calcium imaging to measure the activity of populations of individual PPC neurons. The PPC population represented task-relevant features, including choice and accumulated evidence, as a code distributed across groups of neurons that for the most part had heterogeneous and variable activity patterns. Using population-level analyses based on clustering of trials with similar activity patterns, we found that population activity trajectories were highly variable across trials, even for trials with identical stimuli. This variability was not entirely due to noise; rather, it contained structure that predicted past and future activity patterns over seconds, as well as future behavioral outcomes. Information about past events, including the previous trial's choice and the sequence of past stimuli, was represented in the population activity and contributed to the trial-trial variability. Our results suggest that, in the PPC, incoming stimuli are incorporated into a distributed code that contains a signal for past events, such that variability in a stimulus's representation in part reflects an ongoing historical record. These dynamics could allow the readout of accumulated evidence for decision-making and, more generally, any combination of internal activity and incoming stimuli.

T-33. Normalization and urgency cooperate in optimal multi-alternative decisions

Satohiro Tajima SATOHIRO.TAJIMA@GMAIL.COM
Daniel Robles Llana DANIEL.ROBLESLLANA@GMAIL.COM
Jan Drugowitsch JDRUGO@GMAIL.COM
Alexandre Pouget ALEX.POUGET@GMAIL.COM

University of Geneva

In the real life, we often make decisions among multiple alternatives. While studies with binary choice paradigms have demonstrated the canonical mechanism of decision-making in simplified situations, much less is known about the computational principle of decisions with more than two options. A simple implementation of multi-alternative decisions is the 'race model (RM)', where the momentary choice preference is signaled by a competition among ramping-up neural activities, which are terminated as soon as one of them reaches a decision threshold. On the other hand, physiological recordings have suggested puzzling properties of neural dynamics that require extending standard RMs; in perceptual and value-based decision tasks, neurons often show activity normalization over the neural units (e.g., Louie et al., J. Neurosci., 2011) and a time-dependent bias input to the network (urgency signal) that urges rapid decisions (e.g., Churchland et al., Nat. Neurosci., 2008). Although some ad-hoc models have been proposed to fit neural behavior, why the nervous system requires the normalization and urgency and how they relate to each other remains poorly understood. To address these problems, we theoretically derive the normative strategies in general multi-alternative decisions, and identify the optimal stopping-rules for perceptual or value-based evidence accumulation. This reveals nonlinear and time-dependent decision-boundaries in a high-dimensional belief space, which appear intractable by nervous systems in situ. However, we find that a geometric symmetry in those decision boundaries allow the optimal strategies to be reduced to a remarkably simple neural mechanism, which is interpreted as a novel extension of RM ('urgency-constrained race model') that features a time-dependent activity-normalization controlled by the urgency signal. The model explains why the nervous system requires the activity normalization and urgency signal: they are necessary to implement optimal decisions under multi-alternative choices. The model predicts a time-dependent normalization that constrains neural population activity during decision-making.

T-34. Optogenetic dissection of descending behavioral control in Drosophila

Jessica Cande[1,2] CANDEJ@JANELIA.HHMI.ORG
Gordon J Berman[3] GORDON.BERMAN@EMORY.EDU
Shigehiro Namiki[1,2] NAMIKIS@JANELIA.HHMI.ORG
Wyatt Korff[1,2] KORFFW@JANELIA.HHMI.ORG
Gwyneth Card[1,2] CARDG@JANELIA.HHMI.ORG
Joshua Shaevitz[4] SHAEVITZ@PRINCETON.EDU
David Stern[1,2] STERND@JANELIA.HHMI.ORG

[1]Janelia Farm Research Campus
[2]Howard Hughes Medical Institute
[3]Emory University
[4]Princeton University

In most animals, the brain sends signals through local neural circuitry in the nerve chord to produce behaviors. Despite the central importance of these signals as an informational and anatomical bottleneck, little is known about how these signals are encoded at the neuronal level or how they control aspects of behavior. In insects, signals from the brain to the ventral nerve chord are carried by an estimated 350 pairs of bilaterally symmetric descending neurons (DNs). To date, only a handful of these descending neurons have known functions. In order to understand how DNs can control insect behaviors, we developed a method to identify descending interneuron function in an unbiased and systematic fashion in the model insect D. melanogaster. Using the red-shifted channelrhodopsin CsChrimson (Klapoetke, 2014), we activated neurons in a collection of ~ 200 lines, most of which individually target single neurons out of a collection of DNs that fall into 60 distinct neuro-anatomical classes. Using techniques from Berman, 2014, we created a two-dimensional behavioral space based on the underlying postural dynamics of freely moving flies with and without red light activation. In this map, stereotyped behaviors are represented by local probability density maxima, and distinct behavioral motifs are easily distinguished. We then looked for map regions that were upregulated in CsChrimson activated animals. Using this technique, we were able to assign phenotypes to 90% of the DNs in our collection. We find that (1) DN control of stereotyped behaviors appears to be modular, (2) much DN function can be correlated with neuro-anatomy, and (3) the nature of this correlation hints at centralized control of locomotory activities. These findings, which are only apparent in a dataset of this size, have wide-ranging implications for how complex signals from the brain are encoded by descending neurons.

T-35. Muscle and kinematic representations for arm and BMI control exist in orthogonal subspaces

Hagai Lalazar[1] HAGAI_LALAZAR@YAHOO.COM
Larry Abbott[1] LFABBOTT@COLUMBIA.EDU
Eilon Vaadia[2] EILON.VAADIA@ELSC.HUJI.AC.IL

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

Brain-Machine Interfaces enable high-performance 3D cursor control. However, the structure in neural activity that drives BMIs is unknown. In our experiment, monkeys previously trained using arm control, succeeded in using our BMI within 75 seconds on their very first day. Since tool learning takes much longer, we hypothesized that some components of neural activity must reflect intended cursor kinematics and must be shared between arm and BMI control. While components related to muscle activity must be different, because the arm was at rest during BMI. We tested this hypothesis by analyzing the same M1 neurons recorded while monkeys performed a target-to-target reach and hold task in blocks of arm or BMI control. Using dimensionality reduction techniques, we found that M1 population activity parcellates into three orthogonal subspaces. One subspace is occupied only during arm control but not during BMI. Single-trial decoding of EMG shows that this subspace includes neural activity for controlling the muscles. A second subspace is shared between arm and BMI control, and contains an invariant representation of cursor kinematics. During target-hold, this representation matched the geometry of the targets in Cartesian space. During reaching, neural trajectories for each type of reach are similar for arm and BMI control, and we could decode cursor trajectory from the activity in this subspace. Finally, we found a third subspace, a one-dimensional binary context subspace, where arm and BMI activity form two distinct clusters, but each overlaps for different movements. These results explain why M1 neurons correlate with both kinematic and force/muscle variables, and why both can be decoded from M1 populations. Moreover, they elucidate how neurons involved in arm control can drive the BMI when the arm is at rest, and why BMIs work quickly for rehearsed tasks. Finally, orthogonal subspaces may serve as a general mechanism for selective gating.

T-36. A dynamic Bayesian observer model reveals origins of bias and variability in path integration

Kaushik J Lakshminarasimhan[1] JKLAKSHM@BCM.EDU
Marina Petsalis[1] MARINA.PETSALIS@BCM.EDU
Gregory DeAngelis[2] GDEANGELIS@MAIL.CVS.ROCHESTER.EDU
Xaq Pitkow[3,1] XAQ@RICE.EDU
Dora Angelaki[1,3] ANGELAKI@BCM.EDU

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

That the brain performs optimal probabilistic inference has been established primarily through binary decision tasks with time-invariant stimuli. However in complex, dynamic environments where evidence is often non-stationary, optimal performance requires perfect integration of instantaneous estimates over task-relevant timescales in addition to optimal estimation at each instant. To test this, we asked human subjects to perform path integration: subjects had to use a joystick to steer themselves to a cued target location in a virtual environment devoid of any recognisable landmarks. This task challenges subjects to compute their position by integrating self-motion velocity estimates obtained from sparse optic flow. Consistent with other studies of path integration, behavioural responses were found to be systematically biased: subjects generally travel beyond the target location. Such a pattern of results has previously been attributed to leaky integration of evidence that ultimately leads to an underestimation of distance travelled. Here we considered an alternative hypothesis: bias in path integration stems from a prior favouring slower speeds that causes subjects to underestimate their travel velocity. We tested these mutually exclusive hypotheses in the framework of a Dynamic Bayesian observer model. The model computes the posterior distribution over position by integrating the posterior over velocity obtained by combining noisy sense data with a prior belief. The distinction between the two hypotheses was realised by manipulating both the prior (flat or exponential) and the nature of integration (leaky or perfect). Fitting these models to data, we found that the model with an exponential prior and perfect integration accounts well for subjects' biases and has a much higher likelihood than the leaky integration model. This suggests that humans can perform optimally even in the presence of dynamical evidence, and that behavioural inaccuracies under these conditions are more likely to be due to wrong beliefs rather than suboptimal evidence integration.

T-37. Noise correlations support a feedback model of motion prediction in V1

Till Hartmann TILL@HMS.HARVARD.EDU
Richard Born RICHARD_BORN@HMS.HARVARD.EDU

Harvard Medical School

Our brains routinely predict the trajectories of moving objects in order to, for example, catch a thrown ball. While it is clear that the motor commands sent to the muscles compensate for neural delays and anticipate the future location of an object, it is not known where along the neural pathways from vision to action that these functions begin. We thus asked whether V1 responses to drifting bars have a predictive component. To do this, we recorded neuronal activity with a multi-electrode array in V1 of a fixating monkey and compared responses under two conditions: 1) a white bar (75% contrast) drifted smoothly across the receptive fields (drifting-bar condition) or 2) an identical bar flashed at random locations along the same motion trajectory (space-time receptive field, or STRF). By convolving the STRF with the representation of the drifting bar stimulus we produced a linear model prediction of each unit's response profile. If V1 neurons respond linearly (i.e. without a predictive component), the linear model prediction should be identical to the drifting-bar response; however, we found that it significantly lagged the drifting-bar response (median difference 14 ms, $p<<0.001$, Wilcoxon signed-rank test). A second model, which included a negative feedback component, did a much better job of predicting the drifting-bar response profiles than the linear model (median improvement 11 ms, $p<<0.001$, Wilcoxon signed-rank test). The negative-feedback model predicts negative noise correlations between pairs of neurons whose receptive fields lie at adjacent, but offset, locations along the bar trajectory, but only when the drifting bar is between the receptive fields of the two neurons. We found noise correlations were dynamically modulated, revealing negative correlations at the time and place predicted by the feedback model. This signature of feedback could be generated within V1 or produced by feedback from higher areas, such as MT.

T-38. Temporal expectations in reward prediction: “what” and “when” computations in the basal ganglia

Angela Langdon[1] ALANGDON@PRINCETON.EDU
Yuji Takahashi[2] YTAKA001@UMARYLAND.EDU
Geoffrey Schoenbaum[2] GEOFFREY.SCHOENBAUM@NIH.GOV
Yael Niv[1] YAEL@PRINCETON.EDU

[1]Princeton University
[2]NIDA/NIH

Work in recent years has leveraged the computational framework of temporal-difference reinforcement learning (TDRL) to unveil the neural substrates of how we learn to predict rewards and to choose actions that will obtain them. Within this framework, dopaminergic neurons are thought to signal reward prediction errors that drive learning of reward predictions in the striatum. Importantly, these learned predictions then input back to dopamine neurons to allow the computation of prediction errors. We recorded activity of putative dopaminergic neurons in the ventral tegmental area while rats with neurotoxic (or sham) lesions of the ipsilateral ventral striatum performed a simple odor-guided choice task in which the timing or size of rewards was manipulated. Firing patterns in sham-lesioned animals were consistent with reward prediction error signals. However, dopamine neurons in the lesioned animals failed to signal reward prediction errors to changes in reward timing, while prediction errors to changes in reward size were intact. These results suggest a functional dissociation between predicting the timing (‘when’) and magnitude (‘what’) of rewards in the striatum. To account for these findings we developed a TDRL model based on a partially observable semi-Markov decision process that explicitly dissociates learning of temporal expectations from learning of magnitude. This model obviates the need to assume a temporally precise state representation over which learning occurs, instead allowing the duration of each state to be learned concurrently with the expected value of that state. We model lesions of the ventral striatum as an inability to learn precise temporal expectations, and show how this critically changes state value estimation and thus reward prediction-error signals, mimicking the experimental results. This model makes precise the role of temporal expectations in shaping reward prediction errors and suggests that the ventral striatum is critical for forming and exploiting temporal expectations during reward prediction and learning.

T-39. A computational role for cortical feedback in odor detection from complex scenes

Gonzalo Otazu GHOTAZU@GMAIL.COM
Paul Masset PMASSET@CSHL.EDU
Dinu F Albeanu ALBEANU@CSHL.EDU

Cold Spring Harbor Laboratory

Rodents use a repertoire of $\sim 1,100$ odorant receptors (ORs) to represent several orders of magnitude more olfactory objects. Despite large changes in concentration and turbulent odor flow, they readily identify weak target odors in rich sensory scenes, where several strong background odors can be present. To date, the neural mechanisms underlying this complex computation remain unknown. Here we present a novel algorithm that creates an estimate of odor input identity using as few elements (non-zero contributions) as possible from a large, previously learned dictionary representing possible odor sources. Our algorithm uses as cost function the sum of the squares of the contributions (L2 minimization) to find a sparse solution, as opposed to the widely used sum of absolute values (L1 minimization). It is implemented as a real-time predictive coding scheme, where the current estimate of the sources present, combined with the current odor input create iteratively a new estimate of the odor input. The resulting estimation error is further used to update the estimate of which sources are present. The model mirrors biologically the olfactory bulb (OB)-to-piriform cortex circuit, and assigns a critical role to cortical-bulbar feedback signals. The model predicts: 1) existence of two distinct feedback channels that differ in response polarity to odor stimulation (enhanced vs. suppressed); 2) existence of two OB output channels, one that represents the estimation error, and is suppressed by cortical feedback, and a second channel that broadcasts incoming sensory input to the cortex, and is independent of cortical feedback. We successfully cross-validated these predictions using multiphoton calcium imaging in awake head-fixed mice and monitoring: 1) corticalbulbar feedback boutons and 2) dynamics of mitral and tufted cells before and after suppression of cortical feedback. Our model represents a major advance in understanding cortical feedback, creating a computational framework that is closely corroborated by experiments.

T-40. In vivo characterization of synaptic reliability at the thalamocortical synapse of 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

A goal of sensory neuroscience is understanding the contribution of thalamic inputs to receptive fields (RFs) of layer-4 (L4) neurons. In V1, L4 RFs resemble the spatial sum of their thalamic LGN inputs, despite the fact that LGN afferents account for only $\sim 10\%$ of synapses. Because of their disproportionate contribution, it is assumed that thalamocortical (TC) synapses are robust. We used in vivo measurements combined with computational modeling to study the reliability of synaptic transmission at the TC synapse of cat V1. Using paired recordings of monosynaptically connected LGN (extracellular) and V1 (intracellular) neurons, we present the first in vivo characterization of the TC synapse in V1, including amplitude and probability of monosynaptic EPSPs. We used our observations to constrain a simple integrate-and-fire (IF) model of the TC circuit, and used it to test the effect of probabilistic synapses on L4. We combined biological and computational approaches via dynamic clamp (DC), which we used in vivo to simulate TC synapses between many visually driven LGN neurons and a L4 V1 neuron that does not receive any direct visual input. We found that TC synapses are on average not very reliable, with $44 \pm 14\%$ of LGN spikes failing to produce an EPSP in the post-synaptic L4 neuron. We report the novel use of DC to create simple-cell RFs in L4 V1, with functionally relevant properties such as orientation tuning, from spatially aligned LGN inputs. We found that probabilistic TC synapses lead to a several-fold increase in the variance and

firing rate of L4 neurons, without overtly affecting RF properties. Our findings suggest that, although TC synapses in cat V1 are not robust, probabilistic TC synapses nonetheless lead to amplification of thalamic input. Our novel use of DC may clarify the contribution of thalamic inputs to L4 neurons across sensory systems.

T-41. A synaptic and circuit switch for control of flexible behavior

Kishore Kuchibhotla KUCHIBH@GMAIL.COM
Jonathan Gill JVG219@NYU.EDU
Eleni Papadoyannis ESP304@NYU.EDU
Tom Hindmarsh Sten HINDMARSH.STEN@NYU.EDU
Robert Froemke ROBERT.FROEMKE@MED.NYU.EDU

New York University

Sensory perception and sensorimotor behaviors enable animals to interact with the external world. Sensory stimuli convey critical information about various types of opportunities and threats, including access to nourishment, the presence of predators, or the needs of infants. The same stimulus, however, can have different meanings based on previous associations and behavioral context (1-3). For example, in language processing, the same words often have multiple meanings. Humans determine the meaning of these words by integrating prior knowledge with context to converge on a relevant interpretation. How does the brain enable such flexible interpretation of sensory cues based on behavioral context? There are at least three requirements for flexible processing of external stimuli by neural circuits: a stable and high fidelity representation of the sensory stimulus independent of context, a dynamic adjustment when context changes, and an instructive signal to convey the global context. To determine how these requirements translate into synaptic and circuit mechanisms, we trained mice to perform an auditory task using the same sounds in two different contexts, one that produced a learned sensorimotor response (active context) and one that did not (passive context), while we performed calcium imaging or whole-cell voltage-clamp recordings. Neural activity in auditory cortex (AC) was rapidly modified when switching from passive hearing to active engagement. Synaptic inhibition gated these contextual changes, with parvalbumin- and VIP-positive interneurons showing distinct activity profiles. Axonal calcium imaging showed that cholinergic modulation communicated contextual information to AC; correspondingly, contextual changes in inhibition were blocked by atropine and mimicked by optogenetic activation of cholinergic axons. Therefore, excitatory feed-forward drive provides the high-fidelity stimulus representation, synaptic inhibition provides a dynamic switch to adjust network output, and cholinergic tone conveys the relevant global contextual signal. This synaptic and circuit switch provides a rapid mechanism for control of flexible behavior.

T-42. Ring attractor dynamics in the Drosophila central brain

Sung Soo Kim[1,2] SUNGSOO@JANELIA.HHMI.ORG
Herve Rouault[1,2] ROUAULTH@JANELIA.HHMI.ORG
Johannes Seelig[1,2] SEELIGJ@JANELIA.HHMI.ORG
Shaul Druckmann[1,2] DRUCKMANNS@JANELIA.HHMI.ORG
Vivek Jayaraman[1,2] VIVEK@JANELIA.HHMI.ORG

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

Recent results from two-photon calcium imaging in the central brain of head-fixed behaving flies, Drosophila melanogaster, identified a set of neurons that share similarities with mammalian head direction (HD) cells. These wedge neurons, so named because their dendrites each innervate a single wedge of the ellipsoid body (EB), a torus-shaped fly brain structure, together display a bump-like localized activity whose position on the torus follows the fly's body orientation in the presence of landmarks or in darkness. The anatomical arrangement of

this neural population is suggestive of a ring attractor, a network structure hypothesized to generate compass-like activity in mammalian HD cells. In this study, we first used optogenetics to show that simple feedforward connectivity cannot account for the observed EB activity patterns. Then, we used two-photon calcium imaging to test two ring attractor models in wedge neurons of tethered flies. When a visual stimulus abruptly changed its position, bump activity in the EB quickly jumped to the matching position on the torus. The classical ring attractor model, which features global recurrent connectivity with sinusoidal weights between neurons, cannot explain such dynamics, which instead favor a model with local recurrent excitation and global uniform inhibition. We observed that the bump activity often smoothly flowed in response to the same visual stimulus when the tethered fly was in motion. Simulation and physiological results suggest that internally generated velocity cues contribute to bump flow, consistent with observations from mammalian HD cells, whose activity depends on vestibular signals generated by head movements. Ring attractor dynamics are ubiquitously described in the HD cell and orientation tuning literature, yet have only been inferred from observing small subsets of the relevant neurons. Here, we imaged from complete neural populations and performed selective optogenetic perturbation to characterize ring-attractor-like dynamics in the fly brain.

T-43. Neural ensemble dynamics underlying a long-term associative fear memory

Benjamin Grewe[1] GREWE@STANFORD.EDU
Jan Gruendemann[2] JAN.GRUENDEMANN@FMI.CH
Jesse Marshall[1] JESSE.D.MARSHALL@GMAIL.COM
Jones Parker[1] JONESGRIFFITH.PARKER@PFIZER.COM
Jin Zhong Li[1] JINZHONG@STANFORD.EDU
Andreas Luethi[2] ANDREAS.LUTHI@FMI.CH
Mark Schnitzer[1] MSCHNITZER@GMAIL.COM

[1]Stanford University
[2]Friedrich Miescher Institute

The brain's ability to associate different events and external stimuli is vital to the formation of associative memories. Numerous prior studies have examined the molecular, synaptic and cellular level substrates of associative memory, but a systems-level description of how neural ensemble dynamics encode a long-term associative memory has remained elusive. Here, we studied classical fear conditioning as a model system of associative learning and identified the coding scheme by which neural populations in the basal and lateral amygdala (BLA) represent conditioned and unconditioned stimuli (CS and US), as well as the learned association between the two. To do this, we used a miniature fluorescence microscope and time-lapse imaging of BLA ensemble neural Ca2+ dynamics in freely behaving mice across multiple sessions of habituation, fear learning and extinction. By using population vector and linear discriminant analyses we tracked and decoded the changes in neural ensemble information processing that allow the BLA to form and reliably store a long-term associative memory. In distinction to prior work on cellular correlates of fear memory, which proposed that a small set of neurons support memory by potentiating their CS-evoked responses after training, studies of the ensemble code revealed that up- and down-regulation of individual cells' CS-evoked responses were equally important for reliably storing the learned association. This bi-directional plasticity reshaped the ensemble encoding of the CS to increase its similarity to that of the US. In mice that underwent behavioral extinction following training, the neural ensemble representations of CS and US grew more dissimilar, with similar kinetics as those governing the coding changes during conditioning. Throughout learning and extinction the strength of the ensemble-encoded CS-US association was predictive of the mouse's behavioral performance. Together, these results reveal the fundamental information possessing events that allow neural ensembles in BLA to reliably encode an associative memory over days. Our findings may also generalize to multiple other brain areas and forms of associative learning.

I-1. Using the past to estimate sensory uncertainty

Ulrik Beierholm[1] BEIERH@GMAIL.COM
Tim Rohe[2] TIM.ROHE@TUEBINGEN.MPG.DE
Oliver Stegle[3] STEGLE@EBI.AC.UK
Uta Noppeney[1] U.NOPPENEY@BHAM.AC.UK

[1]University of Birmingham
[2]University of Tuebingen
[3]Max Planck Institute for Biological Cybernetics

Combining multiple sources of information requires an estimate of the reliability of each source in order to perform optimal information integration. The human brain is faced with this challenge whenever processing multisensory stimuli, however how the brain estimates the reliability of each source is unclear with most studies assuming that the reliability is directly available. In practice however reliability of an information source requires inference too, and may depend on both current and previous information, a problem that can neatly be placed in a Bayesian framework. We performed three audio-visual spatial localization experiments where we manipulated the uncertainty of the visual stimulus over time. Subjects were presented with simultaneous auditory and visual cues in the horizontal plane and were tasked with locating the auditory cue. Due to the well-known ventriloquist illusion responses were biased towards the visual cue, depending on its reliability. We found that subjects changed their estimate of the visual reliability not only based on the presented visual stimulus, but were also influenced by the history of visual stimuli. The finding implies that the estimated reliability is governed by a learning process, here operating on a timescale on the order of 10 seconds. Using model comparison we found for all three experiments that a hierarchical Bayesian model that assumes a slowly varying reliability is best able to explain the data. Together these results indicate that the subjects' estimated reliability of stimuli changes dynamically and thus that the brain utilizes the temporal dynamics of the environment by combining current and past estimates of reliability.

I-2. Optimal decision making in social networks

Simon Stolarczyk[1] SPSTOLAR@MATH.UH.EDU
Kevin Bassler[1] KBASSLER@CENTRAL.UH.EDU
Manisha Bhardwaj[1] SWEETMANI23@GMAIL.COM
Wei Ji Ma[2] WEIJIMA@NYU.EDU
Kresimir Josic[1] JOSIC@MATH.UH.EDU

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

Humans and other animals integrate information across modalities and across time to perform simple tasks nearly optimally. However, it is unclear whether humans can optimally integrate information in the presence of redundancies. For instance, different modalities, or different agents in a social network can transmit information received from the same or related sources. What computations need to be performed to combine all incoming information while taking into account such redundancies? Moreover, if information propagates through a larger network, does locally optimal inference at each node permit optimal inference of all available information downstream? To address these questions we study a simple Bayesian network model for optimal inference. We first investigate feedforward networks where nodes (agents) in the first layer estimate a single parameter drawn from a Gaussian distribution. The agents pass their beliefs about these estimates on to nodes in the next layer where they are optimally integrated, accounting for redundancies. The information is then propagated analogously across other layers until it reaches a final observer. We give a simple criterion for when the final estimate is nonoptimal, showing that redundancies can significantly impact performance even when information is integrated locally optimally by every agent. This gives us a benchmark to compare to the case when observers do not account for such correlations. We also show that when connections between layers are random, the probability that the final observer can perform optimal inference approaches 1 if intervening layers contain more nodes than the first. We

also examine other factors in the network structure that lead to globally suboptimal inference, and show how the process compares to the case of parameters that follow non-Gaussian distributions, and how information propagates through recurrent networks. This work has the potential to account for how individual performance can be detrimental for group intelligence.

I-3. Approximately Bayesian inference can be implemented by pop. vectors based on point process inputs

Josue Orellana[1] JOSUE@CMU.EDU
Jordan Rodu[1] JORDAN.RODU@GMAIL.COM
Steven Suway[2] SBS45@PITT.EDU
Andrew Schwartz[2] ABS21@PITT.EDU
Robert Kass[1] KASS@STAT.CMU.EDU

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

There is considerable interest in understanding how the brain might use populations of spiking neurons to encode, communicate, and combine sources of information optimally, as specified by Bayesian inference. For instance, Kording and Wolpert, in a 2004 paper, showed that performance on a sensorimotor task was consistent with optimal combination of sensory input and the statistics of the task that were learned during training, while Ma and colleagues in a 2006 paper proposed a neural modeling framework according to which Bayesian inferences could be computed. These works focused on the form in which inputs were combined to produce the posterior mean and variance. We show that population vectors (PV) based on point process inputs combine evidence in a form that closely resembles Bayesian inference, with each input spike carrying information about the tuning of the input neuron. We are investigating performance of population vector-based inference with various tuning functions. While it is exactly Bayesian for von Mises tuning functions, it remains approximately Bayesian for many other cases. We also suggest that encoding stability within short epochs of time could lead to nearly optimal sensorimotor integration.

I-4. Low dimensional representations enhance associative memory flexibility and new learning

Anthony DeCostanzo ANTHONY.DECOSTANZO@GMAIL.COM
Tomoki Fukai TFUKAI@RIKEN.JP

RIKEN Brain Science Institute

The computational power of the brain is thought to reside in its immense scale—the large pool of neurons at its disposal can generate high-dimensional representations that allow complex, nonlinear processing of sensory stimuli and subsequent behavior. Despite the potentially high dimensionality available, low-dimensional representations are commonly found. This raises two general questions. How are the low-dimensional representations achieved? And, why should the brain reduce dimensionality if the chief benefit of having so many neurons is the ability to do computation in high dimensional space? We present a biologically plausible learning rule based on a simple principle of competition among neurons. This rule derives the underlying dimensionality of a pattern classification task, thereby enhancing performance. It emerges from this model that sparse coding among the neurons of the intermediate representation is necessary for optimal performance. Extracting the low, inherent dimensionality of the task benefits the network by freeing dimensions of neuronal population activity for future use. We show, firstly, how this allows the network to rapidly change the class of an existing associative memory, secondly, how the learning of novel, unrelated stimuli is accelerated. Meanwhile, the enhanced dimensionality reduction offered by sparse activity improves the tradeoff between the learning/forgetting of novel/remote stimuli. Therefore, it seems that neural representations that match the dimensionality of behavior offer more than an enhanced input-output mapping. Rather, the freeing of population-activity dimensions enhances the flexibility of existing memories and the acquisition of new ones. Such a learning rule might be expected in higher areas of the brain in which dimensionality would match that of planned or executed behavior as opposed to lower areas in which features are encoded.

I-5. Constraints on the accuracy of auditory discriminations in rats

Alfonso Renart ALFONSO.RENART@NEURO.FCHAMPALIMAUD.ORG
Mafalda Valente MAFALDA.VALENTE@NEURO.FCHAMPALIMAUD.ORG
Dmitry Kobak DMITRY.KOBAK@NEURO.FCHAMPALIMAUD.ORG
Christian Machens CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Jose Pardo-Vazquez JOSE.PARDOVAZQUEZ@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Centre for the Unknown

Does variability of sensory neurons constrain the accuracy of perception? This question has vexed researchers for decades. Answering it requires making sure that an animal's behavioral accuracy in a sensory discrimination task is exclusively limited by sensory factors, something which to the best of our knowledge has not yet been shown in rodents. Here we first provide evidence that this is the case in a 2AFC inter-aural level difference (ILD) discrimination task in rats. We show that, in this task, behavioral accuracy cannot be improved by manipulations of motivation, and that accuracy is unaffected by past trials in most subjects. Given that performance is apparently limited by 'sensory noise', we then studied the structure of neural variability in auditory cortex. As a first step in this direction, we have quantified neural variability in the absence of behavior. Specifically, we simultaneously recorded the evoked responses of large neural populations ($N \sim 100$) from the auditory cortex of rats under urethane anesthesia, using the same set of stimuli as in the ILD discrimination task. We made three key observations: (1) although approximately half of the neurons are ILD-selective, single-cell neurometric functions systematically underperform compared to behavior, except for $<$ 1% of the neurons. (2) Population-level neurometric functions can match behavioral performance, but show a strong brain-state dependence, with performance increasing with the level of desynchronization. Finally, we show that this brain-state dependence arises from a reorganization of where signal and noise reside in the high-dimensional space of population firing rates. Our results establish a behavioral paradigm in rats for the investigation of a long-standing question in neuroscience, and constitute one of the first empirical investigations of the relationship between neural variability and behavioral accuracy at the population level.

I-6. Risk aware control

Terence Sanger TERRY@SANGERLAB.NET

University of Southern California

Risk is a ubiquitous element of human existence, and the species has survived only because our brains guard our relatively unprotected bodies from harm. So it should be no surprise that risk controls the nature of movement and the reflex response to perturbations. The theory of risk-aware control provides a mathematical framework that models human behavior under risk and that also provides simple algorithms that explain this behavior. I derive the mathematical theory and show that a simple implementation can be created using populations of integrate-and-fire spiking neurons.

I-7. Weber's law in disparity estimation is predicted by the statistics of natural stereo-images

Arvind Iyer ARVINDIY@SAS.UPENN.EDU
Johannes Burge JBURGE@SAS.UPENN.EDU
University of Pennsylvania

Human precision in a broad class of sensory-perceptual tasks is well-described by Weber's Law: discrimination thresholds increase systematically with stimulus magnitude. Does this law arise from optimal neural processing in the face of natural image variability? To address this question we examine neural encoding and decoding for the estimation of binocular disparity, a stimulus property for which human performance is known to obey Weber's Law. First, we collected a large database of calibrated stereo images with precisely co-registered laser distance data, using a Nikon d700 DSLR camera and a Riegl VZ-400 range scanner mounted to a robotic gantry. Next, we developed a procedure for sampling binocular stereo-pairs from the dataset with arcsec precision and created a ground-truth labeled training set for a range of disparities (1deg/patch; 1000patch/disparity). Then, using Accuracy Maximization Analysis, we learned a small population model neurons having linear receptive-fields (RFs) optimized for disparity estimation. The population responses were obtained by projecting the contrast-normalized stereo-pairs onto the RFs and adding multiplicative neural noise. These population responses optimally encode the disparity information in the natural stimuli. These responses to natural stimuli also specify the optimal non-linear (quadratic) pooling rules for decoding disparity. We quantified the Fisher Information in the population response, and measured disparity estimation performance with an optimal Bayesian decoder. Both measures predict, consistent with psychophysical findings, that the precision of disparity estimation decreases as predicted by Weber's Law. Similar results hold for other specific tasks (e.g. speed estimation, motion in depth estimation) suggesting these findings are general to an important class of tasks in early- and mid-level vision. Thus, whereas Weber's Law was originally formulated to summarize empirical psychophysical observations, our study offers a novel computational explanation relating task-specific neural processing with the variability of natural images.

I-8. Novelty and uncertainty as separable exploratory drives

Jeffrey Cockburn JEFF.COCKBURN@GMAIL.COM
John O'Doherty JDOHERTY@CALTECH.EDU
California Institute of Technology

Despite the real-world importance of balancing exploration and exploitation, the computational mechanisms brought to bear on the problem by humans and other animals are poorly understood. Strategies for motivating exploration in computational reinforcement-learning include boosting the value of novel stimuli to encourage sampling in new regions of the environment, or augmenting the value of a given option according to the degree of uncertainty in the estimate of predicted reward. While there is preliminary evidence of both uncertainty and novelty directed exploration in humans, the nature of the relationship between them is unknown. In the present study we sought to address how these variables relate to each other. We also sought to address the paradox of why uncertainty driven exploration can co-exist alongside the frequently observed contrarian behavioral imperative of uncertainty avoidance. To this end we tested human participants on a bandit task where these variables were systematically manipulated. We found clear evidence of both novelty and uncertainty driven behavior. Uncertainty and novelty driven exploration were found to evolve differently over time. Approximately half of our sample exhibited uncertainty-seeking early on when there was ample opportunity to exploit what was learned. As participants approached the end of the sampling period all participants became increasingly uncertainty-averse. Conversely, the majority of our participants exhibited novelty-seeking response patterns throughout the session. Moreover, we found a negative correlation between the degree to which participants were influenced by novelty and uncertainty, suggesting an antagonistic relationship between the two exploratory drives. These results support the existence of separable valuation processes associated with novelty and uncertainty as motivations to explore, and provide one possible account for why two competing attitudes toward uncertainty can co-exist in the same individual.

I-9. Pre-perceptual grouping accounts for contextual dependence in the perception of frequency shift

Vincent Adam VINCENTA@GATSBY.UCL.AC.UK
Claire Chambers CHAMBERS.CLAIRE@GMAIL.COM
Maneesh Sahani MANEESH@GATSBY.UCL.AC.UK
Daniel Pressnitzer DANIEL.PRESSNITZER@GMAIL.COM

Gatsby Computational Neuroscience Unit, UCL

Perception is the product of a complex and multi-scale inferential process. For instance, the perceived attributes of an object may depend on how its basic components group together. Here we show that such pre-perceptual grouping underlies a recently reported effect of acoustic context on the perceived direction of pitch shift[1]. Observers judged whether the second of a pair ($T1$, $T2$) of Shepard tones (complex tones with components spaced whole octaves apart around a base frequency, and with amplitudes following a Gaussian spectral envelope [2]) was higher in pitch than the first. The octave-spacing of components makes this judgement ambiguous; the components of $T1$ may be mapped to components of $T2$ that either rising or fall in frequency. When the base frequencies of $T1$ and $T2$ differ little, listeners report the percept associated with the smaller frequency shift between components. But when $T2$ components are equidistant from adjacent $T1$ components (half-octave interval), this cue is removed and listeners report either an upward or downward shift with equal probability. In this case, preceding contextual complex tones (C) bias the response towards the transition that encompasses the frequency region of C. We model these results as arising from pre-perceptual grouping influenced by proximity and temporal expectation. Components are ordinarily grouped by (log-)frequency proximity giving rise to unambiguous shifts between components when this grouping is unique. However, the component groups are not heard individually; rather, the within-group shifts are heard as a shift in the overall pitch of the complex tone. Contextual tones C establish probabilistic expectations about the frequency-locations of groups. A quantitative Bayesian model, based on a factorial Markov process, accounts for characteristics of context dependence: how the bias varies as a function of the frequency of a context stimulus and its relative invariance to spectral and time scale.

I-10. Ecological, rather than physical, features determine object salience

Ali Ghazizadeh IRONICAN@GMAIL.COM
Okihide Hikosaka OH@LSR.NEI.NIH.GOV

National Institutes of Health

We are surrounded by many objects that compete for our attention and influence our decision making. The outcome of this competition can be determined by unique perceptual features of an object (physical salience) or by its enhanced ecological relevance such as its past reward association (ecological salience). The relative contribution of ecological vs physical salience is critical for studying the mechanisms of decision making but has remained controversial. Previous studies with naturalistic stimuli could not independently manipulate perceptual vs ecological features of stimuli. Furthermore the ecological relevance of stimuli for individual subjects was not fully controlled. These factors have impeded the dissociation of the unique contribution of each domain to attention. We have addressed these issues by using macaque monkeys in a design that independently manipulated perceptual and reward history of objects. Salience of each object was tested during a free viewing task. Free viewing before value association showed a preference for physically salient objects. However following value training physical salience ceased to determine attentional bias between objects. Instead we found a significant role of past reward that was present from the first saccade ($p<0.001$). Notably, any putative assignment of salience to objects based on their perceptual features alone is dwarfed by value when determining gaze between objects (Monte-Carlo simulation). We show that physical salience enables rapid object vs background detection. However upon detection of potential objects, ecological salience will be the main determinant for initial attention orienting as well as subsequent exploration of objects. Such a mechanism can provide a certain evolutionary advantage in primates by suppressing attention to ecologically unimportant but physically salient distractors.

I-11. Network-level model of feed-forward inhibition simulates attentional symptoms of schizophrenia

Nathan Insel NATHAN.INSEL@UTORONTO.CA
Blake Richards BLAKE.RICHARDS@UTORONTO.CA

University of Toronto

Among the most prominent symptoms in schizophrenia is patients' inability to filter-out irrelevant stimuli. Decades ago it was hypothesized that this could arise from a disruption of feed-forward inhibition, a process the nervous system may use to adaptively and selectively regulate responses according to an input's magnitude. We developed a network-level computational model to demonstrate how attention deficits can emerge from disrupted feed-forward inhibition. Specifically, we show how the basic principles of two behavioral phenomena affected in schizophrenia, latent inhibition and blocking, can be accomplished in a simple network with Poisson excitatory units and linear, divisive inhibition. We further describe mathematically how combining such a network with a competitive learning output layer can replicate many of the behavioral effects of GABAergic manipulations to the rodent frontal cortex. The network is designed using several principles that may be important for understanding brain circuit function and its disruption in schizophrenia. First, the number of active excitatory units in the middle layer, which we equate with neurons in the medial frontal (rostral cingulate) cortex, can provide a scalar code for attentional relevance, read-out by an efferent unit that we equate with a catecholamine salience signal. Second, the relevance signal can be compared against the actual value of the situation (whether or not an unconditioned stimulus is present) to create a delta signal that feeds-back onto the middle (medial frontal cortex) layer. Third, the delta signal specifically trains synaptic weights between the input layer and the inhibitory unit of the middle layer, thus maintaining homeostatic excitation except when inputs are novel or signal the presence of a learned, relevant stimulus. These theoretical results confirm that attentional symptoms found in schizophrenia could arise from dysfunction of feed-forward inhibitory systems, and further point to a potentially critical role of inhibitory neuron plasticity.

I-12. A unified dynamic model for learning, replay and sharp-wave/ripples

Raoul-Martin Memmesheimer[1,2] RM3354@CUMC.COLUMBIA.EDU
Sven Jahnke[3] SJAHNKE@NLD.DS.MPG.DE
Marc Timme[3] TIMME@NLD.DS.MPG.DE

[1]Columbia University
[2]Radboud University Nijmegen
[3]Max Planck Institute for Dynamics and Self-Organization

Hippocampal activity is fundamental for episodic memory formation and consolidation. During phases of rest and sleep, it exhibits sharp-wave/ripple (SPW/R) complexes – short episodes of increased activity with superimposed high-frequency oscillations. Simultaneously, spike sequences reflecting previous behavior, such as traversed trajectories in space, are replayed. Whereas these phenomena are thought to be crucial for the formation and consolidation of episodic memory, their neurophysiological mechanisms are not well understood. We present a unified model showing how experience may be stored and thereafter replayed in association with SPW/Rs. We propose that replay and SPW/Rs are tightly interconnected as they mutually generate and support each other. The underlying mechanism is based on nonlinear dendritic computation due to dendritic sodium spikes that have been prominently found in the hippocampal regions CA1 and CA3, where also SPW/Rs and replay are generated. Besides assigning SPW/Rs a crucial role for replay and thus memory processing, the proposed mechanisms explain their characteristic features such as the oscillation frequency and the overall wave form. Further, our simulations indicate that for often assumed, standard spike-timing dependent plasticity the SPW/R and replay events may rather erase than enhance learned hippocampal network structures. The results shed a new light on the dynamical aspects of the mechanisms of hippocampal circuit learning.

I-13. The role of orbitofrontal cortex in model-based planning in the rat

Kevin Miller[1] KJMILLER@PRINCETON.EDU
Matthew Botvinick[1,2] MATTHEWB@PRINCETON.EDU
Carlos Brody[1,3] BRODY@PRINCETON.EDU

[1]Princeton University
[2]Google Deep Mind
[3]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 many trials of planned behavior within a single experimental session, opening the door to experiments which characterize the neural correlates of planning. We have adapted the two-step task for use with rats, and shown that inactivations of the dorsal hippocampus, prelimbic prefrontal cortex, and orbitofrontal cortex impair planning behavior on the task. Here, we use extracellular recordings to characterize the neural correlates of planning in the orbitofrontal cortex (OFC). We find that single units in the OFC encode planning-related variables, including model-based value signals. In particular, we find model-based value signals associated with outcomes, but not with choices, arguing in favor of the idea that it acts as part of a model-based learning system (Schoenbaum, et al., 2009), and against the view that OFC plays a key role in model-based decision-making (Padoa-Schioppa, 2011). Ongoing work will further test this idea using temporally precise optogenetic inactivations.

I-14. Hippocampal sharp-wave ripples influence selective activation of the default mode network.

Raphael Kaplan[1,2] RAPHAEL.KAPLAN@UPF.EDU
Mohit Adhikari[1] MOHIT.ADHIKARI@UPF.EDU
Rikkert Hindriks[1] RIKKERT.HINDRIKS@UPF.EDU
Dante Mantini[3] DANTE.MANTINI@HEST.ETHZ.CH
Yusuke Murayama[4] YUSUKE.MURAYAMA@TUEBINGEN.MPG.DE
Nikos Logothetis[4] NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE
Gustavo Deco[1] GUSTAVO.DECO@UPF.EDU

[1]Universitat Pompeu Fabra
[2]University College London
[3]Katholieke Universiteit Leuven
[4]Max Planck Institute for Biological Cybernetics

Hippocampal ripples occur during restful periods and are associated with circuit-level memory consolidation. In parallel, the default mode network (DMN), a prominent network observed during the resting-state, includes brain regions involved in memory consolidation, but has unclear behavioral correlates. Large-scale neocortical fMRI activations and subcortical deactivations have been observed in monkeys specifically after hippocampal sharp-wave ripple ($\sim 80 - -180Hz$ in monkeys) events, yet it is still unclear whether they and other hippocampal neural events influence cortical resting-state networks (RSNs) differently. Investigating fMRI datasets from two anesthetized monkeys with simultaneous hippocampal electrophysiology recordings, we implemented a recently developed technique that uses spatial independent component analysis (ICA) to define correlated fMRI signal fluctuations measured across multiple scan experiments/sessions and subjects into component brain networks. We then isolated the non-human primate equivalents of the DMN and another prominent RSN, the ventral somatomotor network. We first investigated whether there were positive DMN blood-oxygen-level-dependent (BOLD) responses after 2,831 hippocampal ripple events and whether these positive responses also occurred after the onset of 2,004 hpsigma ($\sim 8 - -22Hz$) and 1,740 gamma ($\sim 25 - -75Hz$) hippocampal events. Second, we investigated whether these three different types of hippocampal events, also co-occurred with BOLD signal fluctuations in the ventral somatomotor network, a RSN not implicated in memory consolidation. Consequently, we could determine whether RSN BOLD responses were network and neural-event specific. We observed a dramatic increase in the DMN BOLD signal following ripples, but not other electrophysiological events in the hippocampus. Notably, we found BOLD increases in the DMN after hippocampal ripples, but not in a prominent ventral somatomotor RSN. Our results relate endogenous fluctuations in the DMN BOLD signal to the onset of hippocampal ripple events' linking resting-state fMRI network fluctuations with behaviorally relevant circuit-level neural dynamics.

I-15. Stimulus detectability in confidence judgments: a normative account and recurrent neural network

Megan Peters MEGANAKPETERS@UCLA.EDU
Hakwan Lau HAKWAN@GMAIL.COM

University of California, Los Angeles

Perceptual decision-making is well described by models that calculate the 'balance of evidence' favoring each stimulus alternative relative to the other(s), including drift diffusion models and Bayesian ideal observers. However, a 'balance of evidence' approach cannot account for multiple recent findings, including (a) higher proportion of 'seen' (high-confidence) trials for unattended over attended stimuli, despite similar discrimination performance; (b) differential metacognitive sensitivity for response alternatives due to unbalanced evidence; and (c) decreased metacognitive sensitivity when decoded neurofeedback (DecNef) is used to change confidence ratings. It has proven difficult to explain these data without relying on a heterogeneous set of heuristics that differ by situation. Here we propose a single, normative framework to account for these findings, and additionally provide a recurrent neural network implementation. It was recently proposed that the absolute magnitude of response-congruent evidence (i.e., stimulus energy/detectability favoring the chosen stimulus alternative) is disproportionately weighted in confidence judgments. We also recently showed that confidence judgments can be well described both by an ideal observer and a Bayesian observer that takes into account stimulus detectability. Following this work, here we present the Confidence as Detectability (CaD) model: a normative Bayesian framework accompanied by a recurrent neural network implementation that utilizes tuned normalization to represent the degree to which a neuron codes for balance of evidence versus stimulus energy magnitude, which contribute differentially to perceptual decisions and confidence judgments. In a series of simulations, we show that this model can explain all of the above-mentioned findings with a single set of parameters. Importantly, using the 'detectability' of a stimulus in judging confidence may not be suboptimal in the real environment, when the task is to judge not only which of innumerable stimulus alternatives is most likely to be present, but also whether a stimulus is present at all.

I-16. A collicular mechanism for flexible sensorimotor gating during task switching

Chunyu Duan[1] ANNDUAN2@GMAIL.COM
Marino Pagan[1] MPAGAN@PRINCETON.EDU
Charles Kopec[1] CKOPEC@PRINCETON.EDU
Jeffrey Erlich[2,1] JERLICH@NYU.EDU
Alexander J Riordan[1] ARIORDAN@PRINCETON.EDU
Athena Akrami[1] AAKRAMI@PRINCETON.EDU
Carlos Brody[1,3] BRODY@PRINCETON.EDU

[1]Princeton University
[2]New York University, Shanghai
[3]Howard Hughes Medical Institute

Flexible sensorimotor gating based on current environmental context is a fundamental component of adaptive behavior, but its underlying neural mechanisms are still largely unknown. While previous studies have mostly focused on the role of frontal regions such as prefrontal cortex (PFC), we recently demonstrated that in the rat the midbrain superior colliculus (SC) also plays an important role. Here, we use electrophysiological recordings, optogenetics, and computational modeling to further examine the role of SC in this executive function. On each trial, rats were first presented with an auditory cue indicating the current task rule ('Pro' or 'Anti'), followed by a short memory delay period, and finally a choice period during which rats were required to either orient toward ('Pro') or away ('Anti') from a visual stimulus. Using a neural population decoding approach, we were surprised to find that SC neurons contained significantly more information about task identity and earlier information about upcoming choice than PFC, thus arguing against a model in which the decision is first computed in PFC and then relayed to SC. During the memory delay, a subpopulation of SC neurons displayed strong task information that was disrupted on error trials. Consistent with this observation, bilateral optogenetic silencing of SC neurons during the delay period resulted in a behavioral deficit. The same subpopulation of neurons had a short-latency choice signal, and was thus compatible with a role in both maintaining and implementing the current task rule. However, no decrease in accuracy was observed upon bilateral SC inactivation during the choice period. The electrophysiology and inactivation data were reconciled in a model of SC as a 4-way mutually-inhibitory circuit that is sufficient to perform this flexible task-switching behavior. Overall, our findings suggest that SC might play a central role in the execution of flexible behavior.

I-17. Humans exhibit discrete confidence levels in perceptual decision-making

Matteo Lisi MATTEO.LISI@PARISDESCARTES.FR
Gianluigi Mongillo GIANLUIGI.MONGILLO@PARISDESCARTES.FR
Andrei Gorea ANDREI.GOREA@PARISDESCARTES.FR

Universite Paris Descartes

Animals (including humans) are able to assess the quality of incoming sensory information and act accordingly while taking decisions. The computations underlying such ability are unclear. If neuronal activity encodes probability distributions over sensory variables, then uncertainty—hence confidence—about their value is explicitly represented and, at least in principle, readily accessible. On the other hand, if neuronal activity encodes point-estimates, then confidence must be obtained by comparing the level of the evoked response to fixed (possibly learned) criteria. To address this issue we developed a novel task allowing the behavioral read-out of confidence on a trial-by-trial basis. Each trial consisted of two consecutive decisions on whether a given signal was above or below some reference value, call it zero. The first decision was to be made on a signal uniformly drawn from an interval centered at zero. Correct/incorrect responses resulted into signals uniformly drawn from the positive/negative sub-intervals to be judged when making the second decision and subjects were told so. The task reliably elicited confidence assessments as demonstrated by the finding that second decisions were more fre-

quently correct than first decisions. We compared the ability of Bayesian and non-Bayesian observers to predict the empirically observed pattern of both first and second decisions. The non-Bayesian observer was designed to have discrete confidence levels instantiated by one, two or three second-decision criteria representing different levels of the evoked response. Different confidence levels resulted into different second-decision criteria. Synthetic data-sets reliably discriminated Bayesian from non-Bayesian observers. The non-Bayesian observer with two-three confidence levels systematically (over 9 subjects) outperformed the Bayesian observer in predicting the actual behavior. Hence, contrary to previous claims, confidence appears to be a discrete rather than continuous quantity. Simple heuristics are sufficient to account for confidence assessment by humans making perceptual decisions.

I-18. Encoding of value and choice as separable, dynamic neural dimensions in orbitofrontal cortex

Daniel Kimmel[1] DANIELKIMMEL@GMAIL.COM
Gamaleldin F Elsayed[1] GAMALELDIN.ELSAYED@GMAIL.COM
John Cunningham[1] JPC2181@COLUMBIA.EDU
Antonio Rangel[2] RANGEL@HSS.CALTECH.EDU
William Newsome[3] BNEWSOME@STANFORD.EDU

[1]Columbia University
[2]California Institute of Technology
[3]Stanford University

Orbitofrontal cortex (OFC) has long been implicated in value-based decision-making. However, OFC responses are complex, with individual neurons representing multiple task-relevant signals, such as stimulus value, behavioral choice, and expected reward. It was therefore unclear how mixed responses of a given neuron might contribute to distinct cognitive-behavioral functions theoretically subserved by these various signals. Separately, the tools of dynamical systems theory have provided low-dimensional descriptions of the complex responses of premotor cortices in the context of motor behavior and perceptual decision-making. However, the diverse responses from more abstract association areas, such as OFC, have eluded a compact, satisfying description. Moreover, the application of dynamical systems has largely been descriptive, without rigorous means to evaluate the significance of a given low-dimensional response. Here, we combined dynamical systems theory and statistical hypothesis testing to understand the complex, heterogeneous signals observed in OFC during value-based decision-making. Specifically, we asked how task-relevant signals were encoded at the population level and assessed the stability of this encoding during the trial. We recorded from macaque OFC while monkeys performed a cost-benefit decision-making task that required animals to evaluate an offer and then maintain an effortful response so as to earn the promised reward. Using a novel dimensionality reduction technique, we identified a low-dimensional subspace in which separable patterns of neural activity (i.e., linear combinations of hundreds of serially recorded neurons) in OFC represented distinct task-relevant variables, e.g., value, choice, and expected reward. We went further and developed novel statistical methods – applicable to many high-dimensional datasets – that determined the representations were not only significant in magnitude, they also were stable over discrete temporal epochs that aligned with behaviorally relevant events. The separability and temporal dynamics of these neural representations suggest they may subserve distinct cognitive-behavioral functions essential for cost-benefit decision-making.

I-19. A unifying theory of explore-exploit decisions

Robert Wilson[1] BOB@EMAIL.ARIZONA.EDU
Jonathan Cohen[2] JDC@PRINCETON.EDU

[1]University of Arizona
[2]Princeton University

Many decisions involve a choice between exploring unknown opportunities and exploiting well-known options. Work across a variety of domains, from animal foraging to human decision making, has suggested that animals solve such 'explore-exploit dilemmas' with a mixture of at least two exploration strategies—one driven by information seeking (directed exploration) and the other by behavioral variability (random exploration). Here we propose a unifying theory in which these two strategies emerge naturally from a kind of stochastic planning, known in the machine learning literature as Posterior Sampling for Reinforcement Learning (PSRL) (Strens, 2000). Briefly, the PSRL model assumes that people make explore-exploit decisions by simulating a small number (as low as 1) of random, but plausible, experiences in order to approximate the expected value of taking different actions. Once these values are computed, the decision is made by picking the action with highest approximate value. Random exploration arises naturally from this model because the simulations are stochastic. More subtly, PSRL also yields directed exploration based on the details of how the simulated choices play out. Crucially, PSRL makes a number of predictions about how directed and random exploration change in different situations. We tested the predictions of the model in two experiments in which we manipulated the uncertainty subjects had about the explore and exploit options (Experiment 1) and the time horizon available for exploration (Experiment 2). In line with the model, we found that directed and random exploration have qualitatively different dependence on both uncertainty and horizon. This is perhaps most striking for the horizon manipulation in which directed exploration asymptotes rapidly with horizon, while random exploration does not. These findings are difficult to explain with other models and provide strong initial support for the theory.

I-20. Impulsive or indecisive: decision-making impairment in a cortical model from disrupted E/I balance

John D Murray[1] JOHN.MURRAY@YALE.EDU
Thiago Borduqui[2] TBORDUQUI@GMAIL.COM
Jaime Hallak[2] JECHALLAK@RNP.FMRP.USP.BR
Antonio Roque[2] ANTONIOR@FFCLRP.USP.BR
Alan Anticevic[1] ALAN.ANTICEVIC@YALE.EDU
Xiao-Jing Wang[3] XJWANG@NYU.EDU

[1]Yale University
[2]University of Sao Paulo
[3]New York University

The balance between excitation and inhibition (E/I balance) is a fundamental property of cortical circuits. Disruption of E/I balance is a leading hypothesis for pathophysiologies of neuropsychiatric disorders, such as schizophrenia, which are also associated with cognitive deficits, including impaired decision making. However, it is poorly understood how E/I disruptions at the synaptic level propagate upward to induce cognitive deficits at the behavioral level. To link these levels of analysis, we investigated how E/I perturbations may lead to impaired temporal integration of evidence during decision making in a biophysically-based model of an association cortical microcircuit (Wang, Neuron 2002). Specifically, we tested the effects of hypofunction of NMDA receptors at two key sites: on inhibitory interneurons (thereby elevating E/I ratio via disinhibition), versus on excitatory pyramidal neurons (thereby reducing E/I ratio). We found disruption of E/I balance in either direction can impair decision making performance as assessed by psychometric functions. Nonetheless, these two regimes make dissociable predictions under fine-grained analyses of the time course of evidence accumulation. In the regime of elevated E/I ratio, behavior can be characterized as impulsive: evidence early in time is weighted much more than evidence late in time, compared to a control circuit. In contrast, in the regime of reduced E/I ratio, behavior can be characterized as indecisive: the circuit exhibits weakened evidence integration and reduced winner-take-all competition between options. These regimes are further distinguished from the scenario of impaired upstream coding of sensory evidence. Our findings characterize a role of E/I balance in cognitive functions supported by cortical circuits. The model makes specific predictions for behavior and neural activity that can be tested in humans or animals under manipulation of E/I balance (e.g. via pharmacology) or in disease states.

I-21. Optimal and suboptimal integration of sensory and value information in perceptual decision-making

Hyang Jung Lee HYANGJUNG.LEE@SNU.AC.KR
Issac Rhim ISSACRHIM@GMAIL.COM
Sang-Hun Lee VISIONSL@SNU.AC.KR
Seoul National University

Many existing models offer successful formalisms for perceptual or value-based decision-making separately. However, optimization of decision-making often requires the effective integration of sensory and value information, particularly when sensory inputs are ambiguous and the criterion for successful decision changes stochastically over time. To understand how human individuals adjust their decision under this situation, we developed a Linear-Nonlinear-Poisson model for integrating sensory and value information, inspired by previous animal studies (Corrado et al., 2005; Lau & Glimcher, 2005; Busse et al.,2011). Then we fit our model to the temporal dynamics of choices made by 30 individuals, who classified ring stimuli of 5 different sizes into 'small' or 'large' classes and received either 'correct' or 'incorrect' feedback on each trial. The key manipulation was, unbeknownst to subjects, to induce a subtle amount of bias in feedback, favoring either 'small' or 'large' choices, or staying 'unbiased'. The model was successful at capturing the trial-courses of individual choice behavior. Having confirmed the competence of our model, we carried out an 'ideal' decision-maker analysis to characterize the individuals in terms of how far they deviated from their own, tailor-made, ideal decision-maker both in performance and in model parameter space. Then, we conducted the canonical correlation analysis (CCA) to know which model parameters are associated with the performance in which phases relative to the onset of biased feedback. The CCA identified one significant mode, which links the three model parameters, two determining the tail length of the reward and choice kernels respectively at the linear stage and one determining the slope of the non-linear softmax function, with the performances 'after' the onset of feedback bias. This suggests that suboptimal perceptual decision-making is majorly due to suboptimal translation of reward/ choice histories into a decision variable and suboptimal arbitration of the decision variable into action.

I-22. On the formation of phoneme categories in computational neural network models

Tasha Nagamine TASHA.NAGAMINE@COLUMBIA.EDU
Michael Seltzer MSELTZER@MICROSOFT.COM
Nima Mesgarani NIMA@EE.COLUMBIA.EDU
Columbia University

Computational neural network models have recently gained widespread popularity in speech recognition due to their markedly improved performance over other models. Loosely based on biological network models, these artificial neural networks must perform the same computation as the human auditory pathway when presented with a speech stimulus; that is, mapping a time-varying pattern of acoustic frequencies to relevant perceptual categories. In this study, we propose a deep neural network (DNN) model trained for phoneme recognition as a potential model of neural speech processing and analyze its properties using standard neuroscience techniques in an attempt to better understand possible feature representations in the brain. First, we characterized the model's representational properties at the single node and population level in each layer and found that selectivity to phonetic features organize the activations in different layers of a DNN, a result that mirrors the recent findings of feature encoding in the human auditory system. Furthermore, we found that the network learns an invariant classification scheme by using a representational basis that explicitly models multiple acoustic variations of each phoneme. We also attempted to study the mechanism by which these feature representations are encoded by characterizing aspects of the nonlinearity used in the DNN, since cortical processing is also highly nonlinear. We show that the hidden layer nonlinearities warp the feature space non-uniformly so as to increase the discriminability of acoustically similar phones, aiding in their classification. These analyses can be used as a comparative

baseline for neural data to determine to what extent neural network models can serve as feasible models for human speech perception. Additionally, this study may also provide intuitions as to how neural networks fail to explain complex cortical processing and suggest the addition of biologically-inspired mechanisms such as feedback that can improve these models.

I-23. Spiking neural networks as superior generative and discriminative models

Luziwei Leng[1] LUZIWEI.LENG@KIP.UNI-HEIDELBERG.DE
Mihai Petrovici[1] MPEDRO@KIP.UNI-HEIDELBERG.DE
Roman Martel[1] ROMAN.MARTEL@KIP.UNI-HEIDELBERG.DE
Ilja Bytschok[2] ILJA.BYTSCHOK@KIP.UNI-HEIDELBERG.DE
Oliver Breitwieser[1] OLIVER.BREITWIESER@KIP.UNI-HEIDELBERG.DE
Johannes Bill BILL.SCIENTIFIC@GMAIL.COM
Johannes Schemmel SCHEMMEL@KIP.UNI-HEIDELBERG.DE
Karlheinz Meier[1] MEIERK@KIP.UNI-HEIDELBERG.DE

[1]University of Heidelberg
[2]Kirchhoff-Institute for Physics

An increasing number of experiments suggest that the brain performs stochastic inference when dealing with incomplete and noisy sensory information. This, in turn, has led to the development of various theoretical models that attempt to explain how this could be achieved with spiking neural networks. One candidate theory interprets spiking activity as sampling from distributions over binary random variables (Buesing et al., 2011) and has been shown to be compatible with the ensemble dynamics of noise-driven LIF neurons in the high-conductance state (Petrovici et al., 2013, 2015; Probst et al., 2015). Based on this theory, we constructed hierarchical LIF networks that sample from restricted Boltzmann distributions and compared their performance with conventional restricted Boltzmann machines (RBMs) on a commonly used dataset (MNIST). An important result is that LIF networks can achieve similar classification rates (95.1 % with 1994 neurons) as their machine-learning counterparts of equal size (95.2 %). In classical RBMs however, statistics are typically gathered by Gibbs sampling. This algorithm has a distinct disadvantage when dealing with high-dimensional multimodal distributions, where it often gets trapped in a local minimum due to deep troughs in the energy landscape that appear during training. It is for this reason that conventional RBMs that may perform very well as discriminative models are, at the same time, rather poor generative models of the learned data. While various methods exist that alleviate this problem (such as AST, see Salakhutdinov, 2010) they usually come at a highly increased computational cost. In the second part of our study, we show how short-term plasticity enables LIF networks to travel efficiently through the energy landscape and thereby attain a generative performance that far surpasses the one achievable by conventional Gibbs sampling. This distinct advantage of biological neural networks allows them to simultaneously become good generative and discriminative models of learned data.

I-24. An attractor neural network architecture with an ultra high information capacity

Alireza Alemi ALIREZA.ALEMI@ENS.FR

Ecole Normale Superieure

Attractor neural network is an important theoretical scenario for modeling memory function in the hippocampus and in the cortex. In these models, memories are stored in the plastic recurrent connections of neural populations in the form of ‘attractor states’. The maximal information capacity for conventional abstract attractor networks with unconstrained connections is 2 bits/synapse. However, an unconstrained synapse has the capacity to store infinite amount of bits in a noiseless theoretical scenario: a capacity that conventional attractor networks cannot achieve. To address this challenge, I propose a hierarchical attractor network that can achieve an ultra high information capacity. The network has two layers: a visible layer with Nv neurons, and a hidden layer with Nh neurons. The visible-to-hidden connections are set at random and kept fixed during the training phase, in which the memory patterns are stored as fixed-points of the network dynamics. The hidden-to-visible connections, initially normally distributed, are learned via a local, online learning rule called the Three-Threshold Learning Rule. Note that there is no within-layer connections. Random, uncorrelated patterns were stored at the dense regime in a network of binary units. The results of simulations suggested that the maximal information capacity grows exponentially with the expansion ratio Nh/Nv. As a first order approximation to understand the mechanism providing the high capacity, I simulated a naive mean-field approximation (nMFA) of the network. The exponential increase was captured by the nMFA, revealing that a key underlying factor is the correlation between the hidden and the visible units. Additionally, it was observed that, at maximal capacity, the degree of symmetry of the connectivity between the hidden and the visible neurons increases with the expansion ratio. These results highlight the role of hierarchical architecture in remarkably increasing the performance of information storage in attractor networks.

I-25. Versatile predictive estimator without weight copying

David Xu D44XU@UWATERLOO.CA
Cameron Seth CJMPSETH@UWATERLOO.CA
Jeff Orchard JORCHARD@UWATERLOO.CA

University of Waterloo

In predictive coding, cortical circuits called predictive estimators (PEs) participate in a hierarchy, passing predictions to lower layers, which send back the prediction error. These PE units learn a set of connection weights that translate between the layers. However, the standard implementation is not biologically plausible, since it uses the same weight matrix for both feed-forward and feed-back projections. We propose a more general PE unit that can learn more complex, non-linear transformations in a biologically plausible manner. Our connection weights are adjusted using an error signal that is local to the connections themselves, and there is no need for connection-weight copying. Our new architecture is demonstrated by learning a Cartesian-to-Polar transformation, and exhibits non-classical receptive fields akin to end-stopping.

I-26. Full-rank regularized learning in recurrently connected firing rate networks

Brian DePasquale[1] BRIAN.DEPASQUALE@ICLOUD.COM
Christopher Cueva[1] CJC2197@CUMC.COLUMBIA.EDU
Raoul-Martin Memmesheimer[1,2] RM3354@CUMC.COLUMBIA.EDU
Larry Abbott[1] LFA2103@CUMC.COLUMBIA.EDU
Sean Escola[1] SEAN.ESCOLA@GMAIL.COM

[1]Columbia University
[2]Radboud University Nijmegen

Trained recurrent neural networks (RNNs) are powerful tools for modeling neural representation and computation in an artificial neural circuit. We present a least-squares based method for training the full connectivity matrix of RNNs that allows small networks of firing-rate units to reliably perform tasks that evolve over behaviorally relevant timescales (on the order of seconds) but that does not require the extensive training time of back-propagation or other gradient based methods. Our method is an alternative learning algorithm in the 'reservoir computing' framework. In this framework, an initial connectivity matrix J is chosen randomly so that the untrained network activity is rich and therefore forms a suitable basis for constructing a desired target function F(t). Typical training

approaches, such as FORCE learning, add a rank-1 matrix to J such that a linear readout of the network activity Z(t) matches the target F(t). Our method learns the full recurrent connectivity matrix so that a linear projection of the network activity Z(t) matches the desired target function F(t) without restricting that the recurrent connectivity matrix be a rank-1 perturbation to the initial connectivity. Instead, we perform learning at the input into each network unit while regularizing the full, learned matrix. Learning the full connectivity matrix removes superfluous and potentially harmful modes of the initial connectivity in addition to incorporating the standard rank-1 connectivity. We found that our method was successful at training small networks of firing-rate units to perform a variety of tasks that rank-1 methods were unable to reliably solve with the same number of units. Since our method does not require the computation of gradients it converges faster than back-propagation or other gradient based methods while performing tasks with similar numbers of units at similar performance levels.

I-27. Correlation-based model predicts efficacy of artificially-induced plasticity in motor cortex by a bidirectional brain-computer interface

Guillaume Lajoie[1] GLAJOIE@UW.EDU
Nedialko Krouchev[2] NEDIALKO.KROUCHEV@MCGILL.CA
Adrienne Fairhall[1] FAIRHALL@U.WASHINGTON.EDU
Eberhard Fetz[1] FETZ@UW.EDU

[1]University of Washington
[2]McGill University

Experiments on macaque monkeys reveal that neurons in Motor Cortex (MC) display a variety of activities correlated to their co-activated muscles and the motor task being performed. Generally, MC neurons with overlapping muscle fields are spatially grouped together and may have enhanced synaptic connections as opposed to more distant neurons. Such connections are believed to be simultaneously a source and a consequence of correlated neural activity among MC neurons, mediated by Spike-Time-Dependent Plasticity (STDP) mechanisms. Consistent with this paradigm, spike-triggered stimulation performed with Bidirectional Brain-Computer-Interfaces (BBCI) can artificially strengthen synaptic connections between distant MC sites and even between MC and spinal cord sites, with changes that last several days. Here, a neural implant is triggered by spikes from an MC site and electrically stimulates a secondary target site after a set delay, the value of which is critical in determining the efficacy of the procedure and consistent with experimentally derived STDP windows. As the development of BBCIs progresses, with applications ranging from a science-oriented tool to clinical treatments, it is crucial to develop a theoretical understanding of the interaction between neural implants, recurrent neural activity from cortical sites, and the plasticity mechanisms that modulate synaptic strengths. In parallel with ongoing experiments, we are developing a recurrent network model with probabilistic spiking mechanisms and plastic synapses (STDP) capable of capturing both neural and synaptic activity statistics relevant to BBCI protocols. This model successfully reproduces key experimental results and we use analytical derivations to predict optimal operational regimes for BBCIs. We make experimental predictions concerning the efficacy of spike-triggered stimulation in different regimes of cortical activity such as awake behaving states or sleep. Importantly, this work provides a theoretical framework which is intended as a design testbed for next-generations applications of BBCI.

I-28. Anterior piriform cortex is essential for olfactory working memory

Chengyu Li[1] TONYLICY@ION.AC.CN
Xiaoxing Zhang[2] ZHANGXIAOXING@ION.AC.CN
Wenjun Yan[2] WJYAN@ION.AC.CN
Yulei Chen[2] YLCHEN@ION.AC.CN
Hongmei Fan[2] HMFAN@ION.AC.CN

[1]Chinese Academy of Sciences
[2]Institute of Neuroscience

Working memory (WM) is a critical cognitive ability of actively maintaining and manipulating information over a delay period of seconds. Modulation in activity of sensory regions during delay period has been observed, but the functional role remains unclear. In the previous study with an olfactory delayed non-match to sample (DNMS) task, we have demonstrated that delay-period activity of medial prefrontal cortex was critical for information maintenance only in learning, but not for well-trained phase of the WM task (Liu, et al., Science, 2014). A natural follow-up question is which brain regions are important in well-trained phase. In the current unpublished study, we addressed these questions using optogenetic and electrophysiological methods and demonstrated the causal role of an olfactory sensory region, anterior piriform cortex (APC), in performing the DNMS task in both learning and well-trained phases. In mice well-trained of the task, suppressing APC activity during memory-delay, sensory-delivery, and response-delay periods impaired performance (Figure A, B). Suppressing delay-period APC activity also impaired WM performance in learning. Interestingly, elevating APC delay-period activity in an odor-specific manner during learning of the task enhanced performance in well-trained phase (Figure C). The later-phase delay-period activity appeared to be more important for the bi-directional manipulation in performance. Furthermore, neural activity of APC in delay period was correlated with memory retention (Figure D). The sustained encoding of olfactory information during delay period seems to be acquired through learning of the WM task, because of the lack of sustained delay decoding power in naive mice. Our results therefore uncovered the essential role of APC in both learning and well-trained phases of the olfactory WM task and further underscored the importance of a sensory cortex in WM.

I-29. Identification of a stable STDP rule from spike timing with generalized multilinear modeling

Brian Robinson BSROBINS@USC.EDU
Theodore Berger BERGER@USC.EDU
Dong Song DSONG@USC.EDU

University of Southern California

Spike-timing-dependent plasticity (STDP) is an attractive spike-based learning rule because of its Hebbian and other computational properties as well as its widespread experimental evidence. Characterization of STDP in the mammalian brain during behavior, however, has mostly been indirect, relying largely on analyses of neural responses evoked by environmental stimuli. We have developed a computational framework for quantifying STDP during behavior by estimating a functional plasticity rule solely from spiking activity. First, we formulate a flexible point process spiking neuron model structure with STDP, which includes functions that characterize the stationary and plastic properties of the neuron. One of the key challenges in estimating these plastic properties is addressing a realistic time course for plasticity induction as well as ensuring the stability of the plasticity rule. The STDP model includes a novel function for prolonged plasticity induction as well as a more typical function for synaptic weight change based on the relative timing of input-output spike pairs. Consideration for system stability is incorporated with weight-dependent synaptic modification. Next, we formalize an estimation technique utilizing a generalized multilinear model (GMLM) structure with basis function expansion. The weight-dependent synaptic modification adds a nonlinearity to the model, which is addressed with an iterative unconstrained optimization approach. We demonstrate successful model estimation on simulated spiking data and show that all model functions can be

estimated accurately with this method across a variety of simulation parameters, such as number of inputs, output firing rate, input firing type, and simulation time. Since this approach only requires naturally generated spikes, it can be readily applied to behaving animals performing cognitive tasks.

I-30. Unsupervised learning of neural sequences

Ulises Pereira ULISES@UCHICAGO.EDU
Nicolas Brunel NBRUNEL@GALTON.UCHICAGO.EDU

The University of Chicago

Neuronal networks in the brain learn stereotyped sequences of activity in tasks ranging from motor to cognitive. Often there is no apparent supervision, and sequential neural activity is learned using repeated presentation of a stimulus without an error signal. The neural sequences learned are usually robust, lasting hundred of millisecond to seconds, with little amplitude attenuation. However, it is unclear what are the necessary conditions for a plasticity rule to learn neural sequences (NS) after repeated sequential stimulation. Here we first show that in a purely excitatory network with both recurrent and feed forward connections, there is a range of connectivity parameters for which NS are generated, and their amplitude do not attenuate across populations within the sequence. Moreover, including inhibition and adaptation produce NS that behave in an all-or-none fashion, and once triggered are difficult to stop. By introducing a general class of separable excitatory to excitatory synaptic plasticity rules, we mapped stimulation parameters to the learning dynamics, identifying the region of stimulus parameters where learning NS take place. We found that the sequence can be reliably retrieved after learning by stimulating just the first population within the sequence. However, we demonstrate that for stable learning during multiple stimulus presentation, an additional variable is necessary for controlling coactivation, preventing NS triggering during stimulation. For retrieving the learned NS, this variable must return to values compatibles with NS generation. The exact mechanism for controlling activity can be implemented in different ways, and distinct commonly used homeostatic plasticity rules produce stable unsupervised NS learning. Our work provide experimentally testable predictions, suggesting that during multiple presentations of the stimulus, population dynamics must present fast time scales, followed after learning, by a slow transition to slower time scales consistent with NS generation.

I-31. Estimating short-term synaptic plasticity from paired spikes in vitro

Abed Ghanbari[1] ABED.GHANBARI@UCONN.EDU
Vladimir Ilin[2] VLI4@PITT.EDU
Maxim Volgushev[1] MAXIM.VOLGUSHEV@UCONN.EDU
Ian Stevenson[1] IAN.STEVENSON@UCONN.EDU

[1]University of Connecticut
[2]University of Pittsburgh

Synaptic connections between neurons evolve over time and these changes affect transmission and processing of information in neuronal circuits. Synaptic plasticity is traditionally studied using intracellular recording techniques where the synaptic ‘weight’ can be directly estimated from postsynaptic potentials or currents. Here we consider the problem of estimating synaptic weights and short-term plasticity from paired spike recordings. We use a hierarchical generalized-linear-model (GLM) to describe the dynamics of short-term synaptic plasticity with time-varying coupling between a presynaptic neuron and a postsynaptic neuron. We constrain the coupling term to vary according to an extended Tsodyks and Markram (eTM) model. Under this model a set of nonlinear differential equations describe depressing and facilitating synaptic dynamics using four parameters: the baseline release probability, the magnitude of facilitation, and time constants for depression and facilitation. We estimate model parameters using maximum likelihood with a coordinate ascent that alternates between optimizing the GLM parameters and the eTM parameters and use bootstrapping to quantify parameter uncertainty. In order to test

the accuracy of plasticity estimation in a realistic setting we use in vitro current injection where 1024 simulated presynaptic inputs, with different synaptic weights and plasticity, drive a recorded neuron. We recorded spike responses to injection of this artificial postsynaptic current in layer 2/3 pyramidal neurons in slices from rat visual cortex in vitro. In both this controlled experimental setting, as well as in simulation, we find that a model-based approach 1) can recover short-term plasticity parameters from pairs of spike trains and 2) makes more accurate spike predictions than a model without plasticity.

I-32. Neural and neuromodulatory substrates underlying exploration initiation

Timothy Muller TIMOTHYMULLER127@GMAIL.COM
Timothy EJ Behrens BEHRENS@FMRIB.OX.AC.UK
Jill O'Reilly JOREILLY@FMRIB.OX.AC.UK

University of Oxford

Uncertain and changing environments render decision-makers in conflict between exploiting and exploring information. The neural substrates of exploitative and exploratory decisions have been identified (Daw et al, 2006). However, the transition between these states is not understood. We designed a probabilistic learning task in which subjects inferred which one of several locations contained a high instead of low probability of reward. This high reward-probability location moved unpredictably, reliably driving subjects between phases of exploitation and exploration. Taking advantage of whole-brain analyses offered by functional MRI (fMRI), we demonstrate the anterior cingulate cortex (ACC) is activated at the initiation of exploration, and that this activation precedes that of the lateral intraparietal area (LIP). The noradrenergic system can be indexed by its positive correlation with pupil diameter (Nassar et al, 2012) and has been implicated in exploratory behaviour (Jepma & Nieuwenhuis, 2011). Here we show an increase in pupil diameter between the last trial of exploitation and the first of exploration. Further, simultaneous fMRI and pupil measurements allowed us to demonstrate that trial-by-trial ACC activity correlates with change in pupil diameter. This correlation remarkably held when only including the last trial of exploitation. These findings reinforce a role of the ACC in switching away from default courses of action (Kolling et al, 2012) and the revision of beliefs (Karlsson et al, 2012), and further demonstrate this may be via altering activity in LIP—an area known to integrate information for action selection (O'Reilly et al, 2013a; Gottleib, 2007; Yang & Shadlen, 2007). These data also demonstrate the suspected (McClure et al, 2006) link between ACC fMRI activation and pupil-linked arousal systems at exploration-initiation in humans.

I-33. Short-term plasticity amplifies exploration in reinforcement learning tasks

Sara Zannone S.ZANNONE14@IMPERIAL.AC.UK
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK

Imperial College London

Long-term synaptic plasticity is widely believed to be the basis of learning and memory. Experimental evidence indicates that short-term plasticity varies in a manner that strongly depends on long-term plasticity. Nevertheless, the precise functional role short-term plasticity in learning is still unknown. We addressed this open question in a reinforcement learning framework. We hypothesised that short-term plasticity could increase exploration by adding extra variability in a structured way. In order to test this theory, we built a model of short- and long-term reward-modulated synaptic plasticity, according to plasticity experiments. We then studied this model for two different settings. We first analysed, both analytically and numerically, the effect of this plasticity rule on a deterministic network model of the multi-armed bandit problem. In reinforcement learning network models, exploration is typically induced through the explicit injection of external noise. Our findings indicate that short-term plasticity alone can, instead, introduce exploration in a completely deterministic model and successfully solve the multi-armed bandit problem. We then investigated the effect of this combined short- and long-term

plasticity rule in a model of navigation, using a continuous state and action space reinforcement learning model. Our results indicate that, even in a non-deterministic framework where exploration is already present, short-term plasticity increases exploration. This leads to an improvement in performance in real world scenarios where the environment changes. Overall, our findings confirm that short-term plasticity can, indeed, code for exploration. This strongly suggests a novel functional role for short-term in reinforcement learning.

I-34. Single neuron dynamics in primate striatum and prefrontal cortex during classification learning

Yarden Cohen YARDENCSMAIL@GMAIL.COM
Elad Schneidman ELAD.SCHNEIDMAN@WEIZMANN.AC.IL
Rony Paz RONY.PAZ@WEIZMANN.AC.IL

Weizmann Institute of Science

We study dynamics of neural representations for learning of new rules. To do so, we trained monkeys to perform two-alternative forced choice classification of binary visual patterns according to daily changing rules. We then recorded neurons in the dACC and Striatum when the monkeys learned to classify according to novel rules with different complexities. We found that during learning the spiking patterns of neurons in dACC and Putamen became more correlated with the value of the correct category. Classification rules can always be represented as a function of statistically independent features of the stimulus patterns. Accordingly, we characterized the changes in neuronal properties during learning as trajectories of a vector in a high dimensional space whose axes are the correlations of the neural activity and a spanning set of stimulus features. We quantify the dynamics of neuronal encoding during learning in terms of the change in the vector's magnitude and the angle between the rule and the neural activity, and found changes in magnitude or angle to category that were highly significantly correlated with the temporal performance of the monkey in $\sim 25\%$ of the cells. Specifically, we found that in the dACC the learning dynamics is mainly rotation towards the category reflecting a search in the space of available features, whereas in the putamen we found mainly magnitude increase reflecting a policy strengthening and gain of confidence, with some rotation as well. Our results provide a detailed account of the dynamics of neural representations at single cells level in learning and a mathematical framework for characterizing these dynamics that separates the roles of different brain regions.

I-35. A subset of CA1 and subiculum neurons selectively encode rewarded locations

Jeffrey Gauthier JEFF.L.GAUTHIER@GMAIL.COM
David W. Tank DWTANK@PRINCETON.EDU

Princeton University

Lesion studies have demonstrated that the hippocampus is critical for the retention of spatial memories, such as the location of an escape platform in a water maze, as well as transferring those memories to other brain areas. However it remains unclear which features of hippocampal physiology convey this memory. Previous studies have shown that when rewards are presented at fixed locations, a higher density of place fields in CA1 develop near the rewards. We set out to distinguish whether this reflects a generally increased density of place fields, or whether the activity of some cells is specifically related to reward. Mice were trained to run on two virtual linear tracks with rewards at several possible locations, and simultaneous optical recordings were made from two major hippocampal output structures, CA1 and the subiculum. Consistent with previous studies, an increased density of place cells was found near rewarded locations. Interestingly, when environmental changes were made that induced either global remapping or rate remapping, the same subset of neurons maintained firing fields near the reward. These observations reveal a previously undescribed feature of hippocampal remapping: a distinct neural

population that shifts its firing fields to consistently predict a rewarded location, despite unrelated restructuring of simultaneously recorded place fields. A further trial-by-trial analysis revealed that activity of reward-predicting cells correlated with behavioral anticipation of reward, raising the possibility that these cells transmit memory of the rewarded location to other brain areas.

I-36. Predicting human sensorimotor adaptation with synaptic learning rules in a spiking model of cerebellum

Yufei Wu YUFEI.WU13@IMPERIAL.AC.UK
Aldo Faisal ALDO.FAISAL@IMPERIAL.AC.UK

Imperial College London

Sensorimotor adaptation enables us to adjust our movements to external perturbations. The cerebellum is known to play a significant role in adaptation, however, it is still unclear how such learning mechanisms, studied in great detail both experimentally and computationally at the level of human reaching movements, are implemented in the neural substrates of the cerebellum. Here, we present a novel spiking model of the cerebellum using established synaptic learning rules, which predicts human motor adaptation experiments for planar reaching movements, including multi-rate learning in force fields. Our cerebellar model is embodied as it receives sensory-feedback can and can modulate downstream motor signals. The model is made of spiking neuron populations reflecting Purkinje, Granule, Golgi as well as neurons of the Cerebellar nuclei. Synaptic learning rules are based on LTP/LTD and Hebbian learning operating between and within these populations. We model two cerebellar learning mechanisms that modulate this upstream motor command: 1. cerebellar predictive learning (Wolpert et al. 1998) and 2. memory formation (Nagao et al, 2013) - see detailed description below. We test our model and make it perform in simulation the same force-field reaching tasks as used in human studies. We show that our model reproduces key behavioural features of sensorimotor adaptation, including aftereffects (compensatory response persists after switching off perturbations and gradually 'wash-out'), savings (relearning a forgotten perturbation is faster than the first time it was encountered) and multi-rate learning (motor adaptation can be decomposed into a slow learning process that is resistant to wash-out and a fast learning process that compensates rapidly but also forgets rapidly, Smith et al. 2006). Our model reproduces learning curves when applying external force perturbations. We then use this model to predict how patients with specific types of cerebellar injury perform in behavioural experiments.

I-37. A probabilistic theory of deep learning

Ankit Patel ABP4@RICE.EDU
Tan Nguyen MN15@RICE.EDU
Richard Baraniuk RICHB@RICE.EDU

Rice University

Understanding the brain mechanisms underlying perception is a key goal in neuroscience and neuroengineering. While the functioning of an individual neuron has been well studied, much less is understood about how populations of neurons are organized to perform perceptual inference. Recent research studies have shown that the best model of neural activations in ventral stream (V1 to IT) is a Deep Convolutional Network (DCN). This finding promises that a better understanding of DCNs will lead directly to a better understanding of the ventral stream. DCNs are deep learning algorithms for high-nuisance inference tasks; they are constructed from many layers of alternating linear and nonlinear processing units and are trained using large-scale algorithms and massive amounts of training data. The success of deep learning systems is impressive—they now routinely yield pattern recognition systems with near- or super-human capabilities—but a fundamental question remains: Why do they work? Intuitions abound, but a coherent framework for understanding, analyzing, and synthesizing deep learning architectures has remained elusive. We answer this question by developing a new probabilistic framework for deep learning based on a generative probabilistic model that explicitly captures variation due to nuisance variables. The graphical structure of the model enables it to be learned from data using classical expectation-maximization (EM) techniques. Furthermore, by relaxing the generative model to a discriminative one, we can recover two of the current leading deep learning systems, DCNs and random decision forests (RDFs), thus unifying hierarchical generative models with deep convolutional networks. Our work provides insights into the successes and shortcomings of current deep architectures, as well as a principled route to their improvement. In particular, we define a new class of generative DCNs that can learn unsupervised from unlabeled data, and can also execute top-down inference. The implications for neuroscience will also be discussed.

I-38. Emergence of coordinated neural dynamics supports neuroprosthetic skill learning

Vivek Athalye[1,2] VIVEK.ATHALYE89@GMAIL.COM
Karunesh Ganguly[3] KARUNESH.GANGULY@UCSF.EDU
Rui Costa[2] RUI.COSTA@NEURO.FCHAMPALIMAUD.ORG
Jose Carmena[1] JCARMENA@BERKELEY.EDU

[1]University of California, Berkeley
[2]Champalimaud Centre for the Unknown
[3]University of California, San Francisco

Motor learning studies have found that subjects initially produce trial-to-trial variability in both movement and neural activity which decreases with training, resulting in the consolidation of particular movement and neural activity patterns. How does a task-relevant neural population explore and consolidate spatiotemporal patterns supporting skill learning? We analyzed data from an operant Brain-Machine Interface (BMI) study (Ganguly and Carmena, 2009) in which stable recordings from ensembles of primary motor cortex neurons in macaque monkeys were paired with a constant decoder that transformed neural activity to prosthetic movements. Over long-term training, two subjects increased accuracy and speed and reduced trial-to-trial variability of cursor movements. This BMI learning paradigm allows us to model neural variability changes and test how they causally relate to neuroprosthetic skill acquisition via the decoder. To distinguish between private and shared sources of variability which produce independent and coordinated activity, we used Factor Analysis as a generative probabilistic model of observed spike counts. We asked how private and shared variability sources contribute to trial-to-trial firing rate variability, fine-timescale neural pattern consolidation, and ultimately movement. We found that private input signals produced large initial trial-to-trial firing rate variability which reduced significantly over training. Concomitantly, the ratio of shared-to-total trial-to-trial variability increased, and the shared input structure stabilized. Zooming in to finer timescale, we found that task-relevant shared input signals strengthened and consolidated a structure which aligned with the decoder. Finally, we asked how these two sources contributed to movement. On a trial-to-trial basis, we estimated the expected value of shared and private activity and re-simulated the decoding with each source separately. While both sources adapt to produce successful movements, shared input driven movements were significantly faster and straighter. Our findings indicate that over learning, shared input signals are consolidated which coordinate initially variable, high-dimensional activity to drive skilled movement.

I-39. Generalization in goal-directed learning: independent clustering of action-effect and outcome-values

Nicholas Franklin NICHOLAS_FRANKLIN@BROWN.EDU
Michael Frank MICHAEL_FRANK@BROWN.EDU

Brown University

A hallmark of goal-directed behavior is the ability to combine predictions about the effects of actions with outcome values to plan a novel course of action. Adaptive behavior requires these factors to be flexibly combined, especially in unfamiliar contexts for which habitual actions may be less relevant. However, it is often unclear how knowledge from one context should generalize to another. Previous data suggest that rather than learning about specific contexts, humans build latent abstract structures and then learn to link contexts to these structures, facilitating generalization to novel contexts. Computational models further suggest this process involves context popularity-based clustering, such that task structures most popular across contexts are more likely to be revisited in new contexts. However, in ecological settings, novel contexts often indicate some aspects of task structure—such as the effects actions have on subsequent states—should be generalized from one cluster of previous contexts whereas other aspects—such as the value of those states—should be generalized from another cluster. Here, we consider how a non-parametric Bayesian agent can learn and cluster separate latent structures of action-effects and outcome-values, forming independent clusters that may have different popularity across contexts. We show this leads to qualitatively different behavioral predictions than an agent that considers both together. In settings where partial information can be shared between contexts, an agent that clusters action-effects and outcome-values together is susceptible to interference from the partial transfer of information, whereas an agent capable of separate clustering shows robust generalization, speeded learning and a stronger propensity to generalize in new contexts. We develop a novel multistep task to investigate how people can discover latent structure and generalize in a flexible and goal-directed way. We provide preliminary experimental evidence for this model and show that people generalize transition structure independently of reward value.

I-40. A theory of sequence memory in the neocortex

Jeff Hawkins JHAWKINS@NUMENTA.COM
Subutai Ahmad SAHMAD@NUMENTA.COM
Yuwei Cui YCUI@NUMENTA.COM
Chetan Surpur CSURPUR@NUMENTA.COM

Numenta, Inc.

The ability to recognize and predict temporal sequences of sensory inputs is vital for survival in natural environments. Extensive experimental evidence demonstrates that sequence learning occurs in multiple cortical regions, yet the underlying neural mechanism remains obscure. It has been proposed that non-linear properties of dendrites enable neurons to recognize multiple patterns. Here we extend this idea by showing that a neuron with several thousand synapses arranged along active dendrites can learn to accurately recognize hundreds of unique patterns of cellular activity. We propose a neuron model with distinct synaptic integration zones: patterns recognized by proximal dendrites lead to action potentials and define the classical receptive field, whereas patterns recognized by distal dendritic segments act as predictions by slightly depolarizing the neuron without immediately generating an action potential. We then show that a network model using neurons with these properties learns a robust model of time-based sequences. We evaluate the model on both artificial and real-world sequence prediction problems. Our model not only achieves comparable or better accuracy than state-of-the-art sequence prediction algorithms, including ARIMA and LSTM recurrent neural networks, but also exhibits other properties critical for sequence learning. These properties include online learning, the ability to handle multiple simultaneous predictions and branching sequences, robustness to sensor noise, and high fault tolerance. The neuron and network models we introduce are robust over a wide range of parameters because the network uses a sparse distributed code of cellular activations. Our work represents a theory of sequence memory that integrates many fundamental cellular and physiological properties of cortical neurons, and has profound implications for the neural mechanism of sequence learning in the cortex.

I-41. Conservation of neural events across self-initiated, quasi-automatic and cue-initiated movements

Antonio Lara AHL2143@COLUMBIA.EDU
Gamaleldin F Elsayed GFA2109@CUMC.COLUMBIA.EDU
John Cunningham JPC2181@COLUMBIA.EDU
Mark Churchland MC3502@COLUMBIA.EDU

Columbia University

Neurons in primary motor (M1) and dorsal premotor (PMd) cortex exhibit distinct ‘preparatory’ and ‘movement’ responses during delayed-reach tasks, suggesting distinct preparatory and movement-related processes. Most real-life movements, however, are not preceded by a delay period and can be initiated internally (e.g., reaching for a cup of tea) or quasi-automatically (e.g., reaching to catch a falling glass). Whether the lessons of the delayed-reach task generalize to such real-word contexts remains an open question. We recorded from PMd and M1 of two monkeys during a novel task where the same set of reaches was executed in three contexts. In the cue-initiated context, a target appeared and a delayed go-cue indicated when to reach, as in the standard delayed-reach paradigm. In the self-initiated context, the decision regarding when to move was made by the monkey, based on evolving reward size. In the quasi-automatic context, monkeys intercepted a rapidly moving target with no delay or hesitation. At the single-neuron level it was not possible, across all contexts, to confidently distinguish preparatory and movement-related response features. We therefore used a novel dimensionality reduction approach that isolates preparatory and movement-related response components at the level of the population. Remarkably, we observed that the same preparatory state was achieved prior to movement onset, regardless of context. Yet, there were marked timing differences. In the cue-initiated context, preparation lasted most of the delay. In the self-initiated context, preparation was deferred until shortly before the choice to move. In the quasi-automatic case, preparation consumed only a few tens of milliseconds. Yet in all cases the same preparatory events were followed by the same pattern of movement-related dynamics. Thus, preparation appears to be an obligatory processing stage that always precedes movement, regardless of tasks demands. However, the timecourse of preparation is strongly context-dependent and can be remarkably rapid.

I-42. Circuit mechanisms for predicting the sensory consequences of motor sequences

Conor Dempsey CD2656@COLUMBIA.EDU
Nathaniel B Sawtell NS2635@COLUMBIA.EDU
Larry Abbott LFA2103@CUMC.COLUMBIA.EDU

Columbia University

Predicting the sensory consequences of motor commands is believed to be critical for both sensory processing and motor control. However, generating such predictions may be challenging, because the sensory consequences of a motor command depend on the sequence of actions in which the motor command occurs. Here we demonstrate how this problem is solved in the electrosensory lobe (ELL) of mormyrid fish – a system in which mechanisms for predicting sensory consequences of isolated commands is already well understood. Recent work has shown that a brief motor command to fire the electric organ discharge (EOD) leads to temporally extended granule cell corollary discharge responses. Principal cells in ELL form predictions of the sensory consequences of the EOD by sculpting, via anti-Hebbian plasticity, these responses into a negative image of the sensory input caused by the EOD. We demonstrate that ELL learns predictions of the sensory consequences of EOD sequences and that this learning can generalize across sequences. For the system to generalize the granule cell responses must have a history dependence matched to that of the sensory response. Intracellular recordings show that granule cells do indeed have such a dependence. We use an input inference method to reveal that a major factor in shaping the granule cell response to EOD sequences is the dependence of their mossy fiber inputs on recent EOD history. A model of the granule cell population including this feature reproduces the ability of the system to

predict sensory consequences of EOD sequences across frequencies. Our work results in a relatively complete picture of how ELL performs an important computation: learning predictions of the sensory consequences of motor sequences.

I-43. Pharmacogentic silencing of the substantia nigra disrupts control of effort in reward-seeking mice

Kathleen Martin[1] KAM953@NYU.EDU
Jennifer Brown[2] BROWNJ@BERKELEY.EDU
Joshua Dudman[3,4] DUDMANJ@JANELIA.HHMI.ORG

[1]New York University
[2]University of California, Berkeley
[3]Janelia Farm Research Campus
[4]Howard Hughes Medical Institute

Animals learn to perform arbitrary actions with appropriate timing and vigor (velocity and amplitude) to efficiently obtain desired outcomes. The basal ganglia (BG) are thought to control purposive movements. Most models focus solely on the role of BG in movement initiation by feed forward circuitry. However, dysfunction of BG activity results in impairments in timing, initiation, and vigor of movements. It remains unclear how the BG feed forward circuitry contributes to this wide array of properties. Work in vitro has demonstrated that feedback via axon collaterals of the substantia nigra projection neurons (SNr) can modulate the gain of the BG output. These results suggest that both feed forward and feedback connections in the BG could be essential for the control of purposive action; however, the functional role of SNr feedback in behaving animals remains untested. To study the function of SNr feedback in behaving animals, we designed a task in which head-fixed mice receive reward as a function of the vigor of lever pressing. Then, using cell-type specific expression of a pharmacogenetic inhibitor, we selectively inhibited synaptic transmission at either all SNr output targets (globally) or just axon collaterals within SNr (locally) during task performance. Global SNr suppression impaired the tuning of vigor to task demands by producing a large increase in the peak velocity and variance of movement. In contrast, local SNr suppression did not alter the average movement velocity, only variance increased. This finding is consistent with our computational model that feedback inhibition should increase velocity variance (due to a change in the slope of the input/output function) with minimal changes in the mean (minimal change in offset of the input/output function). Together, the modeling and behavioral results indicate that feedback gain control mediated by intranigral collateral synapses could be critical for regulating the vigor of goal-directed behavior.

I-44. Experimentally guided modeling of thalamic DBS to selectively reduce tremor in Essential Tremor

Shane Lee SHANE_LEE@BROWN.EDU
David Segar DAVID_SEGAR@BROWN.EDU
Wael Asaad WAEL_ASAAD@BROWN.EDU
Stephanie Jones STEPHANIE_JONES@BROWN.EDU

Brown University

Essential Tremor (ET) is the most common movement disorder, manifesting in an oscillatory tremor (4–8 Hz) often of the upper limbs, which increases in amplitude upon intentional movement. For medication-refractory ET, deep brain stimulation (DBS) of the cerebellar motor thalamus—the ventral intermediate nucleus (VIM)—can provide therapeutic benefit by reducing tremor amplitude or eliminating tremor altogether. DBS is also used as a treatment for Parkinson's Disease and other disorders and requires invasive, chronically implanted electrodes that stimulate the target at high amplitudes (1–4 V) and frequencies (110–180 Hz). The primary mechanism of therapeutic benefit of DBS is not fully understood, and while DBS provides relief from tremor, it is not universally effective and often results in side effects such as slurred speech and gait imbalance. Why tremor amplitude is increased upon intentional movement and how DBS in VIM works to reduce tremor are not well understood. Here, we combined a center-out joystick task with intraoperative electrophysiological recordings of VIM during DBS electrode implant surgery to understand how dynamic changes in physiological tremor were reflected in tremor frequency local field potential in VIM. We also created a biophysically principled computational model of oscillatory activity in VIM to understand how tremor oscillatory activity could be dynamically modulated by inputs from the cerebellum. We demonstrated how DBS may simultaneously reduce tremor and interfere with physiologically relevant cerebellar inputs, which may result in the observed side effects of DBS in ET. Finally, we simulated a modified DBS protocol that produced a reduction in tremor frequency activity, while improving fidelity of transfer of cerebellar synaptic inputs in the model, which could in principle alleviate stimulation-induced side effects.

I-45. Songbird respiration is controlled by multispike patterns at millisecond temporal resolution

Caroline Holmes[1] CAROLINE.HOLMES@EMORY.EDU
Kyle Srivastava[2,1] KYLE.SRIVASTAVA@GMAIL.COM
Michiel Vellema[3] VELLEMA@BIOLOGY.SDU.DK
Coen Elemans[3] COEN@BIOLOGY.SDU.DK
Ilya Nemenman[1] ILYA.NEMENMAN@EMORY.EDU
Sam Sober[1] SAMUEL.J.SOBER@EMORY.EDU

[1]Emory University
[2]Georgia Institute of Technology
[3]University of Southern Denmark

Although the importance of precise timing of neural action potentials (spikes) is well known in sensory systems, approaches to neural coding in motor control typically focus on firing rates. Here we examined whether the precise timing of spikes in multispike patterns has an effect on motor output in the respiratory system of the Bengalese finch, a songbird. By recording from single motor units (i.e., the muscle fibers innervated by a single motor neuron) in the muscles that control breathing, we find that the spike trains are significantly non-Poisson, suggesting that the precise timing of spikes is tightly controlled. We further find that even a one-millisecond shift of an individual spike in a multispike pattern predicts a significant difference in air sac pressure (the motor parameter controlled by the respiratory muscle). Finally, we provide evidence for the causal relation between precise spike timing and the motor output in this organism by stimulating the motor system with precisely timed patterns of electrical impulses. We observe that shifting a single pulse by as little as two milliseconds elicits differences in resulting air sac pressure. These results demonstrate that the precise timing of spikes does play a role in respiratory motor control.

I-46. Continuous parameter working memory in a balanced chaotic neural network

Nimrod Shaham NIMROD.SHAHAM@MAIL.HUJI.AC.IL
Yoram Burak YORAM.BURAK@ELSC.HUJI.AC.IL

The Hebrew University of Jerusalem

The brain is able to accurately store values of continuous parameters, such as the direction of a bar, the color of an object, or the frequency of a tone, for several seconds. One of the main theoretical proposals for modeling continuous parameter working memory is that local circuits in the brain exhibit continuous attractor dynamics, where different values of a stimulus can be represented by different locations along a continuum of steady states.

It is unclear whether a continuum of steady states can be sustained within the balanced network model of cortical circuitry, which offers an explanation for the asynchronous, Poisson-like activity of single neurons (as observed, for example, during short-term memory tasks in prefrontal cortical areas). Furthermore, it is unclear how the chaotic dynamics, which are generated in a balanced network, may affect memory storage. In this work we find an architecture for which a balanced network can sustain slow dynamics in a certain direction in the mean activities space, making it a candidate for a working memory network. To quantitatively understand the information degradation in the network, we use analytical methods as well as large scale numerical simulations to study the effect of the chaotic noise on the dynamics. We find that the irregular activity of single neurons, arising from the chaotic state of the network, drives diffusive motion along an approximate attractor. We analytically calculate the coefficient of diffusion along the attractor and find that it is inversely proportional to the network size. Thus, for large enough (but realistic) neural population sizes, and with suitable tuning of the network connections, it is possible to obtain time scales which are larger by several orders of magnitude than the single neuron time scale, allowing for accurate information storage over several seconds.

I-47. Emergence of global structures through high order synaptic interactions

Neta Ravid[1] RAVIDNETA@GMAIL.COM
Yoram Burak[2] YORAM.BURAK@ELSC.HUJI.AC.IL

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

Evidence from anatomical and physiological studies indicates that local circuitry in cortical areas is often non-random, and involves correlations in the connectivity of different cells. An important question is whether large scale correlations in the connectivity can be generated autonomously by ongoing plasticity mechanisms, or should necessarily be genetically hard wired. We study the formation of two kinds of structures that contain large scale correlations in the connectivity of different cells: wide synfire chains and self connected assemblies. It has been unclear whether these structures can be autonomously generated via Spike Timing Dependent Plasticity (STDP), which is a local plasticity mechanism that depends on the activity of the pre and post synaptic cells. Most theoretical works that have studied formation of these structures via STDP, included structured external input that induced correlations between the connectivity pattern of different neurons. Here we show that STDP can intrinsically generate large scale correlations and can promote formation of synfire chains and self connected assemblies in an initially unstructured network, without a structured external input. We demonstrate this result using numerical simulations, and a precise analytical theory for STDP dynamics in recurrent networks that we developed. The different structures emerge due to contributions from high order motifs in the network connectivity, that were neglected in previous theoretical treatments. The contribution of structural motifs to plasticity can be quantified given the shape of the synaptic current and the STDP function. Depending on these biophysical properties, different motifs are more dominant relative to others, and different structures emerge. Based on the theory we can predict from the local biophysical parameters which structures will emerge in a stochastic network of spiking neurons.

I-48. Inter-areal balanced amplification for signal propagation in a large scale cortical circuit model

Madhura Joglekar MJ98@NYU.EDU
Xiao-Jing Wang XJWANG@NYU.EDU

New York University

Stable transmission of signals in a multi-area feedforward network represents a long-standing challenge: as a signal propagates across areas, it may either die out or explode (Diesmann et al., Nature 1999, Moldakarimov et al., PNAS 2015). In previous studies on signal transmission, areas are identical and inter-areal connection strengths are not constrained by data. We re-examined the problem of signal propagation using a newly developed large-scale network model of the primate cortex (Chaudhuri et al., Neuron 2015). The model is based on the recently published dataset of directed and weighted connectivity of the macaque cortex from Henry Kennedy's group at INSERM, France (Markov et. al., Cereb. Cortex 2014). This network displays complex feedforward and feedback projections with a wide range of connection strengths, posing new questions about signal propagation. Effective signal transmission is quantified in terms of the amount of attenuation of the peak firing rate as the signal propagates to areas higher in the brain's hierarchy. We found two regimes of stable signal propagation. The underlying principle governing both the regimes is strong long-range excitation that boosts inter-areal signal transmission, followed by strong delayed inhibition that stabilizes the network dynamics. In this sense, this work represents an extension of the balanced amplification mechanism (Murphy et al., Neuron 2009) from local circuits to large-scale cortical systems. Consistent with Murphy et al. showing that this mechanism is characterized by the non-normality of the network connection matrix, in our large-scale model, an improvement in propagation is correlated with an increase in the non-normality measure (Henrici, Numer. Math. 1962) of the large-scale connectivity matrix. Our regimes improve signal propagation by 100 fold, from an attenuation of the order of 1e-4 (Chaudhuri et al., Neuron 2015) to an order of 1e-2.

I-49. The statistics of cortical activity is beneficial for extensive memory storage

Alessandro Barri[1] ABARRI@PASTEUR.FR
David Hansel[2] DAVID.HANSEL@PARISDESCARTES.FR
Gianluigi Mongillo[2,3] GIANLUIGI.MONGILLO@PARISDESCARTES.FR

[1]Institut Pasteur
[2]Universite Paris Descartes
[3]CNRS

Memories are thought to be stored in the cortex by long-term modifications of the pattern and strengths of synaptic connections. They are recalled through self-sustained patterns of neuronal activity (induced by that synaptic structure) triggered by transient, selective external inputs. In this framework, an active memory corresponds to a steady, spatial distribution of firing rates, which is an attractor of the collective network dynamics. A rather puzzling aspect of patterns of cortical activity is their high levels of spatial inhomogeneity and temporal irregularity: Population-averaged firing rates are quite low, distributions of single-cell firing rates are unimodal and right-skewed with a long tail, and spiking is temporally irregular resembling a Poisson process. Though several mechanistic accounts, most notably the balance of excitation and inhibition, have been offered to explain such a regime of activity, its functional advantages are presently unclear. Are the statistical features of neuronal activity in the cortex somehow instrumental to long-term storage? We show, by using a combination of analytical techniques and numerical simulations, that those features are a natural consequence of storing a large number of memories in a distributed neural architecture where each neuron receives a large number of connections. In particular, we show that: (i) The balance of excitation and inhibition is a necessary condition for extensive memory storage in recurrent neuronal networks; (ii) Storage capacity is maximized for right-skewed, long-tailed distributions of firing rates with low averages; (iii) Excitatory neurons must fire at a lower rate than the inhibitory neurons. Importantly, a network of integrate-and-fire neurons operating in this regime, necessarily produces irregular, Poisson-like spiking. Our results suggest, somehow counterintuitively, that efficient memory storage results in network activity characterized by strong spatio-temporal irregularity.

I-50. Criticality signatures in a self organizing recurrent neural network

Bruno Del Papa[1,2] DELPAPA@FIAS.UNI-FRANKFURT.DE
Viola Priesemann[3] VIOLA@NLD.DS.MPG.DE
Jochen Triesch[1] TRIESCH@FIAS.UNI-FRANKFURT.DE

[1]Frankfurt Institute for Advanced Studies
[2]Max Planck Institute for Brain Research
[3]Max Planck Institute for Dynamics and Self-Organization

Recent experiments have suggested that the brain operates close to a critical state, but in a subcritical regime, based on signatures of criticality such as power-law distributions of neuronal avalanches. Although critical systems have been argued to possess exceptional information processing properties, the criticality hypothesis is highly controversial. Several neural network models exhibiting criticality signatures have been proposed, but these typically have highly simplified connectivity structures and do not combine advanced information processing with learning abilities. Here, we investigate criticality signatures in the activity of a self organizing recurrent neural network (SORN), which is known to exhibit spatio-temporal pattern learning, such as grammar learning, while being able to reproduce experimentally observed statistics and fluctuations of synaptic efficacies and various findings on neuronal variability and spontaneous activity. The network consists of excitatory and inhibitory threshold units with synapses evolving due to a combination of spike-timing dependent plasticity and homeostatic mechanisms. We observe power-law distributions for both duration and size of neuronal avalanches, suggesting that the SORN self-organizes to a critical state when the units are driven by weak noise. We also verify that ongoing plasticity is not necessary to maintain these criticality signatures once the network has reached a stationary state. Furthermore, we show that while random noise might contribute to the maintenance of a critical regime, extra structured input sequences can drive the network towards super-criticality matching recent experiments. Overall, our results show for the first time that signatures of criticality are present in the activity of a recurrent network model with advanced sequence learning abilities.

I-51. Touch responses in layer 4 in the barrel cortex break the balance between excitation and inhibition

David Golomb[1] GOLOMB@BGU.AC.IL
Diego Gutnisky[2,3] GUTNISKYD@JANELIA.HHMI.ORG
Jianing Yu[2,3] YUJ10@JANELIA.HHMI.ORG
S Andrew Hires[4] SHIRES@USC.EDU
Karel Svoboda[2,3] SVOBODAK@JANELIA.HHMI.ORG

[1]Ben-Gurion University of the Negev
[2]Janelia Farm Research Campus
[3]Howard Hughes Medical Institute
[4]University of Southern California

The microcircuit of each layer-4 (L4) barrel is a tractable cortical circuit: excitatory (E) and inhibitory neurons within the same cortical column are connected with probabilities of 25-50% within the barrel and rarely to other barrels. The only prominent long-range input originates in VPM thalamic (T) nucleus, which innervates the excitatory and inhibitory fast-spiking (I-FS) L4 neurons. Using extra- and intra-cellular recordings from mice performing an object-localization task, we observed that during transitions from non-whisking to whisking epochs, T and I-FS neurons double their spike rates νT and νI (to 14 and 19 Hz respectively) whereas E neuron spike rates remain unchanged (1.4 Hz). Despite the large difference in whisking response, both T and E neurons fire about 0.6 spikes per contact with short latency (T, 3 ms; E; 6-10 ms). This means that the L4 circuit actively suppresses self-movement signals and enhances touch representation. To explain the dynamical mechanisms underlying the observed circuit responses, we investigate a conductance-based model of a single barrel and its thalamic input. Cell numbers (1600 E, 150 I-FS, 200 VPM), synaptic conductances, connection probabilities and cellular biophysical parameters are based on published anatomical and neurophysiological data. Networks under whisking or non-whisking states behave similarly to large networks with strong synapses, in which νE and νI depend linearly on the νT. In our circuit, however, νE is weakly dependent on νT (as experimentally observed) if the inhibitory-to-inhibitory total synaptic conductance gII is small enough. Excitatory neurons can respond to touch as strongly as thalamic neurons only if there is a "window of opportunity", during which fast thalamic excitation is not balanced by feed-forward inhibition from I-FS neurons. This scenario demands synaptic delays of $\sim 1ms$. Transmission of brief touch events in the thalmocortical circuit results from breaking the excitation-inhibition balance.

I-52. Conservation of cortical network properties throughout development

Matthew Colonnese COLONNESE@EMAIL.GWU.EDU
Hirofumi Watari WATARI@EMAIL.GWU.EDU
Jing Shen SHENJING0619@GWU.EDU

The George Washington University

During development neural networks must grow from a state of zero connectivity to the interconnected circuits that underlie complex dynamics in adults. The basic principles the brain uses to accomplish this are poorly understood. One prominent hypothesis is that exuberant connectivity and weak inhibition render early networks hyper-synchronous, allowing for a rough 'first-draft' connectivity. In this model maturation of synaptic properties and the integration of new cell classes (eg. inhibitory interneurons) into the circuit causes activity to become progressively more precise and less-synchronous allowing for synaptic refinement. A contrary hypothesis is that immature activity provides a model of mature function allowing for formation of correct local connectivity from the outset. In this model synaptic maturation and new cell integrations occurs not to transform network properties, but rather maintain them in the face of increasing connectivity. To distinguish between these hypotheses we have examined early cortical networks for evidence of hyper-synchronization using single units isolated from multi-electrode arrays in the visual cortex of awake mice from the onset of synaptic connectivity at P(ost-natal) day five, through the acquisition of mature activity patterns (P24). We find that, while infant animals (P5-11) superficially appear hyper-synchronous because they spend more time in down-states (network silence), when analysis is restricted to active periods, firing rate, synchronization1, distribution of firing vectors2 and population coupling3 are stable across development. Most surprisingly, the distribution of neuron spike-rate covariance suggests that even during a period of development with weak inhibition, immature networks generate an asynchronous state similar to that observed during adult wakefulness. Together our data support the hypothesis that cortical neurons express functionally adult-like patterns of interconnectivity from the beginning of network formation, and subsequent development maintains these relationships. Our data provide important database for models of cortical network development.

I-53. Intrinsically generated up and down states in a sparsely connected network with strong inhibition

Nicolas Brunel[1] NBRUNEL@UCHICAGO.EDU
Elisa Tartaglia[2] ELISANITAMARIA@GMAIL.COM

[1]The University of Chicago
[2]Institut de la Vision

Electrophysiological recordings have revealed that cortical networks exhibit in some circumstances transitions between two very distinct states called up and down states. Up states are active states, in which neurons are depolarized compared to their resting potential, receive a large amount of excitatory and inhibitory inputs and emit spikes at rates of a few spikes per second. Down states are essentially quiescent states in which most neurons have their membrane potential close to the resting potential and fire very few spikes if any. Typical durations of up and down states in vivo are of the order of 1s. Durations of up states in in vitro preparations are of the same order, but down states tend to be much longer. For the last 15 years, the dominant model to explain the mechanisms of such dynamics relies on strong recurrent excitation leading to bistability between active and inactive states, and adaptation in excitatory neurons. While this scenario reproduces the basic up/down state alternation, it fails to account for several observations: i) average firing rates during up states in such models are typically much higher than those seen in experiments; ii) inhibition plays little role in these models, while a growing body of evidence indicates that inhibition dominates the dynamics of local cortical networks. Here, we show that both observations are naturally reproduced in a sparsely connected network of LIF neurons in an inhibition-dominated regime, in which bistability is induced by the variance of the synaptic inputs, provided the excitatory synaptic time constants are large enough. Transitions between up and down states occur due to noise, and are facilitated by the introduction of adaptation in excitatory neurons. This regime is present in sparsely connected networks, but not in firing rate models nor in large fully connected networks of spiking neurons.

I-54. Neural population dynamics during saccadic behavior in the macaque prefrontal cortex

Aniruddh Galgali[1,2] GALGALIA@INI.UZH.CH
Valerio Mante[1,2] VALERIO@INI.UZH.CH

[1]University of Zurich
[2]ETH Zurich

Single neuron responses in frontal areas of the primate brain exhibit a vast diversity in temporal dynamics and response tuning. Understanding the nature of this response diversity and its role in behavior requires the characterization of neural activity at the level of an entire population. Previous studies in motor cortex have suggested that a dynamical systems approach may capture the nature of neural population activity during limb movement preparation and execution. In this study, we assess the validity of a specific class of dynamical models, the latent variable linear dynamical system (LDS), in characterizing the nature of neural population activity in pre-arcuate cortex during movements of a simpler nature, specifically saccadic eye movements. We fit LDS models to neural population activity that was recorded on an electrode array while monkeys performed a motion discrimination and reported their choice with a saccade to one of two targets. A relatively large number of latent dimensions ($\sim 6 or 7$) are required to explain the patterns of average population activity during saccades. The latent dynamics estimated through the model reflected a gradual rotation across time in state space, spanning different regions depending on the saccade direction. To assess whether the fitted model reflected the true underlying dynamics in the data, we further used the estimated model parameters to generate simulated latent trajectories and compared them to the latent dynamics estimated from the neural data. The simulated trajectories approximately reflected the effects of saccade direction on the estimated latent trajectories, but failed to reproduce how the neural dynamics during movement is affected by reaction time or the state of the neural population before movement onset. Hence, the LDS model revealed surprisingly rich population dynamics during saccades, but failed to capture prominent aspects of the responses, possibly revealing non-linear contributions to the dynamics.

I-55. Homeostatic control of neuronal dynamics: scaling of synaptic strength with network size

Jeremie Barral JEREMIE.BARRAL@GMAIL.COM
Alex Reyes REYES@CNS.NYU.EDU

New York University

Features of sensory input are represented as the spatiotemporal activities of neuronal population. This network dynamics depends on the balance of excitatory and inhibitory drives to individual neurons. Maintaining balance in the face of continuously changing nervous system is vital for preserving the response properties of neurons and preventing neuropathologies. While homeostatic processes are in place to maintain excitatory level, the conditions for maintaining stable responses are yet unclear. Here, we used a culture preparation of cortical neurons to systematically vary the density of cells in the network. Using optogenetic techniques to stimulate individual neurons in the network with high spatial and temporal resolution, we were able to systematically vary the number and correlation of external inputs to drive the network. We found that the average firing rate and the correlation structure of synaptic inputs are invariant with network size. Finally, we used paired recordings to measure the synaptic strengths and connection probability between excitatory and inhibitory neurons. We confirmed experimentally a long standing theoretical assumption that synaptic strength scales with the number of connections per neuron (N) closer to $\sqrt{N}$ than to 1/N in order to maintain network dynamics.

I-56. Towards bifurcation theory for rhythmogenesis in neural networks

Andrey Shilnikov ASHILNIKOV@GSU.EDU
Deniz Alacam DALACAM1@STUDENT.GSU.EDU
Jarod Collens JARODCOLLENS@GMAIL.COM
Aaron Kelley AARNKELLEY@GMAIL.COM
Drake Knapper DKNAPPER1@STUDENT.GSU.EDU
Krishna Pusuluri PUSULURI.KRISHNA@GMAIL.COM
Justus Schwabedal JSCHWABEDAL@GMAIL.COM

Georgia State University

We identify and describe the key qualitative rhythmic states in various network motifs of a multifunctional central pattern generator (CPG). Such CPGs are neural microcircuits of cells whose synergetic interactions produce multiple states with distinct phase-locked patterns of bursting activity. To study biologically plausible CPG models, we develop a suite of computational tools that reduce the problem of stability and existence of rhythmic patterns in networks to the bifurcation analysis of fixed points and invariant curves of a Poincare return maps for phase lags between cells. We explore different functional possibilities for motifs involving symmetry breaking and heterogeneity. Our findings provide a systematic basis for understanding plausible biophysical mechanisms for the regulation of rhythmic patterns generated by various CPGs in the context of motor control such as gait-switching in locomotion. Our analysis does not require knowledge of the equations modeling the system and provides a powerful qualitative approach to studying detailed models of rhythmic behavior. Thus, our approach is applicable to a wide range of biological phenomena beyond motor control.

I-57. A thalamocortical neural mass model of evoked potentials during NREM sleep

Arne Weigenand[1] WEIGENAND@INB.UNI-LUEBECK.DE
Michael Schellenberger Costa[1] SCHELLENBERGER@INB.UNI-LUEBECK.DE
Hong-Viet Ngo[2] HONG-VIET.NGO@UNI-TUEBINGEN.DE
Matthias Molle[1] MOELLE@INE.UNI-LUEBECK.DE
Lisa Marsall[1] MARSHALL@INE.UNI-LUEBECK.DE
Jens Christian Claussen[3] J.CLAUSSEN@JACOBS-UNIVERSITY.DE
Thomas Martinetz[1] MARTINETZ@INB.UNI-LUEBECK.DE

[1]University of Luebeck
[2]University of Tuebingen
[3]Jacobs University Bremen

Much is known about the dynamics of the thalamocortical system during natural sleep, anesthesia and in slice preparations. However, its interaction with sensory stimuli is not fully understood. Recently, auditory stimulation during natural sleep has gained attention, as it can improve memory consolidation by manipulating cortical slow oscillations and sleep spindles. Here, we present a thalamocortical neural mass model which is generates K-complexes, slow wave activity ($<$4 Hz) and fast spindles (12-15 Hz). We incorporated a slow firing rate adaptation into the cortical neural mass and mechanisms for rebound bursts into the thalamic neural mass to account for sleep specific dynamics. The model allows the investigation of responses to auditory stimulation during wake and non-REM sleep. In particular, it reproduces EEG data from phase-independent and closed-loop auditory stimulation of recent sleep studies in humans (Ngo et al., 2013). In the model, transitioning from wake to non-REM sleep corresponds to approaching a Hopf bifurcation, which manifests in the EEG by the slowing of frequencies and an increase in amplitudes. A canard phenomenon and the associated homoclinic orbit determine the shape of K-complexes and slow oscillations. Importantly, a K-complex is a transient event, which corresponds to a single excursion along the homoclinic orbit. In contrast, slow oscillations are noise-driven oscillations around a stable focus. The significance of the work is that it introduces a differentiated view on cortical slow oscillation dynamics in common experimental conditions. The model suggests that during natural sleep the cortex remains predominantly in the active state and only transiently assumes the silent state. A stable silent state, hyperregular relaxation oscillations or genuine bistability are rather found in anesthesia, coma and slice preparations and belong to other parameters of the same system. This view is supported by the model's ability to capture the phase-coupling between spindles and slow oscillations, and evoked responses.

I-58. Multiple mechanisms of theta rhythm generation in a model of the hippocampus

Ali Hummos HUMMOSA@GMAIL.COM
Satish S Nair NAIRS@MISSOURI.EDU

University of Missouri

Hippocampal theta oscillations (4-12 Hz) are consistently recorded during memory tasks and spatial navigation. While computational models suggested specific mechanisms for theta generation, experimental inactivation of these mechanisms did not disrupt theta, precluding definitive conclusions about their roles. We investigated this discrepancy using a biophysical model of the hippocampus that included several of the components implicated in rhythm generation, all constrained by prior experimental results. The CA3 network model included recurrently connected pyramidal cells, and inhibitory basket cells (BC) and oriens-lacunosum moleculare (OLM) cells. The model was developed by matching experimental results characterizing neuronal firing patterns, synaptic dynamics, short-term synaptic plasticity and the three-dimensional organization of the hippocampus. Reciprocal connections to OLM cells generated theta through two mechanisms, and reciprocal connections to BCs generated theta rather than gamma oscillations. The firing dynamics of individual CA3 pyramidal cells strongly influenced the network oscillatory behavior and caused a theta power spectrum signal even in the absence of any connectivity. Diverging input from the Entorhinal cortex produced theta oscillations despite lack of local connectivity, by exploiting the firing dynamics of pyramidal cells. Another novel finding was that the low and high cholinergic states differentially recruited theta generating mechanisms. Consistent with experimental results, inactivation of any single mechanism did not disrupt the rhythm. These results suggested that the theta rhythm is an intrinsic property of the network, and any experimental manipulation or brain state that enhances or suppresses excitation might also, therefore, non-specifically enhance or suppress theta oscillations.

I-59. Robust information propagation in noisy neural population codes

Joel Zylberberg[1] JOEL.ZYLBERBERG@UCDENVER.EDU
Peter Latham[2] PEL@GATSBY.UCL.AC.UK
Alexandre Pouget[3] ALEX.POUGET@GMAIL.COM
Eric Shea-Brown[4] ETSB@UW.EDU

[1]University of Colorado
[2]Gatsby Computational Neuroscience Unit, UCL
[3]University of Geneva
[4]University of Washington

Sensory neurons respond noisily to stimulation, and that noise limits the amount of information available to downstream circuits. Much work has investigated the factors that affect the amount of information encoded in noisy population responses, leading to insights about the role of covariability among neurons, tuning curve shape, etc. However, the informativeness of neural responses is not the only relevant feature of population codes: of potentially equal importance is the issue of how robustly that information propagates to downstream structures. For concreteness, consider the case of the retina transmitting visual information to the cortex via the thalamus. To quantify the retina's performance, one must consider not only the informativeness of the optic nerve responses, but also the ability of LGN neurons to transmit visual information to the cortex based on the inputs they receive from the optic nerve. In other words, only that information which survives the LGN's spike-generating nonlinearity and/or noise corruption will propagate to downstream cortical structures. Our study identifies the set of noise structures for the upstream cells that optimize the ability of information to propagate through noisy and potentially nonlinear circuits. Within this optimal family are noise structures with "differential correlations", which are known to reduce the information encoded in neural population activities. This contrast emphasizes that the adaptations that maximize information in neural population codes, and those that maximize the ability of this information to propagate, can be very different. Our results suggest a new way to interpret the impact of coordinated neural activity on neural population coding: the robustness of encoded information may put important constraints on the set of population activity patterns that encode stimulus information.

I-60. Mapping conceptual knowledge in the human brain with a grid cell code

Alexandra Constantinescu ALEXANDRA.CONSTANTINESCU@MAGD.OX.AC.UK
Jill X O'Reilly JILL.OREILLY@NDCN.OX.AC.UK
Timothy EJ Behrens BEHRENS@FMRIB.OX.AC.UK

University of Oxford

Humans have a remarkable capacity to generalize experiences to novel situations. It has been hypothesized that this capacity relies on a 'cognitive map', allowing conceptual relationships to be navigated in a similar fashion to space (Tolman 1948; O'Keefe and Nadel 1978; Buzsaki and Moser 2013; Eichenbaum and Cohen 2014). Grid cells use a hexagonally symmetric code to organize spatial knowledge (Hafting 2005). Human grid cells have been identified during intra-operative recordings (Jacobs 2013), and are the likely source of a precise six-fold (hexagonal) symmetry in the functional magnetic resonance imaging (fMRI) signal as a function of movement direction during spatial navigation (Doeller 2010; Kunz 2015). This hexagonally symmetric signal can be found in a network of brain regions that are also activated during abstract reasoning processes, such as memory (Binder 2009), imagination (Schacter 2012), valuation (Clithero and Rangel 2014) and social cognition (Saxe 2004). Whether conceptual knowledge is also organized by grid-like codes remains unknown. Here we found that humans (n=28, 94 independent datasets) navigating in an abstract space show the same fMRI hexagonally symmetric signal in a strikingly similar set of brain regions to those activated during spatial navigation. The abstract space was continuous, two-dimensional, and consisted of bird-shaped stimuli, with the legs and neck lengths being the features of the two orthogonal dimensions. Subjects navigated through this space by morphing these features to generate bird shapes previously associated with rewards, by analogy to a standard stimulus-outcome task. The fMRI hexagonal signal was consistent across sessions acquired hours and more than a week apart. These findings suggest that grid-like codes, known to organize spatial knowledge, can also organize abstract knowledge that is difficult to study in nonhuman species. These results raise the possibility that such global relational codes provide the organizing principles for other types of knowledge.

I-61. Uncertainty coding in a model of auditory localization

Ruben Coen-Cagli RUBEN.COENCAGLI@UNIGE.CH
Guillaume Dehaene GUILLAUME.DEHAENE@GMAIL.COM
Alexandre Pouget ALEX.POUGET@GMAIL.COM

University of Geneva

There is accumulating evidence that the nervous system encodes the uncertainty of the variables it represents. But which aspects of the neural response encode this uncertainty? Experimentally, this question has often been investigated by manipulating the uncertainty of the stimulus while monitoring the tuning curves of the neurons. In the auditory system, the width of the tuning curves to azimuth has been found to increase as the uncertainty of the stimulus azimuth increases, a result that has been reported in other systems as well. This code is appealing because it would suggest that tuning curves are proportional to the posterior distribution, or likelihood function of the encoded variable. We show here that this encoding approach can be misleading. We use a standard cross-correlation model for computing the azimuth of a sound source based on inter aural time difference (ITD) which we show implements the ideal observer. In this model, the width of the tuning curves does increase in proportion to the average uncertainty of the azimuth but contributes only marginally to the representation of the ITD. Indeed, ITD log-likelihood can be linearly decoded from the population activity, even when the decoder assumes the wrong tuning width. If we instead ignore the covariance between neurons, the performance of the decoder is considerably reduced which demonstrates the critical contribution of pairwise correlations. This remains true if we only decode the width of the likelihood function (as opposed to the full likelihood function): a linear decoder of population activity achieves almost perfect performance while a decoder of population activity width fails. This work illustrates the danger of the encoding approach to uncertainty while showing that, in a realistic model of auditory localization, uncertainty can be extracted linearly from the neural response, a type of code known as a linear probabilistic population code.

I-62. Low trial-to-trial variability in stimulus-encoding dimensions in macaque primary visual cortex

Benjamin Cowley[1] BCOWLEY@ANDREW.CMU.EDU
Douglas Ruff[2] RUFFD@PITT.EDU
Marlene Cohen[2] COHENM@PITT.EDU
Tai Sing Lee[1] TAI@CNBC.CMU.EDU
Adam Kohn[3] ADAM.KOHN@EINSTEIN.YU.EDU
Matthew Smith[2] SMITHMA@PITT.EDU
Byron Yu[1] BYRONYU@CMU.EDU

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

The trial-to-trial variability of neural activity can limit the fidelity of information encoded about a sensory stimulus, but only if that variability lies along the same axes by which the sensory stimulus is decoded (i.e., stimulus-encoding dimensions). One way to assess the extent to which trial-to-trial variability lies in the stimulus-encoding space is to compare population information in raw data, and data in which variability has been altered (e.g.

shuffling; Averbeck et al., 2006; Moreno-Bote et al., 2014). Here, we adopt a complementary multi-dimensional statistical perspective to directly identify stimulus-encoding dimensions within a high-dimensional population activity space, and measure the amount of trial-to-trial shared variability (i.e., variability shared across neurons) within and outside of the stimulus-encoding space. This approach is equally applicable to stimuli that can be simply parameterized (e.g., drifting gratings) and to stimuli for which there is no simple parameterization (e.g., natural movies). We applied this approach to the population activity of dozens of neurons recorded simultaneously in the primary visual cortex (V1) of anesthetized and awake macaque monkeys in response to drifting gratings and natural movies. We found that even though trial-to-trial shared variability was large, the single-trial population activity was reliable in the stimulus-encoding dimensions, but not in other dimensions. This was true both in anesthetized and awake animals and for both drifting gratings and natural movies. These empirical observations suggest that underlying circuit mechanisms may shape the trial-to-trial variability of neural activity, such that it is minimized in dimensions that matter for stimulus encoding.

I-63. Investigating the structure of the retinal ganglion cell population response probability landscape

Adrianna Loback ADRIANNA@PRINCETON.EDU
Jason Prentice JASONSP@PRINCETON.EDU
Mark Ioffe MIOFFE@PRINCETON.EDU
Michael Berry BERRY@PRINCETON.EDU

Princeton University

Throughout the brain, local circuits encode information using large populations of neurons; the nature of this population code remains a central question in neuroscience. Advances in multi-electrode array technology have enabled simultaneous recording from large ($>$100 neurons), complete neural populations in the vertebrate retina. Past work modeling the joint probability distribution of retinal ganglion cell (RGC) population activity found a proliferation of local maxima, suggesting a potential role of the corresponding response configurations in neural coding. However, these studies focused on responses to highly repeated stimuli. We have found that the qualitative structure of the RGC population response probability landscape strongly depends on the stimulus ensemble. Specifically, for mildly- and non-repeated ensembles, the landscape is dominated by sparseness, largely precluding the formation of true local maxima. Here we use a recently-developed form of hidden Markov model to investigate the geometric structure of the RGC population response probability landscape for the case of non-repeated stimuli. We find that the modeled probability landscape is comprised of a proliferation of 'ridges,' which consist of sets of adjacent saddle points. Moreover, our results suggest that these ridges correspond with neuron 'communities,' a notion from computer science (Fortunato 2010). Specifically, neurons in the population cluster into groups that are specific to each ridge. Population responses on a ridge exhibit active neurons that are subsets of this identifiable group, combined with silence from all the neurons outside of this group. Our findings extend the class of naturalistic stimuli for which statistical structure in the RGC population response landscape has been reported, and suggest that ridges or communities may be an important feature of the output population activity extracted by visual processing areas downstream of the retina. This type of statistical structure could facilitate clustering, and may point to a general neural coding principle.

I-64. Disentangling the contributions of multiple noise sources to neuronal variability

Alison Weber[1] AIWEBER@UW.EDU
Eric Shea-Brown[1] ETSB@UW.EDU
Fred Rieke[1,2] RIEKE@UW.EDU

[1]University of Washington
[2]Howard Hughes Medical Institute

Noise in the nervous system places fundamental limits on the fidelity with which information can be encoded and transmitted. Understanding the sources of noise that contribute to a neuron's responses is therefore critical for understanding the computations performed by a neural circuit. Noise arises from a number of sources, both internal (in noisy biophysical processes) and external (in noisy sensory inputs), occurring at multiple stages of processing. This noise is often modeled as Poisson noise at the output stage of a model (e.g., as in linear-nonlinear-Poisson, or LNP, models). However, assuming that all noise arises at a single stage of processing is problematic. First, an accurate description of where noise enters a circuit is necessary for correctly inferring a neuron's nonlinear input-output function. Second, the location at which noise enters a circuit can have dramatic implications for how that circuit should optimally transform its inputs. For these reasons, it is important to develop statistical models of neural coding that incorporate accurate descriptions of where noise originates. Few existing analytical tools can disentangle the contributions of multiple sources of noise to the variability in a single neuron's responses. Yet the impact of noise on neural coding has been an area of much recent interest. Here, we present a flexible model of neural responses that incorporates multiple sources of noise. We first demonstrate that traditional methods used to estimate nonlinearities fail in the presence of multiple sources of noise, while simultaneous fitting of nonlinearity and noise parameters more accurately recovers the true underlying nonlinearity. Next, we fit this model to mouse retinal ganglion cell (RGC) responses at different mean light levels. This reveals that the contribution of different noise sources changes as environmental conditions change and allows us to test predictions about the optimality of processing in this circuit.

I-65. Firing rate nonlinearity optimizes decoding of orientation under nuisance parameter uncertainty

Gergo Orban[1] ORGERGO@GMAIL.COM
Merse Gaspar[1] GASPARMERSE@GMAIL.COM
Pierre-Olivier Polack[2] POLACK.PO@RUTGERS.EDU
Peyman Golshani[3] PGOLSHANI@MEDNET.UCLA.EDU
Mate Lengyel[4] M.LENGYEL@ENG.CAM.AC.UK

[1]MTA Wigner Research Centre for Physics
[2]Rutgers University
[3]University of California, Los Angeles
[4]University of Cambridge

The contribution of single neuron properties to cortical population codes is a central question bridging cellular and systems neuroscience. For example, the threshold-power law nature of the firing rate nonlinearity (FRN) of V1 neurons has been suggested to contribute to their contrast invariant tuning curves, high response variability and low noise correlations. However, it is unknown how these nonlinearities change the overall format of the population code, i.e. the decoding strategies that may be appropriate for reading it. Here we investigated the effect of the FRN on the linear decodability of a population of V1 simple cells. We focussed on decoding in the presence of so-called 'nuisance' parameters, stimulus features other than the decoded one that still influence neural responses and that are ubiquitous in real-world tasks. A canonical example of this is Fourier phase, which most previous studies of decoding have ignored. We show that linear decoding of orientation from V1 simple cell membrane potentials (MPs) is efficient when phase is held constant, but fails in the face of phase uncertainty. In contrast, the firing rate population code affords high linear decoding accuracy even when phase was unknown. Importantly, this increased efficiency is relatively insensitive to the precise form of the FRN but altered MP variability entails parallel changes in FRN threshold to ensure optimal decoding performance. Using intracellular recordings from simple cells of behaving mice we show that the dependence of firing threshold on MP variability is similar to that predicted to be optimal by the theoretical analysis. These results suggest that FRNs can greatly enhance the robustness of population codes to features to which individual neurons are highly sensitive.

I-66. Computing with a scale-invariant representation of time, space, and number

Marc Howard MARC777@BU.EDU
Karthik Shankar SHANKARK@BU.EDU
Zoran Tiganj ZORANT@BU.EDU

Boston University

The flexibility of cognitive computation presents a deep challenge to traditional connectionist learning models. To understand computation in the brain, we must know how quantities are represented. Functions over physical quantities are often represented by the firing rate of a population of neurons. Each neuron supports a particular range of the external quantity x centered on xi. Many one-dimensional variables obey what we refer to as Weber-Fechner coding, such that $xi{-}x{-}1 \propto xi$. This results in a scale-invariant logarithmic scale internal scale for x. We describe a theoretical rationale for the ubiquity of Weber-Fechner coding and sketch a framework for flexible cognitive computing on these representations. Consider the problem of representing some function f(x) characterized by some unknown scale a. We show that a set of receptors with Weber-Fechner coding equalizes the amount of information carried about the signal independent of a. Neurophysiological evidence suggests that Weber-Fechner coding may hold not only for variables such as retinal position that depend on the form of a sensory organ, but also for 'hidden variables' such as time. These hidden variables, including time, space and number, require another strategy to implement Weber-Fechner coding. We have introduced a framework for representing hidden variables that requires access to the Laplace transform of the to-be-represented function. Because of the computational efficiency of the Laplace domain, this mathematical framework also facilitates flexible cognitive computation. We describe a operators for translation of a function, addition of two functions (implemented via convolution) and subtraction of two functions (implemented via cross-correlation). Taken together, this work suggests a general framework for representing and computing with one-dimensional variables.

I-67. How many cortical neurons must we record?

Marius Pachitariu[1] MARIUS10P@GMAIL.COM
Carsen Stringer[2] CARSEN.STRINGER@GMAIL.COM
Sylvia Schroder[1] SYLVIA.SCHROEDER@UCL.AC.UK
Matteo Carandini[1] MATTEO@CORTEXLAB.NET
Kenneth Harris[1] KENNETH.HARRIS@UCL.AC.UK

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

Recent advances in optical imaging have made it possible to record the activity of thousands of neurons simultaneously. However, modelling analyses suggest that a much smaller set might be sufficient to capture the major structure of neural population activity, due to high levels of correlation in neuronal populations. To understand the dimensionality and structure of population activity in cortex, we simultaneously recorded the spontaneous activity of $\sim 10,000$ neurons at a 2.5 Hz sampling rate in awake mouse V1, using transgenic mouse lines expressing the high SNR indicator GCaMP6s only in pyramidal neurons. We also developed automated methods

for cell identification, available at https://github.com/marius10p/Suite2P/. We found that although neural activity was indeed correlated, key aspects of the structure of correlations could be seen only by recording from large populations. Specifically, in a population of $\sim 10,000$ neurons we found more than 50 intermingled and largely non-overlapping subnetworks of highly correlated neurons. These subnetworks extended over large regions of horizontal and vertical cortical space, and might thus constitute fundamental building blocks of visual cortical circuitry. Arousal level (indicated by pupil size) appeared to turn off some subnetworks while activating others, and the modulated subnetworks changed over the course of the two hour recordings and on each individual pupil dilation. We were able to reproduce this activity structure in a simulated spiking network with clustered connectivity and top-down modulation. The strong and clustered correlation structure we observed meant that the activity of any one neuron could be well predicted from as few as 100 neurons within the same subnetwork. Nevertheless, unless sufficient neurons from the correct subnetwork were recorded, predictions remained poor. Thus, despite its highly correlated nature, tens of thousands of neurons must be recorded simultaneously to characterize the structure of cortical population activity.

I-68. Photon and cortical noises limit what we see

Denis Pelli[1] DENIS.PELLI@NYU.EDU
Manoj Raghavan[2] MRAGHAVAN@MCW.EDU

[1]New York University
[2]Medical College of Wisconsin

Random fluctuations arise at every stage of visual processing, but does this noise restrict what we see? Adapting a technique developed by radio and television engineers in the 1940s, we measure the effect of added visual noise on grating and letter identification, and estimate the amount of added noise that is equivalent to the observer's intrinsic noise. We map out the equivalent noise as a function of the signal size, duration, and luminance, and compare our measurements to the predicted equivalent noise of several intrinsic noises arising at various levels in the visual system. The results show that visual sensitivity is limited by the sum of photon and cortical noises, each dominating a distinct stimulus domain. The visibility of small brief dim signals is limited by photon noise, while the visibility of large prolonged bright signals is limited by cortical noise. The boundary separating photon and cortical noise domains is a critical amount of light within the signal: the product of luminance, area, and duration. In the photon-noise domain, the measured equivalent noise tells us that vision uses only 2% of the light entering the eye. The cortical equivalent noise is self-similar, and scales with luminance, size, and duration, affirming suggestions that cortical processing is scale-invariant. In brighter light, we can see smaller, briefer, and fainter objects, but, even in the brightest light, we will never see an edge whose two sides differ by less than 1% in luminance, the faintest that the cortical neurons can detect in their own noise. These two noises, photon and cortical, largely account for the sensitivity of achromatic photopic foveal vision. Correlated noise may dominate physiological recordings while hardly affecting psychophysical threshold, underscoring our comparison of these psychophysical estimates with recordings from V1 neurons (Lin et al. 2015; Goris et al. 2014).

I-69. Transcriptional profiling of functionally characterized neurons in mouse primary visual cortex

Petr Znamenskiy PETER.CSHL@GMAIL.COM
Thomas Mrsic-Flogel THOMAS.MRSIC-FLOGEL@UNIBAS.CH

University of Basel

A universal hallmark of sensory cortex is the selectivity of individual neurons for particular features of sensory scenes. On one hand, individual neurons might be born with equal potential to encode different features of sensory stimuli and later acquire their unique properties through stochastic and activity dependent mechanisms. On the

other hand, some cells may specialize in processing particular types of sensory inputs predetermined by their molecular makeup. To investigate this latter possibility, we have developed an approach that allows us to probe the relationship between functional properties and gene expression of single neurons. Focusing on excitatory neurons in layer 2/3 of mouse primary visual cortex, we characterize the responses of neurons to visual stimuli using two-photon calcium imaging, including tuning to orientation, direction, spatial, and temporal frequency. We then identify the imaged neurons in acute brain slices, harvest the cellular contents by microaspiration, and measure gene expression by single cell RNA sequencing. This approach allows us to probe how visual response properties relate to gene expression patterns in an unbiased manner and will help identify molecular markers associated with selectivity for specific sensory features.

I-70. Frames of reference in multisensory spatial perception: how eye position influences spatial priors

Brian Odegaard ODEGAARD.BRIAN@GMAIL.COM
Jason Carpenter JCARPENTER542@GMAIL.COM
Ladan Shams LSHAMS@PSYCH.UCLA.EDU

University of California, Los Angeles

As we observe the surrounding world, each sensory modality initially encodes spatial position with respect to a different frame of reference: vision encodes information with respect to the retina, while audition encodes information with respect to the head. As stimuli in the world are generally multisensory, the brain must somehow decide upon a common representational space to use in order to integrate sensory signals (Pouget, Deneve, & Duhamel, 2002). In principle, three possibilities exist: the brain may employ (1) an eye-centered reference frame, (2) a head-centered reference frame, or (3) a hybrid reference frame (Deneve & Pouget, 2004; Kopco et al., 2009). In order to shed light on this question, we examined the reference frame used by the brain for encoding the prior bias in perception of space. In two experiments, subjects were asked to fixate either a central position (aligning the two reference frames) or one of two eccentric positions (misaligning the reference frames) and were required to localize visual, auditory, or audiovisual stimuli. In Experiment 1, subjects could move their eyes during response, but in Experiment 2, they were required to continue to fixate during response. We quantitatively estimated the position of the spatial prior bias for each individual observer under different gaze conditions using the Bayesian Causal Inference model (Wozny et al., 2010). As with previous studies (Kording et al., 2007), results revealed a prior bias for the straight-ahead position when gaze and head were aligned. However, when the gaze was misaligned with the head (by $13\hat{}\circ$) the spatial prior was partially shifted in the direction of gaze. These results provide evidence that priors involved in the perception of auditory/visual stimuli are encoded in a hybrid reference frame, suggesting that both head and eye-centered coordinates influence the reference frame used by the brain for encoding object locations.

I-71. Feature-coding transitions to conjunction-coding with progression through visual cortex

Rosemary Cowell[1] RCOWELL@UMASS.EDU
John Serences[2] JSERENCES@UCSD.EDU

[1]University of Massachusetts, Amherst
[2]University of California, San Diego

Abundant evidence suggests that the ventral visual object processing pathway in cortex analyzes incoming information in a staged, hierarchical manner: Neurons in early stages are tuned to simple stimulus features whereas neurons in later stages demonstrate selectivity for increasingly complex stimulus attributes. However, it is not clear how the complex object representations in later stages emerge from earlier feature representations. Are complex representations simply a summation of the simple features represented upstream (feature-coding)? Or do complex representations combine simple features into conjunctions for which the whole is greater than the sum of its feature-level parts (conjunction-coding)? We present a novel method for multivariate pattern analysis (MVPA) of functional brain imaging data that measures both feature-coding and conjunction-coding and pits them against each other for a single set of visual stimuli. Scanned participants viewed 'conjunctive' visual objects, composed of simple features known to be represented in early visual cortex. We trained classifiers to predict what a scanned participant was viewing, based on the activation patterns in visual cortex; one classifier used the patterns to predict which whole object was presented, whereas another set of classifiers used the patterns to predict which features the presented object contained. We measured conjunction-coding throughout visual cortex by asking whether the predictions of the object-classifier were more accurate than the prediction obtained by combining the separate feature-classifiers. If so, conjunction-coding was inferred; if the reverse was true, feature-coding was inferred. The results provide the first direct demonstration of a continuous gradient from feature-coding in primary visual cortex to conjunction-coding in inferior temporal and posterior parietal cortices. The method enables the use of experimentally controlled visual features such as orientation, spatial frequency and contour to investigate population-level conjunction-codes throughout human cortex.

I-72. Emergence of transformation-tolerant representations of visual objects in rat visual cortex

Sina Tafazoli[1] TAFAZOLISINA@GMAIL.COM
Houman Safaai[1] HOUMAN.SAFAAI@IIT.IT
Gioia De Franceschi[1] G.FRANCESCHI@UCL.AC.UK
Federica B Rosselli[1] FROSSELL@SISSA.IT
Margherita Riggi[1] MARGHERITA.RIGGI@SISSA.IT
Federica Buffolo[1] FEDERICA.BUFFOLO@IIT.IT
Stefano Panzeri[2] STEFANO.PANZERI@IIT.IT
Davide Zoccolan[1] ZOCCOLAN@SISSA.IT

[1]SISSA (International School for Advanced Studies)
[2]Istituto Italiano di Tecnologia

Recent behavioral studies have uncovered unexpected visual object recognition abilities in rats, thus arguing for the existence of cortical machinery that is specialized for pattern vision in this species. The latest anatomical and neurophysiological findings about mouse visual cortex support this notion, and point to a succession of lateral extrastriate areas (V1->LM->LI->LL) as a candidate shape-processing pathway in rodents. Yet, evidence that these areas are able to support the core requirement of object vision (i.e., the recognition of visual objects despite identity-preserving transformations) is very limited. In this study, we addressed this issue, by recording hundreds of neurons along the V1->LM->LI->LL progression of anesthetized rats, presented with many transformations (or views) of different objects. Consistently with the existence of a processing hierarchy, we found a gradual increase of receptive field size and response latency along this progression, with the median of both properties being approximately twice as large in LL than in V1 (i.e., respectively, $\sim 15\hat{}\circ$ in V1 vs. $\sim 30\hat{}\circ$ in LL, and $\sim 40ms$ in V1 vs. $\sim 75ms$ in LL). More importantly, for each neuron, we trained a binary decoder to discriminate two randomly sampled views of a pair of objects, and then we tested its ability to discriminate other views of the same objects (this generalization test was repeated across many combinations of object pairs and views). This analysis revealed a gradient in the ability of the four visual areas to support transformation-tolerant recognition, with LI and LL neurons yielding median generalization performances that were, respectively, ~ 1.5 and ~ 2.5 times larger than in V1 and LM. These results reveal a specialization of rat lateral extrastriate areas for processing object information in a way that becomes increasingly tolerant to variation in object appearance, thus suggesting a functional homology between the V1->LM->LI->LL progression and the primate ventral stream.

I-73. Organization of ON and OFF inputs in visual cortex enables an invariant columnar architecture

Kuo-Sheng Lee[1,2] KUO-SHENG.LEE@MPFI.ORG
Xiaoying Huang[1] SHARONHUANG263@GMAIL.COM
David Fitzpatrick[1] DAVID.FITZPATRICK@MPFI.ORG

[1]Max Planck Florida Institute for Neuroscience
[2]Florida Atlantic University

Circuits in visual cortex integrate the information derived from separate ON and OFF pathways to construct orderly columnar representations of orientation and visual space. How this transformation is achieved to meet the specific topographic constraints of each representation remains unclear. Here we report several novel features of ON/OFF convergence visualized by mapping the receptive fields of layer 2/3 neurons in tree shrew visual cortex using two-photon imaging of GCaMP6 calcium signals. The spatially separate ON and OFF subfields of simple cells in layer 2/3 were found to exhibit topologically distinct relationships with the maps of visual space and orientation preference. The centers of OFF subfields for neurons in a given region of cortex were confined to a compact region of visual space and displayed a smooth visuotopic progression. In contrast, the centers of the ON subfields were distributed over a wider region of visual space, displayed significant visuotopic scatter, and an orientation-specific displacement consistent with orientation preference map structure. As a result, cortical columns exhibit an invariant aggregate receptive field structure: an OFF-dominated central region flanked by ON-dominated subfields. This distinct arrangement of ON- and OFF- inputs enables continuity in the mapping of both orientation and visual space and the generation of a previously unrecognized columnar map of absolute spatial phase. Analysis of layer 2/3 cell receptive field surround effects reveals that importance of absolute spatial phase in specifying the impact of contextual modulation.

I-74. Synergistic adaptation by synaptic transmission and spiking

Bongsoo Suh BSSUH@STANFORD.EDU
Stephen Baccus BACCUS@STANFORD.EDU

Stanford University

To use their dynamic range more efficiently, neural systems adapt to the strength of their inputs by changing their amplification. It is unknown, however, how neural mechanisms such as synaptic transmission and spiking work together to cause adaptive changes in gain. Retinal ganglion cells adapt to temporal contrast by decreasing their gain at high contrast to avoid saturation in two consecutive stages, the bipolar cell synaptic terminal and ganglion cell spiking. To understand how these two adaptive mechanisms interacted, we developed a four stage Linear-Nonlinear-Kinetic-Spiking (LNKS) model consisting of a linear temporal filter, a static nonlinearity and first-order kinetic model that captures the presynaptic threshold and synaptic vesicle release, (Ozuysal and Baccus 2012) and a spiking stage with feedback. The model accurately predicted the membrane potential and spiking of ganglion cells measured intracellularly, and captured both stages of adaptation. In the model, we computed gain changes at the synapse and in spiking during contrast adaptation, and found paradoxically that spiking alone was anti-adaptive, in that it increased its gain with higher contrast. Furthermore, the synapse and spiking showed great synergy, in that if gain changes at the synapse and spiking were to act independently, the two stages together would show no adaptation at all, whereas in fact the two stages together showed twice the adaptation of the synapse alone. We explain this paradox by dynamic changes in the coincidence of synaptic transmission and spiking at different contrasts. At low contrast, prolonged synaptic release continued while the cell crossed spiking threshold, whereas at high contrast rapid synaptic depletion caused the synapse to act at low gain when the cell exceeded spiking threshold. As the ingredients for these effects were fast synaptic depletion and a spiking threshold, we expect that synergistic adaptation of synapses and spiking is widespread in the nervous system.

I-75. Rapid gain adaptation optimizes pursuit accuracy

Bing Liu LIUBING@UCHICAGO.EDU
Matthew Macellaio MACELLAIO@UCHICAGO.EDU
Leslie Osborne OSBORNE@UCHICAGO.EDU

The University of Chicago

In a rapidly changing world, the statistics of sensory stimuli can fluctuate across a wide range. Theoretically, in order to maximize the information sensory neurons can transmit, they should rescale their sensitivity to input fluctuations dynamically, allocating their limited response bandwidth to the current range of inputs. Such adaptive coding has been observed in a variety of systems, but the premise that adaptation optimizes behavior has not been tested. Here we show that adaptive rescaling maximizes information about visual motion in cortical MT neurons and, importantly, in pursuit eye movements guided by that cortical activity. We use time-varying motion signals that transition between different levels of direction variance and record isolated, extrastriate cortical area MT neurons and we record pursuit eye movements in monkeys. We find that adaptation drives a rapid ($<$100ms) recovery of motion information after steps in variance because both neurons and behavior rescale sensitivity to compensate for differences in direction variance. We find that MT neurons adopt a response gain, a change in firing rate per degree of direction change that maximizes information about motion. We find that pursuit also adapts to a response gain that maximizes the mutual information between eye and target movements and that minimizes tracking errors. Thus efficient sensory coding is not simply an ideal standard but rather a compact description of real sensory computation that manifests in improved behavioral performance.

I-76. Signals in IT reflect visual familiarity memories acquired after single image viewings

Travis Meyer TRMEYER@SAS.UPENN.EDU
Nicole Rust NRUST@SAS.UPENN.EDU

University of Pennsylvania

Our ability to remember the tens of thousands of objects and scenes that we have encountered before is remarkable, particularly given that we robustly store these memories after single image exposures. Previous investigations into the neural mechanisms underlying single-trial visual familiarity memory have been limited in two ways: (1) traditional (largely single-neuron) analyses rely on averaging noisy neuronal responses across many trials whereas a ‘novel’ condition cannot be repeated multiple times; and (2) because no published study to date has explored the relationship between putative familiarity signals and behavior, it remains unclear whether these signals relate to familiarity percepts. To address these issues, we applied single-trial, population-based approaches to analyze neural responses in inferotemporal cortex (IT) collected as monkeys performed a visual familiarity task. Similar to humans, monkeys were highly capable of reporting whether individual images were ‘novel’ (never seen before) or ‘familiar’ (seen exactly once). To quantify the IT familiarity signal, we computed the cross-validated performance of a linear population read-out to classify images presented as ‘novel’ versus ‘familiar’, applied to 24 simultaneously recorded neural sites in each experimental session (across 21 recording sessions). We found that on average, novel and familiar conditions were correctly classified on trials in which the monkey reported the correct answer, whereas novel and familiar conditions were significantly misclassified on trials when the monkey made errors. Additional analyses revealed that the IT familiarity signal was nearly exclusively carried by neural responses that were lower for familiar versus novel images. These results demonstrate that signals exist within IT that co-vary with familiarity judgments based on memories acquired after single image viewings. Additionally, the fact that these familiarity memory signals resemble the well-documented phenomenon of ‘sensory adaptation’ provides support for the compelling proposal that a functional role of sensory adaptation is familiarity memory storage.

I-77. Predictive information in the retina depends on stimulus statistics

Jared Salisbury[1] JARED.SALISBURY@GMAIL.COM
Stephane Deny[2] STEPHANE.DENY.PRO@GMAIL.COM
Olivier Marre OLIVIER.MARRE@GMAIL.COM
Stephanie Palmer STEPHANIE.E.PALMER@GMAIL.COM

[1]The University of Chicago
[2]Institut de la Vision

Predicting the future state of the environment is a central challenge for neural systems, both for overcoming sensory delays due to signal transduction and for guiding future behavior. The task is so crucial that we find evidence for prediction at the sensory periphery, in the form of motion anticipation in the spiking activity of retinal ganglion cells, suggesting that the retina may be optimized for prediction. To explore this hypothesis, we use information theory to analyze ganglion cell responses in a highly simplified visual environment, consisting of a moving bar whose trajectory has components that are both predictable and random. The statistics of the trajectory determine a bound on the information that neural responses to the past stimulus can contain about the future stimulus. We have previously demonstrated that the responses of groups of neurons in the salamander retina saturate this bound for one particular statistical environment. Here we explore the limits of the retina's capacity for prediction, by adjusting the parameters of the stimulus trajectory, this time using the rat retina. We find that increasing the timescales of correlations in the stimulus dramatically increases the retina's performance as a predictor, presumably by allowing built-in predictive mechanisms to exploit these statistical regularities. We observe a substantial decrease in the effective information processing lag in response to the long correlation time stimulus, as well as a saturation of the aforementioned bound that is not seen in response to the short correlation time stimulus. These results indicate that the retina does indeed encode information about the future optimally, but only for a restricted range of stimulus statistics, pointing to the importance of the animal's environment and behavior in constraining the prediction problems its brain must solve.

I-78. Models of disparity computation in the visual cortex: Computational-level analysis and electrophysiology

Junkyung Kim JUNKYUNG_KIM@BROWN.EDU
Thomas Serre THOMAS_SERRE@BROWN.EDU

Brown University

The past decades of research in visual neuroscience have generated a large and disparate body of literature on binocular disparity computation in the primary visual cortex. Several models and theories have been proposed to account for specific biological mechanisms or phenomena, yet a theoretical framework that would link neural circuits to computational stages of the stereo correspondence problem (Marr & Poggio, 1976, 1979) is still lacking. Here, we consider a repertoire of elementary operations to derive a minimal set of circuit-level models to solve the stereo correspondence problem. These models effectively extend the classical disparity energy model (Ohzawa et al., 1990, Qian 1994) via diverse combinations of excitatory and suppressive mechanisms (Tanabe et al., 2011). We first assess the consistency between models and several neurophysiological studies on binocular gain control (Trushard et al. 2000), suppressive mechanisms (Tanabe et al. 2011) and attenuation to anti-correlated stimuli (Cumming & Parker 1997). We then consider the models' ability to satisfy a set of mathematical desiderata derived from Marr & Poggio's ideal correspondence detector. Lastly, we use a standard stereo computer vision dataset (Sharstein et al., 2014) to further examine the models' capacity to signal true disparities under naturalistic viewing conditions. Overall, we find that a computational model which combines both subtractive and divisive suppressive mechanisms constitutes an optimal (local) disparity detector which is also consistent with existing neurophysiological data. Interestingly, the corresponding micro-circuit constitutes a special instance of a more general class of 'sub-field normalization' models. We describe a general form for this class of models, discuss several ways they might be implemented in cortex and what specific functions they may provide.

I-79. A geometric theory of untangled image representations in sparse networks

James Golden[1,2] JAMESGOLDEN1@GMAIL.COM
Kedarnath Vilankar[2] KPV9@CORNELL.EDU
David Field[2] DJF3@CORNELL.EDU

[1]Stanford University
[2]Cornell University

The conceptual approach to understanding how the primate brain accomplishes object recognition is increasingly moving toward the notion of a neural representation that is untangled as it is processed by the ventral stream. The retinal image is put through a number of nonlinear transforms such that, in principle, a linear decision boundary in the representation space can be used to classify the retinal image as containing a particular object. This is a fundamentally geometric view, and although this theory has been supported by physiological evidence as well as machine learning image classification results, it is defined at a high level of abstraction. Following previous work at a similar level of abstraction, we report direct measurements of the geometry of neural responses in simulated networks. We demonstrate that the response manifolds of neurons in the original sparse coding network have curvature in only a small fraction of the dimensions of image space and are otherwise flat. The principal curvatures of the iso-response surfaces indicate exclusively selective responses, and their magnitudes are a function of the proximity of neighboring neurons in the image space. We find that manifolds of neurons in the two-layer variance components network exhibit both negative and positive curvature, indicative of balanced selectivity and invariance to an array of image features. The magnitude of curvature in these neurons is much greater than the neurons in the first layer. These findings provide insight into how networks begin to untangle image representations and lead to a number of questions that could be probed in physiology. We argue that nonlinear responses of visual cortical neurons in physiology already support this view, and propose several experiments that could confirm or reject this type of curvature in neural response manifolds.

I-80. Visual processing of motion-selective information in the larval zebrafish brain

Clemens Riegler[1,2] RIEGLER@FAS.HARVARD.EDU
Drago Guggiana-Nilo[1] DGUGGIAN@FAS.HARVARD.EDU
Florian Engert[1] FLORIAN@MCB.HARVARD.EDU

[1]Harvard University
[2]University of Vienna

The vertebrate retina extracts spatiotemporal features of the visual environment and diverse retinal ganglion cell (RGC) types carry extracted information into the brain. We studied how and where the direction of object motion, a behaviorally relevant feature, is represented, and further processed, in the retinorecipient arborization fields(AFs) of larval zebrafish. First, we created a functional map that describes the response properties of the RGCs that project to each AF. Of the 10 AFs, the largest is the optic tectum (homologous to the mammalian superior colliculus) and in addition there are 9 smaller AFs. Using 2-photon microscopy and GCaMP6 targeted to synaptic terminals, we recorded from all RGC terminals while showing behavior-relevant directional visual stimuli. One such visual stimulus was whole-field motion, which triggers the optomotor response. In this behavior fish turn their body and swim in the direction of perceived motion. Direction selective information is thus necessary for calculating the direction of perceived motion. Similarly, visually driven escape responses, triggered by an approaching black edge, lead to directed escape turns away from the edge. Therefore both behaviors require direction selective information. We found that the majority of OFF, ON and ON-OFF direction selective terminals are located in the optic tectum and one extratectal area termed arborization field 6 (AF6). Given this exclusivity, we asked if the same population of RGCs projects into both retinorecipient areas. Targeted laser ablation of individual AFs shows at least one subset of direction selective RGCs that both, projects to the posterior AF6, and is also

responsible for conferring directional selectivity to the posterior optic tectum. Finally, the same ablation impairs all turning behaviors to whole-field visual motion but has a lesser effect on directed escape responses. These findings establish the entry into two different behavioral circuits that require similar but not identical directional information.

I-81. Inferring hidden structure in multi-layered retinal circuits

Niru Maheswaranathan NIRUM@STANFORD.EDU
Stephen Baccus BACCUS@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Stanford University

Sensory circuits contain multiple cell layers that successively shape neural response properties. Traditionally, understanding the sensory code at a given layer involves building quantitative encoding models that directly map stimulus to response at that layer. However, between stimulus and response often lie multiple intervening layers of circuitry that are neither simultaneously recorded nor explicitly modeled. Excluding such layers limits both biophysical interpretations and computational capacity of encoding models, obscuring our ability to understand how neural circuits give rise to perception. This problem already occurs in the retina, where signals flow from photoreceptors through bipolar and amacrine cells to ganglion cells. We approximate these transformations with successive stages of linear filtering and nonlinear thresholding, thus modeling the retina as a two layer linear-nonlinear (LN) model, or LN-LN model. We asked whether the parameters of LN-LN models fit to retinal ganglion cells would learn structure resembling properties of the intervening, unrecorded retinal circuitry. To answer this, we developed novel computational methods for learning LN-LN models with very little data. In contrast to previous work, we make no assumptions about the number of hidden first-layer subunits or the structure of subunit nonlinearities. Using these new methods, we find that LN-LN models yield a better description of the retinal response, demonstrating robust improvement ($\sim 53\%$) in prediction performance over single layer LN models. Moreover, we find a striking resemblance between spatiotemporal filters learned in the model's first layer, and quantitative properties of bipolar cell receptive fields measured experimentally using intracellular recording. These results suggest that our methods for learning LN-LN models are sufficient to uncover biophysical mechanisms underlying nonlinear response properties of the sensory code. In general, our methods simultaneously infer properties of unrecorded neurons feeding into a population of recorded neurons, from which significant insights into multi-layered sensory computation may be extracted.

I-82. A new human cortical map for temporal analysis of the natural auditory speech scene

Liberty Hamilton LIBERTY.HAMILTON@UCSF.EDU
Erik Edwards ERIK.EDWARDS4@GMAIL.COM
Edward Chang EDWARD.CHANG@UCSF.EDU
University of California, San Francisco

Speech perception requires integrating information at multiple time scales. The superior temporal gyrus (STG), including Wernicke's area, is fundamental to this process, however, how responses are spatially organized remains unclear. Employing unsupervised and supervised computational methods, we reveal a new spatial map for the temporal integration of natural speech sounds across the human auditory cortex. Using high-density intracranial recordings from 20 human participants undergoing surgical treatment for intractable epilepsy, we obtained simultaneous, direct cortical recordings to speech from the entire auditory cortical hierarchy, including the higher-order STG and middle temporal gyrus, as well as the primary auditory cortex (temporal plane and Heschl's gyrus). We used convex non-negative matrix factorization to perform soft clustering on the local field potential time series

and found two spatially-segregated clusters that divided the primary auditory cortex and belt into two processing streams: a posterior one dominated by "transient" responsivity, and an anterior one dominated by "sustained" responsivity. The "transient" region exhibited fast, temporally-limited responses to sentence onsets. The "sustained" region showed responses with significantly slower peak latencies and longer temporal integration. We fit linear receptive field models to electrodes from these clusters using a spectrogram stimulus representation and a phonetic feature stimulus representation to predict the response. "Transient" electrodes were better fit by a spectrotemporal model, whereas some "sustained" electrodes were better fit by the phonetic feature model, suggesting feature extraction and invariance in the auditory cortical hierarchy. Both "transient" and "sustained" electrodes encoded consonant/vowel contrasts over short and long timescales, respectively. Compared to STG and MTG, temporal plane electrodes showed selectivity for faster temporal modulations. Finally, we segmented the neural response using syllable boundaries and found evidence for integration of prosodic information in "sustained" electrodes. These results demonstrate a new map for encoding critical temporal landmarks of spectrotemporal and linguistic cues in speech perception.

I-83. A cortical-hippocampal-cortical loop of information processing during memory consolidation

Gideon Rothschild GIDEON@PHY.UCSF.EDU
Loren Frank LOREN@PHY.UCSF.EDU

University of California, San Francisco

Humans and animals are constantly exposed to diverse sensory stimuli and numerous studies have examined encoding of these stimuli in sensory brain regions. The impact of these stimuli goes well beyond sensation, however, as while most stimuli are soon forgotten, some leave long lasting traces which can affect how we perceive and make decisions even many years later. We know relatively little about how sensory representations are transformed into long lasting memories. The dominant model of memory consolidation suggests that a recent experience is first rapidly encoded in the hippocampus, which then repeatedly reactivates this representation during sleep and thereby transfers information to the neocortex for long term storage. In particular, reactivation during hippocampal sharp wave ripples (SWRs) has been suggested as a potential mechanism for this process. To better understand information processing across these structures during consolidation, we recorded spiking activity from neuronal ensembles in the hippocampal CA1 region and auditory cortex (AC) of rats learning a place-sound association task and during interleaved sleep. We first found that single cells and neuronal ensembles in AC show spiking modulation during hippocampal SWRs in sleep. Importantly, AC reactivation onset often preceded SWRs, which contrasts with the model of SWRs as a uni-directional hippocampal-cortical communication channel. We then asked whether this preceding activity reflects information transfer in the cortical-hippocampal direction. We found that AC ensemble activity preceding SWRs could significantly predict firing of CA1 cells during the SWR, suggesting that pre-SWR cortical activity may bias hippocampal reactivation. Lastly, we recognized that if cortical activity influences hippocampal reactivation, this process could be biased by sensory stimuli. Indeed, presenting sounds during sleep biased cortical ensemble reactivation, and this reactivation predicted subsequent CA1 reactivation. Our results suggest a revised model of SWR-mediated memory consolidation during sleep, in which cortical reactivation biases hippocampal reactivation during SWRs.

I-84. Behaviour-dependent stimulus encoding and memory in primary auditory cortex

Sophie Bagur[1] BAGUR.SOPHIE@GMAIL.COM
Martin Averseng MARTIN.AVERSENG@GMAIL.COM
Shihab Shamma SAS@ISR.UMD.EDU
Yves Boubenec[2] YVES.BOUBENEC@ENS.FR
Srdjan Ostojic[2] SRDJAN.OSTOJIC@ENS.FR

[1]ESPCI, Paris
[2]Ecole Normale Superieure

Primary cortical areas are believed to process physical features of sensory stimuli, whereas higher cortical areas transform them into behavioural outcomes that can be maintained in memory until a motor command is executed. Within this framework, primary cortices are expected to efficiently represent stimulus features regardless of the behavioural meaning of these stimuli. Here we show that the type of information present in primary auditory cortex in fact strongly depends on the behavioural state of the animal. We recorded the neural activity in the primary auditory cortex (A1) of awake ferrets when the animals either passively listened or actively discriminated two periodic click trains of different rates (target vs. reference). Paradoxically, a reconstruction analysis showed that neural populations encoded the physical features of the stimuli more accurately in the passive than in the engaged state. Despite the degraded representation of stimulus features in the active state, discrimination of stimulus identity, the behaviourally relevant information, was equally accurate during sound presentation between the two behavioural states. Moreover, this information was persistently represented in the neuronal population activity in A1 during the delay period after the stimulus presentation. When animals were engaged in the task and had to behaviourally respond during the delay period, this memory trace was enhanced and was strongly correlated with the animals' performance. These findings suggest that the primary sensory cortices play a highly flexible and task-dependent role in information representation and cortical computations.

I-85. Neural circuitry underlying contrast gain control in auditory cortex

James Cooke[1] JAMES@OXFORDHEARING.COM
Benjamin Willmore[2] BEN@WILLMORE.EU
Jan Schnupp JAN.SCHNUPP@DPAG.OX.AC.UK
Andrew King[2] ANDREW.KING@DPAG.OX.AC.UK

[1]University College London
[2]Oxford University

While sensory environments can vary dramatically in their statistics, neurons have a limited dynamic range with which they can encode sensory information. In sensory cortex, this problem is in part resolved by the systematic adjustment of neural gain in accordance with the contrast of sensory input. In visual cortex shunting inhibition by parvalbumin (PV) expressing interneurons and contrast-dependent membrane potential variance have been shown to contribute to contrast gain control (CGC), but whether these mechanisms underlie CGC in auditory cortex (AC) is currently unknown. We aimed to investigate the contributions these mechanisms to CGC in mouse AC. In order to investigate the computational role of PV interneuron activity, we performed extracellular recordings of sensory evoked multi-unit responses in AC while we manipulated the activity of the PV interneurons optogenetically using Channelrhodospsin (ChR2) or Archaerhodopsin (Arch). PV interneuron activation with ChR2 did not alter spectrotemporal tuning of neuronal responses in AC but resulted in diverse effects on gain, threshold and baseline firing rate. PV interneuron suppression with Arch resulted a modest increase in temporal tuning bandwidth but left receptive field structure largely unchanged. The strongest effect of PV suppression was an increase in the gain of sensory evoked responses. Thus, PV interneuron activity does appear capable of modulating the gain of AC auditory responses. However, we found that the activity of PV interneurons did not increase systematically with increasing stimulus contrast. Finally, we performed whole cell recordings from AC neurons in

anaesthetized mice in order to assess the contribution of shunting inhibition and membrane potential variance to CGC. We found that CGC exists at the level of the membrane potential responses in AC but neither input conductance nor membrane potential variance appeared to contribute to this. This canonical computation therefore appears to be implemented by non-canonical mechanisms in different cortical areas.

I-86. Cortical adaptation is actively shaped by somatostatin-positive and not parvalbumin-positive neurons

Ryan G Natan RNATAN@MAIL.MED.UPENN.EDU
Cedric Huchuan Xia HXIA@MAIL.MED.UPENN.EDU
Winnie Rao W.RAO1993@GMAIL.COM
Maria Geffen MGEFFEN@MAIL.MED.UPENN.EDU

University of Pennsylvania

Adaptation to repeated stimuli is a ubiquitous property of cortical neurons that is thought to enhance the efficiency of sensory coding. In primary auditory cortex (A1), the vast majority of neurons exhibit adaptation, i.e. firing rate attenuation dependent upon stimulus prevalence, which modulates neuronal frequency tuning properties. Such history-dependent adaptation is thought to support sensory-motor behaviors like stimulus habituation, discrimination and deviance detection, yet little is known about the neuronal mechanisms that control how adaptation is generated. Here, we show that adaptation in A1 is differentially shaped by two inhibitory interneuron subtypes. During tone pip trains that induce varying levels of adaptation across frequencies, we measured firing rates of putative excitatory neurons (Exc), while using optogenetic manipulation to selectively suppress interneurons. Prior to adaptation, i.e. during the first tone of each train, parvalbumin-positive interneurons (PVs) uniformly inhibited responses across all frequencies. In contrast, somatostatin-positive interneurons (SOMs) enhanced responses to preferred frequencies and inhibited responses to non-preferred frequencies. Consequently, both interneurons sharpened Exc frequency preference, albeit through different mechanisms: PVs by reducing overall excitability and SOMs through stimulus-specific modulation. In addition, Exc responded to tone repetitions at different frequencies with different magnitudes of adaptation. After stimulus repetition, SOMs inhibited Exc responses to strongly-adaptive stimuli, but enhanced responses to non-adaptive stimuli. Prior to repetition, SOMs enhanced responses to adaptive stimuli and inhibited responses to non-adaptive stimuli. In contrast, PVs inhibited responses uniformly under all conditions. Thus, over the course of adaptation, SOMs reversed their effect on Exc tuning, while Pvs' effect remained unchanged. Induced by repetition, synaptic facilitation or depression at SOM–PV or SOM–Exc inputs may underlie this phenomenon. This study extends our previous discovery of SOMs' role in stimulus-specific adaptation (Natan et al., eLife, 2015), revealing a previously unknown functional mechanism of SOMs in generating adaptation.

I-87. Two subtypes of interneurons complementarily mediate behavioral detection of deviant sounds

Cedric Huchuan Xia HXIA@MAIL.MED.UPENN.EDU
Ryan G Natan RNATAN@MAIL.MED.UPENN.EDU
Winnie Rao W.RAO1993@GMAIL.COM
Maria Geffen MGEFFEN@MAIL.MED.UPENN.EDU

University of Pennsylvania

Neurons in the primary auditory cortex respond strongly to rare sounds and weakly to common sounds, a phenomenon known as stimulus-specific adaptation (SSA). However, the neuronal mechanisms underlying the relationship between SSA and behavioral detection of deviant sounds remain elusive. We recently found that two populations of cortical interneuron, parvalbumin-positive (PV) or somatostatin-positive (SOM), mediate SSA in the auditory cortex (Natan et al., eLife 2015). Here, we tested whether and how these interneurons contribute to behavioral detection of deviant sounds by combining behavioral assays and optogenetic manipulation. We assessed the ability of mice to detect standard or deviant tones by measuring the amount to which a standard or a deviant pre-pulse tone inhibits the acoustic startle reflex to a loud pulse, i.e. pre-pulse inhibition (PPI). We hypothesized that, following a train of standard tones, a deviant tone would lead to stronger PPI than another standard tone. We optogenetically suppressed PVs or SOMs to test their involvement in the detection of standard and deviant tones. Suppression of either class of interneurons led to a significantly reduced difference in PPI between standard and deviant tones, suggesting that these interneurons not only mediate SSA but also contribute to the behavioral detection of deviant sounds. Importantly, PVs and SOMs differentially contributed to reduction in PPI difference between standard and deviant tones in a stimulus-specific and complementary fashion. Suppression of PVs reduced the difference by significantly increasing PPI in response to standard, but not deviant tones. In contrast, suppression of SOMs reduced the difference by significantly reducing PPI in response to deviant, but not standard tones. Taken together, our results establish a neuronal basis to the heretofore hypothesized relationship between SSA and behavioral detection, by demonstrating that two distinct populations of interneurons differentially mediate behavioral detection of rare vs. frequent tones.

I-88. Fast gain control during naturalistic odor detection

Srinivas Gorur Shandilya COSYNE@SRINIVAS.GS
Mahmut Demir MAHMUT.DEMIR@YALE.EDU
Damon Clark DAMON.CLARK@YALE.EDU
Thierry Emonet THIERRY.EMONET@YALE.EDU

Yale University

A central question in neuroscience is how sensory neurons respond to and adapt to the statistics of natural stimuli. Airborne odorant signals tend to be highly intermittent with intensities rapidly varying over many orders of magnitude. Such complexity raises the question of how animals maintain the olfactory sensitivity while navigating odorant plumes. Addressing this question experimentally is complicated by difficulties in generating and measuring complex time-dependent odor stimuli simultaneously with neuronal output. We combined single unit electrophysiology with a novel odor delivery system capable of generating precise and repeatable stimuli to deliver naturalistic odor signals to Drosophila Olfactory Receptor Neurons (ORNs). We characterized the response and gain control properties of multiple ORNs responding to several odorants. We found that ORN gain varies inversely with odorant intensity over several orders of magnitude, consistent with the Weber-Fechner Law. Gain control was fast, with three-fold changes in gain taking place within 300ms. Strikingly, ORN gain not only changed as a function of the mean of the stimulus, but also as a function of the variance of the stimulus. For stimuli with the same mean, increased variance resulted in decreased response gain. Measurements of the Local Field Potential (LFP) and driving ORNs through light-activatable channels suggested that (1) at least some gain control takes place at the receptor level, but (2) that gain control could also be elicited by a spike-dependent mechanism. In summary, our results show that ORNs change their gain to adapt to stimulus statistics, and reveal similarities (Weber-Fechner Law, fast gain control) and differences (contrast adaptation) with other primary sensory cells. Fast gain control in ORNs is likely to be critical in maintaining sensitivity through periods of rapidly changing odor intensity, such as those experienced by an insect navigating a plume to an odor source.

I-89. Adaptive thalamic gating: A framework of dynamic encoding

Clarissa Whitmire[1,2] CLARISSA.WHITMIRE@GATECH.EDU
Christian Waiblinger[1,2] CHRISTIAN.WAIBLINGER@GATECH.EDU
Cornelius Schwarz[3] CORNELIUS.SCHWARZ@UNI-TUEBINGEN.DE
Garrett Stanley[1,2] GARRETT.STANLEY@BME.GATECH.EDU

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

It has been posited that the regulation of burst/tonic firing in the thalamus could function as a mechanism for controlling not only how much, but what kind of information is conveyed to downstream cortical targets. Yet how this gating mechanism is adaptively modulated on fast time scales by ongoing sensory inputs in rich sensory environments remains unknown. Using single unit recordings in the rat vibrissa thalamus (VPm), we found that the degree of adaptation modulated thalamic burst/tonic firing as well as the synchronization of bursting across the thalamic population along a continuum for which the extremes facilitate detection or discrimination of sensory inputs. Optogenetic control of thalamic state combined with computational modeling of single neuron dynamics further suggests that this regulation of burst/tonic firing may result from an interplay between adaptive changes in thalamic membrane potential and reduced synaptic drive from inputs to thalamus. Consistent with the view that tonic firing relays detailed information while burst firing signals the presence of a stimulus, parsing trials by burst and tonic responses demonstrated that thalamic bursting facilitated detectability while tonic activity facilitated discriminability from the perspective of an ideal observer. Generalized linear model (GLM) fits of the thalamic activity in different optogenetically manipulated thalamic states revealed clear feature selectivity associated with tonic firing, yet the thalamic bursting activity was not well captured by the standard GLM architecture due to an extreme dependence upon the silence between spiking periods, or spike history, beyond that of the standard GLM. As such, we hypothesize that a more accurate method of burst encoding models will require the tonic firing estimate of feature selectivity combined with a state variable to estimate bursting affinity. Taken together, these results suggest that dynamic burst/tonic thalamic encoding sets the stage for an intricate control strategy upon which cortical computation is built.

I-90. Neural relativity principle

Hamza Giaffar HGIAFFAR@CSHL.EDU
Daniel Kepple DKEPPLE@CSHL.EDU
Dima Rinberg[1] RINBERG@NYU.EDU
Alexei Koulakov[2] AKULA@CSHL.EDU

[1]New York University
[2]Cold Spring Harbor Laboratory

Sensory systems are constantly facing the problem of computing the stimulus identity that is invariant wrt several features. In the olfactory system, for example, odorant percepts have to retain their identity despite substantial variations in concentration, timing, and background. This computation is necessary for us to be able to navigate in chemical gradients or within variable odorant plumes. How can the olfactory system robustly represent odorant identity despite variable stimulus intensity? We propose a novel strategy for the encoding of intensity-invariant stimulus identity that is based on representing relative rather than absolute values of the stimulus features. We propose that, once stimulus features are extracted at the lowest levels of the sensory system, the stimulus identity is inferred on the basis of their relative amplitudes. Because, stimulus identity depends on relative amplitudes of features, identity becomes invariant with respect to variations in intensity and monotonous non-linearities of neuronal responses. For example, in the olfactory system, an odorant identity can be represented by the identities of p strongest responding odorant receptor types out of 1000. We show that this information is sufficient to ensure the robust recovery of a sparse stimulus (odorant) via l1 norm or elastic net loss minimization. Such a

minimization has to be performed under the constraints imposed by the relationships between stimulus features, i.e. receptor responses. This problem can be mapped onto a dual problem of minimization of a function of Lagrange coefficients. The dual problem, in turn, can be solved by a neural network whose Lyapunov function represents the dual function. We thus propose that the networks in the piriform cortex that compute odorant identity implement dual computations with the sparse activities of individual neurons representing the Lagrange coefficients corresponding to the relationships between stimulus features. Our theory yields predictions for the structure of olfactory connectivity.

I-91. Role of nonlinear dendritic integration and synaptic cooperativity in neuronal feature selectivity

Aaron Milstein[1,2] MILSTEINA@JANELIA.HHMI.ORG
Christine Grienberger[1,2] GRIENBERGERC@JANELIA.HHMI.ORG
Sandro Romani[1,2] ROMANIS@JANELIA.HHMI.ORG
Jeffrey Magee[1,2] MAGEEJ@JANELIA.HHMI.ORG

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

Experimental neuroscience is continually evolving its views of the dynamics of individual synapses, and the complexity and computational capacity of single neurons with extended dendritic morphologies. Yet, most neural network models today continue to employ neuronal units that are essentially linear integrators of synaptic weights that apply a threshold nonlinearity to produce a digital output. We seek to elucidate how the fundamental computations that underlie neuronal feature selectivity depend on implementational details at the low level of synapses and dendrites, using mouse hippocampal area CA1 as a model system. We have constructed a morphologically and biophysically detailed CA1 pyramidal cell model that explicitly implements 1) stochastic presynaptic dynamics with short-term facilitation and depression, 2) electrical compartmentalization of synaptic currents by spines, and 3) graded distributions of ion channels along dendrites. For example, increases in synaptic AMPA receptors and dendritic HCN channels with distance from the cell body counteract passive filtering along the path from spine to cell body and normalize the shape and size of unitary EPSPs measured at the cell body, as observed experimentally. However, local dendritic summation and cooperativity between inputs varies dramatically depending on the distance from the cell body, due to differential recruitment of voltage-dependent NMDA receptors. We are now using this platform to provide complex spatial and temporal patterns of excitatory and inhibitory synaptic inputs that mimic those received by CA1 place cells during spatial navigation, with simulation parameters constrained by new data from in vivo intracellular recordings. Contrary to predictions of single compartment models that linearly sum synaptic currents, our detailed model is able to reproduce an experimentally observed increase in membrane potential variance in response to an experimental reduction in inhibition, which the model suggests is due to an enhancement in supralinear dendritic integration.

I-92. Dendritic disinhibition as a mechanism for pathway-specific gating

Guangyu Yang[1] GYYANG.NEURO@GMAIL.COM
John D Murray[2] JOHN.MURRAY@YALE.EDU
Xiao-Jing Wang[1] XJWANG@NYU.EDU

[1]New York University
[2]Yale University

There is mounting evidence that dendrites of pyramidal neurons are disinhibited by an interneuronal microcircuit motif, but the functional implications remain unexplored. Here we use computational modeling to show that such a disinhibitory circuit can enable a cortical area to flexibly select one of its input pathways, according to behavioral

demands such as task rule or context. We propose that input pathways converge onto pyramidal neurons and cluster on dendritic branches, allowing control signals to open the gate for a pathway through branch-specific disinhibition. We show the plausibility of our proposal by studying a computational model of the disinhibitory interneuronal-circuit. The model is constrained by published experimental measurements that have become available only recently. Where data is not available, we made the assumptions that correspond to the "worst-case scenario", namely, connections between interneurons, and from interneurons to pyramidal dendrites are completely random, which is mostly likely not the case and any specificity would facilitate our proposed mechanism. Nevertheless, the interneuronal circuit can support the pathway-gating function. We validated this gating mechanism in a neural circuit performing a context-dependent decision-making task. We further showed that clustering of input pathways on dendrites can emerge through synaptic plasticity regulated by disinhibition. These findings suggest a microcircuit architecture that harnesses single-neuron dendritic computation to subserve cognitive flexibility. Our work challenges the implicit dogma that inhibition is non-selective, by showing that the interneuronal circuitry can actually support selective dendritic disinhibition. Our model also allows us to make specific predictions that only become testable with recent advances in cell-type specific recording. A surprising new prediction is that signaling of behavioral rule or context involves top-down projections onto specific types of interneurons.

I-93. Identifying neurons that modulate the acoustic startle response

Kathryn Tabor[1] KATHRYN.TABOR@NIH.GOV
Christopher Harris[2] CHRISTOPHER.HARRIS@NIH.GOV
Mary Brown[3] MARY.BROWN@NIH.GOV
Jennifer Strykowski[3] JENNIFER.STRYKOWSKI@NIH.GOV
Kevin Briggman[2] KEVIN.BRIGGMAN@NIH.GOV
Harold Burgess[3] HAROLDBURGESS@MAIL.NIH.GOV

[1]National Institutes of Health
[2]NINDS/NIH
[3]NICHD/NIH

A sudden noise triggers our startle reflex, but the same sound amidst a crowded street is ignored as a non-salient noise based on our constant monitoring of our surroundings. How do we synthesize sensory information to drive appropriate behaviors? Appropriate modulation of the startle response is studied using a rudimentary acoustic prepulse inhibition (PPI) assay, in which a weak prepulse presented before a strong 'startle' pulse suppresses the startle. Yet the neuronal circuit underlying PPI remains unknown. We are identifying the neurons controlling PPI in the larval zebrafish nervous system. Previously we found that neurons expressing the transcription factor genomic screen homeobox 1 (Gsx1) are necessary for appropriate modulation of the acoustic startle response. Yet 5%of cells in the zebrafish brain express Gsx1; which of these control PPI? Because larval zebrafish are naturally transparent we use 2P calcium imaging to simultaneously monitor the activity of hundreds neurons during behavioral assays. We identified Gsx1 neurons whose activity correlated with inhibition of startle during PPI assays. 3D image registration of individual brains to a single reference brain revealed that these potential PPI neurons are clustered near the otoliths in rhombomeres 2 and 3 of the hindbrain. To address whether these neurons are necessary for appropriate modulation of startle we used a combinatorial genetic approach to ablate a subset of Gsx1 neurons at the intersection of Cre and Gal4 expression patterns. Ablation of Gsx1 neurons in rhombomere 2 disrupted PPI behavior, suggesting these neurons are required for appropriate modulation of acoustic startle. We have now developed triple combinatorial genetic tools to label individual Gsx1 neurons within rhombomere 2 in order to trace their connectivity with the acoustic startle pathway. Identifying the PPI neuronal circuit will enable us to study how the nervous system discriminates irrelevant from salient stimuli at the single cell level.

I-94. Interneurons as gatekeepers of principal neuron activity: a derivation from similarity matching

Cengiz Pehlevan[1,2] CPEHLEVAN@SIMONSFOUNDATION.ORG
Dmitri Chklovskii[1] DCHKLOVSKII@SIMONSFOUNDATION.ORG

[1]Simons Foundation
[2]Simons Center for Data Analysis

What is the role of interneurons in a neural circuit? We address this question in a similarity matching framework, a recently developed approach allowing one to derive both neuronal dynamics and synaptic plasticity in a neural circuit from a principled objective function. Previously, we used similarity matching to model dimensionality reduction in early sensory processing by deriving a neural network from a Multidimensional Scaling (MDS) objective function. Such network suffers from three shortcomings: (a) Whereas the number of informative dimensions in a realistic input to be kept in the output should vary among stimuli, our network has a fixed number of output dimensions set a priori by the number of output neurons. (b) Whereas biological neurons have a limited dynamic range, our network's output power varied widely among output dimensions. (c) Whereas biological circuits comprise two types of neurons (principal and interneurons) our network has only one. Here we overcome these shortcomings by formulating two minimax similarity matching optimization problems and solving them in the offline and online settings. In the offline setting, these problems are solved by projecting the input onto its principal subspace the dimensionality of which is chosen adaptively (a) by hard-thresholding the eigenvalues of the input covariance matrix or (b) in addition, by equalizing the preserved eigenvalues. In the online setting, we derive algorithms that map onto biologically plausible neural networks whose dynamics can be viewed as a gradient descent-ascent on a convex-concave objective function. These networks contain two types of neurons (c) one making Hebbian, the other—anti-Hebbian, synapses allowing us to identify them with principal neurons and interneurons respectively. Therefore, the role of interneurons may be to clamp the power dimensions of principal neuron activity.

I-95. Reconciling neuromodulation and homeostasis

Timothy O'Leary[1] TOLEARY@BRANDEIS.EDU
Guillaume Drion[2] GDRION@ULG.AC.BE
Alessio Franci[3] AFRANCI@CIENCIAS.UNAM.MX
Eve Marder[1] MARDER@BRANDEIS.EDU

[1]Brandeis University
[2]University of Liege
[3]Universidad Nacional Autonoma de Mexico

Neuromodulators can alter the activity state of neurons and circuits drastically and for prolonged periods. For example, during locomotion, motor circuits switch from one rhythmic state to another, or from quiescence to sustained pacemaking activity. Similarly, modulatory states in the brain underlie arousal, attention, sleep and some forms of memory formation. This continual switching poses a potential problem for so-called homeostatic mechanisms, which are traditionally believed to push activity levels toward a 'set point'. Here we show, using a very simple theoretical model of activity-dependent ion channel regulation, how pronounced and long-lived changes in activity can be induced by the action of neuromodulators without interfering with ongoing homeostatic regulation. We use physiologically relevant conductance -based models of central pattern generating neurons to illustrate how these theoretical principles point to a normative framework for reconciling modulation and regulatory control in general. Furthermore, our results reconcile reliable modulation with underlying variability in the ion channels and receptors that are targeted by modulatory substances, and offer a clue as to why single modulator receptors sometimes target multiple currents in a neuron.

I-96. Highly multiplexed mapping of single neuron projections by sequencing of barcoded RNA

Justus Kebschull[1,2] JKEBSCHU@CSHL.EDU
Pedro Garcia da Silva[3] PSILVA@CSHL.EDU
Ashlan P Reid[1] REIDA@CSHL.EDU
Ian D Peikon[1] IAN.PEIKON@GMAIL.COM
Dinu F Albeanu[1] ALBEANU@CSHL.EDU
Anthony M Zador[1] ZADOR@CSHL.EDU

[1]Cold Spring Harbor Laboratory
[2]Watson School of Biological Sciences
[3]Champalimaud Centre for the Unknown

Brain regions transmit information to other brain regions via long-range axonal projections. These connections have been mapped out systematically in the mouse and other model organisms [Allen Atlas; iConnectome]. However, these maps were generated by bulk labeling techniques that obscure the fine structure of single neuron projections. Here we describe Multiplexed Analysis of Projections by Sequencing (MAPseq), a novel approach in which the long-range axonal projections of single neurons are labeled with unique random RNA sequences (RNA 'barcodes'). Barcoding converts connectivity into a form that can be probed by next generation sequencing technology, resulting in an efficient and massively parallel method for determining the projections of large ensembles of individual neurones per brain. We applied MAPseq to the Locus Coeruleus (LC), a small noradrenergic nucleus that projects throughout the brain. In contrast to previous bulk labeling studies that suggested diffuse and non-specific projections from the LC, this higher resolution analysis reveals that individual LC neurones have preferred cortical targets, suggesting differential control over cortical areas. We further show that MAPseq can be multiplexed to several injection sites in a single brain, putting a mesoscale mouse brain connectivity atlas at single cell resolution from a single brain within reach. The use of barcoding and sequencing has the potential to revolutionize the study of connectivity and function in neuroscience.

I-97. Inferring excitatory and inhibitory synaptic parameters from the local field potential

Richard Gao RIGAO@UCSD.EDU
Bradley Voytek BVOYTEK@UCSD.EDU

University of California, San Diego

Our ability to understand how neurons give rise to behaviour and cognition is in many ways limited by what signals we can record and detect. Local field potential (LFP) recordings have been invaluable for dissecting the role of neural ensembles in computation and communication, and our understanding of neural computation has benefitted immensely from recent work linking high frequency activity ($>$ 80 Hz) to population firing rate—a major neural code. In addition to firing rate, the balance of excitatory and inhibitory activity (E-I ratio) also plays a critical role in neural computation. However, outside of precise recordings of single neurons, there is still no way to estimate E-I ratio from the population activity. Here, we introduce a method to decompose the LFP into its synaptic and spiking components, by modeling it as a linear superposition of excitatory and inhibitory synaptic potentials, each with characteristic rise and decay time constants. We show that we can retrieve physiologically consistent values for excitatory and inhibitory synaptic current features, as well as the E-I ratio. We validate the model both empirically and in simulation. Given that E-I balance plays a critical role in neural computation, and that imbalances in E-I ratio are implicated in numerous neurological and psychiatric disorders such as schizophrenia and autism, having a reliable method to quickly and easily infer E-I balance has major implications for studying neural coding.

I-98. A point process approach to inferring connectivity from biophysical simulations of Ca2+ fluorescence

Saurabh Vyas[1] SMVYAS@STANFORD.EDU
Amelia Christensen[1] AMYJC@STANFORD.EDU
Catalin Mitelut[2] CATALINM@ALLENINSTITUTE.ORG
Shrivats Iyer[1] SMIYER@STANFORD.EDU
Sergey Gratiy[2] SERGEY.GRATIY@GMAIL.COM
Scott Delp[1] DELP@STANFORD.EDU
Costas Anastassiou[2] ASYMPTOTICS@GOOGLEMAIL.COM

[1]Stanford University
[2]Allen Institute for Brain Science

While technology advances in two-photon imaging have allowed monitoring of ever- increasing populations of neurons in vivo, a key limitation remains our ability to reliably estimate the connectivity of such a neural network. This problem stems from the fact that two-photon imaging does not offer direct observation of spiking associations but rather somatic calcium dynamics, i.e. a proxy for neural activity. In this report we present a Bayesian framework for estimating the neural connectivity and test it on calcium imaging data produced by biophysically detailed, large-scale simulations of thousands of functionally and morphologically detailed and interconnected neurons in a simulated patch of rodent cortex. We model the neural network as a set of coupled hidden Markov chains with chain coupling being the connectivity within the network. We use matrix factorization based methods to estimate an underlying spike train from the observed fluorescence measurements and model it using Point Processes, which appear as a hidden state within the Markov model. Finally, we estimate the model parameters using a pseudo-Expectation Maximization (EM) algorithm and find an average Dice coefficient of 0.869 +/- 0.021.

I-99. Identifying seizure mechanisms from ECoG data using directed information

Rakesh Malladi[1] RAKESH.MALLADI@RICE.EDU
Giridhar Kalamangalam[2] GIRIDHAR.P.KALAMANGALAM@UTH.TMC.EDU
Nitin Tandon[2] NITIN.TANDON@UTH.TMC.EDU
Behnaam Aazhang[1] AAZ@RICE.EDU

[1]Rice University
[2]University of Texas

Epilepsy is a common neurological disease characterized by repeated, unprovoked seizures. Epilepsy affects about 65 million people worldwide. Resective surgery and neuromodulation treatments like vagus nerve stimulation used in treating epilepsy currently don't have great efficacy. Selective spatial modulation of epileptic networks in the brain represents a possible option for better treatments. A crucial first step in this endeavor is understanding seizure generating mechanisms. We studied the seizure mechanisms via causal connectivity graphs inferred from electrocorticographic (ECoG) data during preictal, ictal and postictal periods from twelve seizures in five patients with epilepsy. Causal connectivity is quantified by directed information (DI) and inferred using two different estimators of DI. Causal connectivity inferred from the first DI estimator assumes ECoG data is derived from multivariate autoregressive process (model-based) and captures linear causal interactions between ECoG channels. Causal connectivity graph obtained from the second DI estimator doesn't impose any parametric model assumptions on ECoG data (model-free) and captures both linear and nonlinear causal interactions. We observed that model-free approach captured more causal information than model-based approach. We also intuitively expect seizure onset zone (SOZ) electrodes to act as strong sources at the beginning of a seizure and drive rest of the brain into seizure. The inference from model-free approach agreed with this intuition, while that from model-based approach did not. In fact, SOZ electrodes are weakly connected when model-based approach is used. These observations imply model-free approach is superior to model-based approach particularly while modeling

seizure mechanisms. Our preliminary analysis using model-free approach also indicated that SOZ continuously tries to drive the rest of the brain into a seizure and becomes deactivated immediately after the seizure ends. We are extending this analysis to larger patient cohorts and this analysis potentially holds the key to develop better treatments for epilepsy.

I-100. Nonlinear statistical state and parameter estimation of neural networks

Nirag Kadakia NKADAKIA@PHYSICS.UCSD.EDU
Eve Armstrong EARMSTRONG@PHYSICS.UCSD.EDU
Daniel Breen DLBREEN@PHYSICS.UCSD.EDU
Henry Abarbanel HABARBANEL@UCSD.EDU

University of California, San Diego

Most methods for parameter estimation in neural networks rely on either a) statistical properties of or b) linear approximations to the system's true dynamics. These approximations often fail to reproduce complex neural behaviors, and typically give predictions no better than phenomenological models. Here, we describe nonlinear methods that, using current computational capabilities, produce highly accurate and verifiable predictions of biophysically realistic neural systems. These techniques we employ have been traditionally developed for meteorological systems, which, like neural networks, are nonlinear, high-dimensional, and sparsely measured. We have extended them to incorporate time-delayed measurements, novel annealing protocols, and, recently, underlying symmetries; these refinements have been shown to stabilize the search for the optimal estimate and enhance the information transfer from measurements to model, reducing the minimal data needed for accurate estimation. This is particularly valuable for neurophysiologists, as we can a priori determine the number of measurements, noise, and resolution sufficient to predict the full system. Often, this sufficient data is less than anticipated, providing an extremely informative tool for experimental design. We apply these methods to the song-production pathway of the zebra finch (HVC), whose value lies in understanding the development of language and hierarchical temporal commands. We show that parameters both nonlinear and inaccessible to measurement, such as gating kinetics, can be estimated to high accuracy, without resorting to ad hoc guessing or heuristic fitting. The HVC network exhibits sparse bursts in neurons projecting to the motor command center—these patterns are known to arise from the interplay of inhibitory and excitatory neurons whose connections are as yet unclear. We propose that our methods may uncover the network topology given only sparse voltages measurements. Further, our results suggest that one can accurately estimate nonlinear synaptic dynamical parameters, responsible for features far more complex than those afforded by linear couplings alone.

I-101. Low rank minimal models for multidimensional neural computations

Joel Kaardal[1,2] JKAARDAL@PHYSICS.UCSD.EDU
Frederic Theunissen[3] THEUNISSEN@BERKELEY.EDU
Tatyana O Sharpee[1] SHARPEE@SALK.EDU

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

High-level sensory neurons are challenging to characterize with general dimensionality reduction techniques because such neurons are often selective to conjunctions of input features and require structured correlated stimuli to adequately probe their responses. Recent dimensionality reduction techniques based on maximizing noise entropy (MNE) have addressed these issues, but at the cost of estimating on the order of D-squared parameters where D is the dimensionality of the stimulus space. As a result, the estimation procedures can be prone to overfitting when the stimulus space is not well explored. Here we describe a low-rank approach to estimating

MNE models that reduces the number of variables down to 2rD where $r > 1$ describes the dimensionality of a multidimensional receptive field. We use these methods to analyze the responses of neurons in the Field L of songbird brain, a region homologous to the mammalian primary auditory cortex. These neural responses were recorded in response to natural sounds (songs). The results show highly-structured multidimensional receptive fields dominated by spectro-temporally localized components.

I-102. Assessing dynamic connectivity from high-dimensional recordings

Jordan Rodu JRODU@ANDREW.CMU.EDU
Robert Kass KASS@ANDREW.CMU.EDU

Carnegie Mellon University

When recordings from multiple electrode arrays, ECoG, EEG, or MEG are used to establish functional connectivity across two or more brain regions, multiple signals within each brain region must be combined. If, for example, signals across MEG sources in each of two regions are examined, the problem is to describe the multivariate relationship between all the signals from the first region and all the signals from the second region, as it evolves across time, during a task. While it is possible to take averages across signals, or across numerical summaries of pairwise interactions, this averaging may lose important information. We have investigated an alternative approach, which is capable of finding subtle multivariate interactions among signals that are highly non-stationary due to stimulus or behavioral effects. For a single stationary time series in each of two brain areas statistical tools such as cross-correlation and Granger causality may be applied. On the other hand, to examine multivariate interactions at a single time point, canonical correlation, which finds the linear combinations of signals that maximize the correlation, may be used. The method we have developed produces interpretations much like these standard techniques and, in addition, (a) extends the idea of canonical correlation to 3-way arrays (with dimensionality number of signals given by number of time points by number of trials), (b) allows for non-stationarity, (c) also allows for nonlinearity, (d) scales well as the number of signals increases, and (e) captures predictive relationships, as is done with Granger causality. We demonstrate the effectiveness of the method through simulation studies and illustrate with MEG data.

I-103. Effects of excitatory versus inhibitory neuron sampling on outputs of dimensionality reduction

Sean Bittner[1] SBITTNER@ANDREW.CMU.EDU
Ryan Williamson[1] WILLIAMSON@CMU.EDU
Adam C Snyder[1] ADAM@ADAMCSNYDER.COM
Ashok Litwin-Kumar[2] AK3625@COLUMBIA.EDU
Brent Doiron[3] BDOIRON@PITT.EDU
Steven Chase[1] SCHASE@CMU.EDU
Matthew Smith[3] SMITHMA@PITT.EDU
Byron Yu[1] BYRONYU@CMU.EDU

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

Dimensionality reduction methods have been used in previous studies to analyze the activity of neural populations during motor control, decision-making, visual attention, and other behavioral tasks. To date, however, dimensionality reduction has been applied to neural populations without consideration of cell type. Since excitatory and inhibitory neurons have different synaptic transmission and connectivity properties, one might expect the types of neurons sampled to affect the measured dimensionality and the extent to which activity fluctuations are shared across the population. In this study, we assess the influence of cell type on dimensionality reduction measurements by applying factor analysis to excitatory and inhibitory neurons in balanced network models and real data. Specifically, we assess how cell type influences the measured dimensionality (the number of latent dimensions needed to describe the population activity), and the percent shared variance (the percentage of a neuron's spike count variance that is correlated with other neurons in the sampled population). We performed these analyses for two types of balanced networks: one in which the excitatory neurons had clustered connections and one with no clustered connections. We found intriguing differences in the outputs of dimensionality reduction depending on the ratio of excitatory and inhibitory neurons sampled and the underlying network structure. Furthermore, we applied the same analyses to population activity recorded in the primary visual cortex, where we labeled the neurons as excitatory or inhibitory using waveform classification. Overall, these results indicate that the two cell types have unique contributions to the dimensionality of network activity.

I-104. Variational inference for nonlinear tuning curves using dimension reduced mixtures of GLMs

Vivek Subramanian VAS11@DUKE.EDU
Jeff Beck JEFF.BECK@DUKE.EDU
Duke University

Brain-machine interfaces (BMIs) are devices that transform neural activity into commands executed by a robotic actuator. For invasive BMIs, neural decoding is performed on spikes measured from motor cortical neurons. Decoding algorithms rely on tuning functions for each neuron, which must accurately characterize the response of the neuron to an external stimulus or upcoming motor action. Tuning curves for all neurons may not take the same parametric form and may be sensitive to aspects of the stimulus that are not intuited by researchers. As a result, performance of BMI decoding may suffer because of a simplified or incorrect mapping from firing rate to kinematic variables. We develop a non-parametric Bayesian model for the identification of non-linear tuning curves with arbitrary shape. The model functions by projecting some high-dimensional stimuli vector down to a lower dimension and then approximates non-linear tuning curves by a piecewise linear function. This function is constructed by clustering the low-dimensional projection of the stimulus and then fitting the observed response with a GLM unique to that cluster. We exploit the Dirichlet process to allow the number of clusters to be learned from the data while simultaneously favoring as few components as possible to maximize the marginal likelihood. Inference is accomplished via the Variational Bayesian (VB) expectation maximization algorithm, which results in fast, (mostly) explicit update rules for the parameters and is easily applied to large data sets. The algorithm also provides a tight lower bound on the marginal likelihood, allowing for model comparisons. We demonstrate our model's capabilities on both simulated and experimental data sets. In simulated data, the model is able to accurately fit the chosen non-linearities. In experimental data, the model discovers non-linear responses to kinematic variables which could not be discovered by current methods used in the BMI literature.

I-105. A robust Bayesian method for decomposing noisy single-trial voltage-clamp traces for circuit mapping

Ben Shababo[1] SHABABO@BERKELEY.EDU
Josh Merel[2] JSMEREL@GMAIL.COM
Alex Naka[1] ALEX.NAKA@BERKELEY.EDU
Hillel Adesnik[1] HADESNIK@BERKELEY.EDU
Liam Paninski[2] LIAM@STAT.COLUMBIA.EDU

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

The combination of optical physiology and whole-cell electrophysiology has proven useful in mapping the structure of neural circuits. In particular, we consider an approach where monosynaptic connectivity is measured by recording postsynaptic neurons in voltage-clamp while optically stimulating and/or observing presynaptic neurons. While these experiments have provided some resolution of neural connectivity, the high-throughput, dense, single-cell resolution mapping of monosynaptic connections remains a challenge. In this work, we address a core data analysis problem in scaling up these experiments. While synaptic mapping that makes use of subthreshold events such as postsynaptic currents (PSCs) has access to information that would be lost if only spiking activity were observed, most PSCs are small compared to noise levels. In addition, the timing of optically evoked spiking can be variable (relative to PSC kinetics) because of the magnitude and speed of opsins or caged-neurotransmitters. Therefore, if we wish to map synaptic connections with optically evoked (or spontaneous) presynaptic activity, performance is dependent on the robust decomposition of single-trial, voltage-clamp recordings into unitary PSCs. We present a Bayesian approach for inferring the timing, strength, and kinetics of postsynaptic currents (PSCs). Importantly, the intuitive, generative single-trace model and straight-forward inference procedure flexibly extend to include presynaptic data (stimulation or recordings), structure for synaptic connectivity, and other experimental details. We demonstrate on simulated and real data that this method performs better than standard methods for detecting spontaneous PSCs. We also demonstrate PSC detection on several types of data relevant to mapping experiments such as voltage-clamp recordings contaminated with optically evoked currents. Finally, we provide an example of an extension of the single-trace model and inference procedure to a mapping scenario which includes calcium imaging observations for a set of neurons - some of which are a subset of the presynaptic sources of PSCs in the voltage-clamp recording.

II-1. Bayesian maggots: near-optimal probabilistic inference in Drosophila larvae

Matthieu Louis[1,2] MATTHIEU.LOUIS@ICLOUD.COM
Andreas Braun[1] ANDREAS.BRAUN@CRG.EU
Ruben Moreno[3] RUBEN.MORENO@UPF.EDU
Daniel Robles Llana[4] DANIEL.ROBLESLLANA@GMAIL.COM
Alexandre Pouget[4] ALEX.POUGET@GMAIL.COM

[1]Centre for Genomic Regulation
[2]EMBL-CRG Systems Biology Unit
[3]University Pompeu Fabra
[4]University of Geneva

Numerous studies have shown that a wide range of behaviors from sensory processing to motor control involve near-optimal probabilistic inference. Most of these studies have focused on vertebrates, suggesting that the ability to perform probabilistic inference requires large nervous systems. Yet, neural theories of probabilistic inference can be implemented with the most basic neural networks. To explore the possibility that organisms with small nervous systems perform near-optimal probabilistic inference, we exploited the ability of Drosophila larvae to achieve robust navigation in sensory gradients. Larvae were tested in two-dimensional arenas comprising two olfactory stimuli: a real- odor gradient and a virtual-odor gradient created through the optogenetic stimulation of a single olfactory sensory neuron insensitive to the real odor. We manipulated the reliability of the virtual-odor cue by superimposing random light flashes to a static light gradient. We considered two models: an optimal Bayesian model in which cues are weighted according to their respective reliabilities, and a suboptimal model in which cues are assigned fixed weights. For both models, we derived behavioral predictions for the real and virtual odor alone with noise, as well as for their combination with and without noise. We report that the fixed-weight model fails to accurately predict performance to single cues with and without noise. In contrast, the Bayesian model provides tight fits to the single-cue conditions and, crucially, predicts the behavior well when both cues are combined. We applied the same paradigm for the combination of an odor and a light gradient, and found similar results. Our findings demonstrate that near-optimal inference is not restricted to integration within a single modality, but that it also applies to multisensory integration. This work sets the stage for a systematic analysis of the neural computations underlying probabilistic inference in an insect brain amenable to genetic manipulations and physiological inspections.

II-2. A universal tradeoff between energy, speed and precision in neural communication

Surya Ganguli SGANGULI@STANFORD.EDU
Jascha Sohl-Dickstein JASCHA.SOHLDICKSTEIN@GMAIL.COM
Subhaneil Lahiri SULAHIRI@STANFORD.EDU
Stanford University

The brain performs myriad computations rapidly and precisely while only consuming energy at a rate of 20 watts, as opposed to typical supercomputers whose consumption rate is in the megawatt range. This remarkable performance suggests that evolution may have simultaneously optimized energy, speed and accuracy in neural computation. What are the fundamental limits and tradeoffs involved in simultaneously optimizing these three quantities, and how close do neural systems come to these limits? Information theory and thermodynamics provide upper limits on the accuracy and energy efficiency of any physical device, including biological systems. However, these bounds ignore the time it takes to achieve such limits and are therefore not obviously relevant to neurobiology, as neural systems cannot wait an infinite time to receive and transmit information. We go beyond classical information theory and thermodynamics to prove a new theorem revealing that the communication of an external signal via any signaling system is subject to a universal tradeoff between (1) the rate of energy consumption by the system, (2) the speed at which the signal varies, and (3) the precision with which the signal is communicated. This applies to arbitrary systems, for example synaptic spines transducing incoming neurotransmitters, retinal opsins transducing light, or cochlear hair cells transducing sound. We find simply that the product of speed and precision is upper bounded by energy consumption. Intuitively, this three-way tradeoff arises because any increase in speed at fixed precision requires the signaling system's physical states to change rapidly, leading to increased energy consumption. Similarly any increase in precision at fixed speed requires a larger change in the probability distribution over signaling system's states, which leads to increased energy consumption. Our newly discovered three-way tradeoff motivates new experiments to assess exactly how close neural systems come to simultaneously optimizing energy, speed and accuracy.

II-3. Speech-trained neural networks behave like human listeners and reveal a hierarchy in auditory cortex

Alexander Kell ALEXKELL@MIT.EDU
Daniel Yamins YAMINS@MIT.EDU
Samuel Norman-Haignere SVNH@MIT.EDU
Josh McDermott JHM@MIT.EDU
Massachusetts Institute of Technology

Auditory cortex remains poorly characterized, particularly in computational terms. Motivated by recent work in the ventral visual stream, we developed an improved model of auditory cortical computation by searching a vast space of biologically-inspired models for one that performs well on a real-world auditory task. We trained a hierarchical convolutional neural network (HCNN) to map speech excerpts, embedded in complex background noise, to word labels. To evaluate the network's plausibility as a model of human hearing, we ran both the HCNN and human listeners in a psychophysical experiment measuring word recognition as a function of the type of background noise and SNR. The HCNN performed as well as humans, and, despite not being trained to fit human behavior, exhibited a similar pattern of performance across background types and SNRs (r2 = 0.94). We then tested the HCNN as a model of auditory cortex. We measured human auditory cortical responses to 165 natural sounds with fMRI and presented those same sounds to the HCNN. We predicted each voxel's responses from each layer's unit activities using cross-validated regularized linear regression. For comparison, we performed the identical procedure with a standard spectrotemporal filter model of auditory cortical computation (Chi et al., 2005). The task-optimized HCNN out-predicted the standard model, particularly in non-primary auditory cortex. Moreover, different portions of cortex were best predicted by different HCNN layers. Primary auditory cortex was best predicted by earlier layers; nearby non-primary areas were best predicted by intermediate layers; and more distant areas were best predicted by later layers. These results suggest a multi-staged hierarchy of cortical computation in human auditory cortex, and indicate that task-optimized models offer a powerful means of understanding sensory systems.

II-4. Models of human fixation search in natural scenes

Jared Abrams JARED.ABRAMS@GMAIL.COM
Wilson Geisler W.GEISLER@UTEXAS.EDU

University of Texas

Visual search is a fundamental natural task. Normative models of search provide a critical benchmark for understanding human behavior in such tasks. We consider six models that vary in complexity, but perform at least as well as human observers in our search tasks. Each model consists of two stages: updating beliefs about target location based upon visual responses, and deciding where to fixate next. In the first stage, models use either (1) global contrast normalization or (2) local contrast normalization when updating beliefs. In the second stage, the models select fixations by either (1) selecting a random location with inhibition of return (RAND), (2) selecting the location with features most similar to the target (Maximum a Posteriori, MAP), or (3) selecting the location maximizing information gained about the location of the target (Entropy Limit Minimization, ELM). The RAND searcher utilizes knowledge of the display region, and where it has already fixated. The MAP searcher selects the location with the highest posterior probability of containing the target. Finally, the ELM searcher blurs the posterior probability map with the square of the detectability map in order to create a map of the expected information gain, and picks the max. To test the models, we ran a double dissociation experiment where we selected image-target location pairs that allowed us to assess the role of contrast normalization in human fixation selection. We found that human fixations are consistent with local contrast normalization, ruling out the simpler models that only use global contrast normalization. Next, we analyzed the human fixations in relation to the predicted model fixations. From those data, we rule out the RAND model, and tentatively accept the MAP model, with some caveats. Thus, humans make rational decisions about fixations and make use of local contrast normalization in order to facilitate search performance.

II-5. Dynamic reorganization of neuronal activity patterns in parietal cortex

Laura N Driscoll LNDRISCOLL@GMAIL.COM
Christopher D Harvey CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

Harvard Medical School

The ability to learn new associations between information in our environment and our actions is vital. We are interested in the balance between stability of learned behaviors and flexibility required for learning. We developed a navigation-based decision task in virtual reality and combined this task with chronic recording of neuronal activity using two-photon calcium imaging. We used these methods to study neural circuit dynamics during learning and during stable behavioral performance post-learning, over the time course of $>$ 4 weeks. We focused on posterior parietal cortex (PPC) because it is a key player in linking abstract sensory representations to behavioral actions, which likely develop in large part through learning. We recently found that although single cells reliably and selectively responded to task-relevant parameters on a single session, such as location in a T-maze and behavioral choice, the activity patterns of individual cells changed from day to day. The PPC on each day formed a different trial-type-specific trajectory of population activity, from which choice and other task parameters could be decoded. During a period of 20 consecutive days, less than 10% of task-modulated neurons had significantly similar activity patterns. The changes in activity could not be explained by variability in behavioral performance. These changes were larger than changes observed in V1 for orientation selectivity. We also tracked changes during learning of a new cue-response association, following expert performance on an initial set of two associations. Trials with the novel cue could be decoded from population activity within few trials of the first presentation. As behavioral performance improved, neuronal activity during the new association became more separable from neuronal activity during previously learned associations. Together these findings suggest that PPC has flexible representations of task-relevant information, which could be useful for balancing the stability of learned behaviors and flexibility required for learning.

II-6. Circuit principles underlying a changing motivational state

Stephen C Thornquist[1,2] THORNQUIST@FAS.HARVARD.EDU
Michael Crickmore[1] MICHAEL.CRICKMORE@CHILDRENS.HARVARD.EDU

[1]Boston Children's Hospital
[2]Harvard Medical School

An animal's actions are not determined exclusively by its surroundings; instead, it tunes its behavior to satisfy its most urgent physiological needs. These physiological needs change with time, and accordingly the tendency to perform distinct behaviors changes too. This computation is facilitated by drives, an abstract representation of the importance of separate behaviors. We present findings demonstrating that the mating behavior of flies is a particularly tractable system for understanding how drives are generated and compared within the brain on a circuit level. In particular, we have found easily generalized circuit principles that control a gradually changing motivational state and how it is influenced by multiple drives. The most noteworthy of these are 1) a peptidergic switch serving as an internal representation of the animal's physiology, and which induces an abrupt transition in the motivational state of the animal, and 2) dual-use circuitry controlling both the slow decline in drive during consummatory behavior and the suppression of drive in the presence of a strong external stimulus. These principles are fundamental to the computations performed by this neural circuit but are underrepresented in computational models of collections of neurons, suggesting an additional and essential element to account for the state dependence of behavior.

II-7. Temporal processing of prediction errors reflects hierarchical representation of beliefs

Christoph Mathys[1] C.MATHYS@UCL.AC.UK
Andreea Diaconescu[2,3] ANDREEAOLIVIANA.DIACONESCU@UZH.CH
Vladimir Litvak[1] LITVAK.VLADIMIR@GMAIL.COM
Karl Friston[1] K.FRISTON@UCL.AC.UK
Sara Tomiello[2,3] SARA.TOMIELLO@BIOMED.EE.ETHZ.CH
Katharina Wellstein[2,3] WELLSTEIN@BIOMED.EE.ETHZ.CH
Gina Paolini[2,3] PAOLINI@INI.PHYS.ETHZ.CH
Klaas Enno Stephan[2,3] STEPHAN@BIOMED.EE.ETHZ.CH

[1]University College London
[2]University of Zurich
[3]ETH Zurich

Electroencephalograms (EEGs) were recorded from 100 healthy male human subjects while they performed a social decision-making task that required them to integrate information and update beliefs. Model comparison revealed that, among the models tested, subjects' behavior was best explained by a variant of the Hierarchical Gaussian Filter (HGF) (Mathys et al., 2011; 2014). According to the HGF, subjects minimize surprise by updating beliefs using uncertainty-weighted prediction errors about a set of hierarchically organized quantities relevant to the task. In the winning model, these were the observed cues, the advice received, the advice's tendency to be correct, the volatility of this tendency, and the uncertainties about the latter two. Simpler variants of the HGF

and Rescorla-Wagner learning had less evidence in their favor. Because subjects' predictions and uncertainty estimates were based on estimates of individual prior beliefs about the environment, each had his/her own set of estimated belief trajectories. The time points from these estimated trajectories were used as the independent variables in a linear model of single-trial EEG responses covering epochs of -500 ms to +550 ms relative to outcome onset. This revealed the temporal structure of modulations of evoked potentials in the space of the 128 EEG channels. For each of the prediction errors and uncertainty weights in the winning model, there was an EEG signature in the form of a spatially localized significant positivity or negativity. Strikingly, the temporal succession of these signatures corresponded to their place in the HGF's hierarchy. Signatures for lower levels were followed by those for higher levels. This is further evidence that the brain makes predictions based on hierarchical belief structures that are updated using uncertainty-weighted prediction errors. By elucidating the temporal structure of this process, this study complements fMRI studies of its spatial aspects (Iglesias et al., 2013; Vossel et al., 2015).

II-8. Unsupervised quantifications of social interactions in fruit flies

Ugne Klibaite[1] KLIBAITE@PRINCETON.EDU
Gordon J Berman[2] GORDON.BERMAN@EMORY.EDU
Qingqing Wang QINGQING_WANG@BERKELEY.EDU
Jessica Cande[3,4] CANDEJ@JANELIA.HHMI.ORG
David Stern[3,4] STERND@JANELIA.HHMI.ORG
Donald Rio[5] DON_RIO@BERKELEY.EDU
Joshua Shaevitz[1] SHAEVITZ@PRINCETON.EDU

[1]Princeton University
[2]Emory University
[3]Janelia Farm Research Campus
[4]Howard Hughes Medical Institute
[5]University of California, Berkeley

Social interactions are complex and are frequently crucial for an animal's survival. These interactions, ranging across many sensory modalities, length scales, and time scales, are often subtle and difficult to quantify. This has limited our ability to dissect the genetic and neurobiological mechanisms underlying social behaviors to all but the most dramatic effects. Thus, we require new analysis methods in order to quantify the more delicate behavioral effects of neural perturbations in socializing animals. Recent advances have made it possible to systematically quantify the behavior of freely moving individuals, measuring an animal's entire behavioral repertoire of stereotyped motions in an unsupervised and high-throughput manner [1]. We extend this framework to multi-individual interactions by tracking, segmenting, and analyzing the simultaneous behavior of multiple interacting fruit flies, separately processing the postural dynamics for each. By quantifying probabilities of displaying certain behaviors, behavioral transitions, and interactions over the course of an experiment, we can investigate the social phenotypes that arise within these contexts, and capture small but significant deviations in behavioral repertoire. This method is readily applicable to studies of behavioral evolution, high-throughput screens of social behaviors, and studies targeting behavioral phenotypes. We apply these methods to fruit flies of the Melanogaster subgroup during courtship. Comparisons of behavior between lone and interacting animals reveal previously uncharacterized behaviors displayed during courtship. Finding temporal and spatial relationships between male and female, we determine which features of courtship lead to greater chances of successful copulation. By measuring behavior of transgenic flies that contain alternative splicing in neuronally expressed genes, we determine subtle behavioral shifts that are caused by these genetic, and ultimately neural, changes.

II-9. Laminar predictors of attentive behavior in visual area V4

Monika Jadi JADI@SALK.EDU
Anirvan S Nandy NANDY@SALK.EDU
Terrence Sejnowski TERRY@SALK.EDU
Saket Navlakha NAVLAKHA@SALK.EDU
John Reynolds REYNOLDS@SALK.EDU

Salk Institute for Biological Studies

Attention is a critical mechanism for overcoming information bottleneck in perception. It is thought to be mediated by feedback signals from attentional control centers that target the laminar circuits of sensory cortices as early as visual area V4. However, due to technical challenges in area V4, the laminar neural correlates of attentive behavior have remained unknown. We have overcome this via a novel means of an optically clear cranial window to advance laminar electrodes down a cortical column (Fig.1A). This allowed us to reliably assign recording channels into supra-granular (SG), granular (G) and infra-granular (IG) layers (Fig.1B). We monitored neural activity as macaque monkeys performed an attention-demanding orientation-change detection task (Fig.2A). We analyzed the sub-set of behavioral conditions (Fig.2B) in which the animals were equally likely to correctly detect ('hit') or fail to detect ('miss') a target stimulus presented at a spatially cued location either inside (IN) or outside (AWAY) the receptive field overlap region of the V4 column. Using machine-learning techniques, we assessed the laminar and spatial profile of task-predictive information in the local field potentials (LFP) in a 200 ms window prior to stimulus onset. The 'hit' vs. 'miss' prediction performance of LFP classifiers for all laminae was overall high (>85%); IG classifiers had the best performance (>95%). Performance of classifiers trained on different frequency bands of LFPs showed a peak in the 15-30 Hz range for all laminae (Fig.3). Classifiers trained with spatially specific (IN, AWAY) LFPs showed similar specificity of laminae as well as LFP frequency bands. Further, we found that broad spiking units in SG exhibit a marked decrease in variability within a 100 ms window prior to stimulus onset. Taken together, our first report of laminar recordings from a V4 column suggests a lamina-specific signature of neural correlates of performance on attention-demanding tasks.

II-10. Hippocampal ensemble dynamics timestamp events in long-term memory

Alon Rubin ALON.RUBIN@GMAIL.COM
Nitzan Geva NITZGEVA@GMAIL.COM
Liron Sheintuch LIRONSHEINTUCH@GMAIL.COM
Yaniv Ziv YANIV.ZIV@WEIZMANN.AC.IL

Weizmann Institute of Science

The capacity to remember temporal relationships between different events is essential to episodic memory, but little is currently known about its underlying neural mechanisms. We performed time-lapse imaging of thousands of neurons over weeks in the hippocampal CA1 of mice as they repeatedly visited two distinct environments. Longitudinal analysis exposed ongoing environment-independent evolution of episodic representations. These dynamics time-stamped experienced events via neuronal ensembles that had cellular composition and activity patterns unique to specific points in time. Temporally close episodes shared a common timestamp regardless of the spatial context in which they occurred. Temporally remote episodes had distinct timestamps, even if they occurred within the same spatial context. Our results suggest that days-scale hippocampal ensemble dynamics could support the formation of a mental timeline in which experienced events could be mnemonically associated or dissociated based on their temporal distance.

II-11. Memory transformation enhances reinforcement learning in dynamic environments

Blake Richards[1] BLAKE.RICHARDS@UTORONTO.CA
Adam Santoro[2] ADAMSANTORO@GMAIL.COM
Paul Frankland[3] PAUL.FRANKLAND@GMAIL.COM

[1]University of Toronto Scarborough
[2]Google DeepMind
[3]The Hospital for Sick Children

To make decisions, the brain uses both episodic memories, which provide detailed accounts of single events, and semantic memories, which provide abstract concepts based on generalizations. Interestingly, studies have found that there is a switch from detailed episodic memory to generalized semantic memory over the course of systems consolidation. This switch is sometimes referred to as 'memory transformation'. It is not known why the brain engages in memory transformation, though it has been postulated that it helps to prevent interference. Here we demonstrate a previously unappreciated benefit of memory transformation, namely, its ability to enhance reinforcement learning in a changing environment. We developed a neural network that is able to store memories both for specifics (episodic memories) and statistical patterns (semantic memories). The network is trained to find rewards in a foraging task where reward locations are continuously changing. We find that over short time frames, episodic memories support efficient foraging. However, over extended time frames, semantic memories are better at predicting reward locations, because the accumulation of change over time renders precise memories obsolete. Furthermore, across multiple time frames a system that engages in memory transformation outperforms systems that use solely episodic or semantic memory. We also show that our network can recapitulate experimental data from mice in a modified water-maze task, wherein the location of the platform changes each day. These results suggest that the brain's strategy of relying on precise memories soon after an experience, and generalized memories later on, aids learning in changing environments. Our work recasts the theoretical question of why memory transformation occurs, shifting the focus from memory interference towards the enhancement of reinforcement learning in dynamic environments.

II-12. Optimal storage capacity associative memories exhibit retrieval-induced forgetting

Andrew M Saxe[1] ASAXE@FAS.HARVARD.EDU
Kenneth Norman[2] KNORMAN@PRINCETON.EDU

[1]Harvard University
[2]Princeton University

Retrieving a memory can, surprisingly, cause forgetting of related competitor memories, a phenomenon known as retrieval-induced forgetting. For example, after studying a list of category-exemplar pairs ("Fruit-Pear," "Fruit-Apple",...), partial practice of one target pair ("Fruit-Pe") can cause forgetting of related competitor pairs ("Fruit-Apple"). This striking behavioral finding has been probed in a wealth of experiments that have delimited four key features of this effect: partial practice yields retrieval-induced forgetting; extra study of the complete item ("Fruit-Pear") yields no RIF despite equivalent target strengthening; reversed practice with incomplete category information ("F-Pear") yields no RIF; and when present, the RIF effect can be elicited using *independent cues* ("Red-A") rather than the specific cues used during learning (Norman et al., 2007; Anderson, 2003). These robust and intricate empirical findings severely constrain models of memory and pose a crucial challenge for theory: what sort of memory system might yield these effects, and why? Here we develop a quantitative theory of retrieval-induced forgetting by deriving new exact solutions to the dynamics of learning for the generalized perceptron learning rule (GPLR) as it embeds memories in a binary recurrent neural network. These solutions reveal the full trajectory of every recurrent weight over the course of learning, and yield closed-form expressions for the amount of RIF seen following partial practice, extra study, and reversed practice. In accord with behavioral findings,

partial practice yields RIF while reversed practice and extra study provably do not. Moreover, we show that the recurrence in the network causes independent cues to exhibit RIF to the same extent as trained cues. Hence the GPLR, which is known to attain optimal storage capacity in recurrent binary networks (Gardner, 1988), naturally exhibits subtle behavioral phenomena linked to retrieval induced forgetting, suggesting that RIF is a hallmark of memory storage using a computationally optimal learning rule.

II-13. Neuronal representation of reward in parietal cortex

Jan Kubanek[1] KUBANEK@STANFORD.EDU
Lawrence Snyder[2] LARRY@EYE-HAND.WUSTL.EDU

[1]Stanford University
[2]Washington University

The burgeoning field of neuroeconomics was born out of the finding (Platt and Glimcher, 1999) that parietal neurons encode the economic value associated with each choice alternative. This finding has been made in tasks in which the value associated with each alternative was relatively predictable. However, it has been unclear how parietal cortex responds in more general foraging settings in which the option values are not easy to predict. We engaged animals in such a foraging task. In blocks of trials, one option offered a larger mean reward than the other option, and the size of the reward obtained in each trial was drawn from an exponential distribution with the particular mean. In this dynamic foraging task, macaque monkeys exhibited Herrnstein's matching behavior: they almost exactly matched the proportion of their choices to the proportion of the mean rewards associated with each option. We found that the size of the obtained reward governed the animals' decision whether to repeat their previous choice or switch. We further found that parietal neurons strongly encoded the reward size: the larger the obtained reward, the less vigorously parietal neurons fired action potentials. This effect was independent of–and of sign opposite to–the effect of value reported in parietal cortex previously. This novel parietal effect shares the same basic properties with signals previously reported in the limbic system that detect the size of the recently obtained reward to mediate proper repeat-switch decisions. The decision between repeating and switching is simpler than a decision based on a computation of value of the choice options, and might be more generally applicable to foraging situations with unpredictable rewards. Parietal neurons strongly represent this simpler decision strategy.

II-14. Using feedback to dissect the hierarchy at which the prior distribution affects perception

Yonatan Loewenstein YONATAN@HUJI.AC.IL
Merav Ahissar MSMERAVA@GMAIL.COM
Ofri Raviv OFRI.RAVIV@GMAIL.COM

The Hebrew University

Perception is a multistage process, in which stimuli are processed in a hierarchy of increasing complexity, from basic physical features to task-specific responses. It is strongly modulated by the statistical distribution of past stimuli. For example, we have previously shown that in pitch perception, acuity is dominated by a tendency of the percept to gravitate towards the center of a rapidly formed estimation of the prior distribution, a phenomenon known as contraction bias. Our goal is to identify at which level of the perception-action hierarchy the prior distribution affects perception, by considering two competing hypotheses: (1) contraction bias emerges from unsupervised learning processes that manifest as Bayesian-like computation; (2) the contribution of the prior distribution to perception is merely a particular manifestation of the operation of supervised learning algorithms. Using numerical simulations, we show that both hypotheses are consistent with current experimental findings. The reason for the similar predictions is that reinforcement learning algorithms are expected to yield Bayesian-like computation even if the prior is not explicitly represented, simply because that is the behavior that optimizes performance. Therefore, to test these competing hypotheses, we designed a series of studies, in which binary feedback was systematically manipulated in a subset of 'impossible' trials. We found that as predicted by the 'unsupervised' hypothesis and in contrast to the 'supervised' hypothesis, feedback designed to enhance or suppress contraction bias has no effect on its magnitude. By contrast, a similar feedback designed to bias participants in favor of a particular response effectively shifts the psychophysical curve. These results indicate that the contribution of prior distribution to perception is due to unsupervised learning processes that are insensitive to supervision. They support a hierarchical view of perception, where high-level task-related supervision has limited access to low level unsupervised processes of prior formation and integration.

II-15. Anterior cingulate cortex-hippocampal interactions during goal-driven behaviors

Jai Yu[1] JAI@PHY.UCSF.EDU
Adrianna Loback[2] ADRIANNA@PRINCETON.EDU
Irene Grossrubatcher[3] IRENE.GROSSRUBATSCHER@GMAIL.COM
Daniel Liu[1] DANIEL.LIU@UCSF.EDU
Loren Frank[1] LOREN@PHY.UCSF.EDU

[1]University of California, San Francisco
[2]Princeton University
[3]University of California, Berkeley

The brain has the remarkable ability to adapt behavioral choices when faced with changes in the world. This process requires interactions among multiple brain areas, including the hippocampus and the prefrontal cortex (PFC). While neural correlates of planning and decision making have been identified separately in the PFC and the hippocampus, it is unclear how these regions interact to facilitate these processes. We examined these interactions in the rat using a novel spatial decision making task with changing spatial reward contingencies designed to engage the hippocampus and the anterior cingulate cortex (ACC), a part of the PFC critical for adaptive decision making. We found inactivation of ACC was sufficient to disrupt trajectory choices without altering motivation. We then recorded from multiple single neurons in both the ACC and the hippocampus of awake, behaving rats to examine the interactions between these structures. We found that network ensemble patterns in the ACC preceding trajectory decisions are dynamic and reflect both the operative reward contingency and location. We also found strong modulation during these periods in the activity of a large population of ACC neurons at the time of hippocampal sharp wave-ripples (SWR) events. These SWRs tend to occur during periods of immobility following reward delivery and can reactivate representations of past or potential future trajectories. Interestingly, ACC neurons that were more active during running between goals momentarily increased their activity coincident with SWRs, while others decreased their activity. These results indicate that representations associated with the execution of a trajectory become briefly reactivated during SWRs while other representations are simultaneously suppressed. Our results suggest spatial representations about trajectory options from the hippocampus are concurrently reinstated with associated task relevant cortical representations, potentially facilitating learning and decision making.

II-16. A learning rule for optimal evidence accumulation in decision-making

Jan Drugowitsch JDRUGO@GMAIL.COM
Daniel Robles Llana DANIEL.ROBLESLLANA@GMAIL.COM
Alexandre Pouget ALEX.POUGET@GMAIL.COM

University of Geneva

To make efficient decisions in noisy and ambiguous circumstances, nervous systems need to not only accumulate evidence—however briefly—but, crucially, also learn how to interpret neural population activity that encodes this evidence. Such learning is important to acquire competence in a task (e.g. Law & Gold, 2008), as well as to track changes in the environment and the neural encoding of evidence (e.g. Mendonca et al., 2015). However, besides various heuristic proposals, little is known about how such learning might be implemented. We address this issue in a Bayesian framework in which neural activity of the evidence-encoding population is accumulated and combined linearly until one of two choice-triggering boundaries is reached. The aim is to learn the linear combination weights from feedback about the correctness of these choices, and to do so for arbitrary large sensory populations. Done optimally, the learning rule incorporates feedback with a strength that is modulated by three factors: the confidence in the choice, its correctness, and the current uncertainty about the learned weights. By comparing this learning rule to alternative heuristics we show that all three factors are important for stable and efficient learning. In particular, the vanilla delta rule never reaches stable performance unless augmented by weight re-normalization, and even then is for all parameter choices significantly out-performed by the optimal rule. Applied to changing environments or varying evidence encoding, the optimal learning rule accounts for the frequently observed higher chance of repeating the same choice after correct feedback—the so-called win-stay lose-switch strategy. The theory predicts that this choice stickiness ought to be stronger for hard than for easy choices, as has also been previously observed. Thus, our theory not only clarifies the factors that should influence learning of evidence accumulation, but also provides normative accounts for previously observed sequential choice behavior.

II-17. Choice probabilities in the presence of nonzero signal stimuli, internal bias, and decision feedback

Daniel Chicharro[1] DANIEL.CHICHARRO@IIT.IT
Stefano Panzeri[1] STEFANO.PANZERI@IIT.IT
Ralf M Haefner[2] RALF.HAEFNER@GMAIL.COM

[1]Istituto Italiano di Tecnologia
[2]University of Rochester

Choice probabilities (CP) quantify the correlation between a neuron's response variability and the animal's choice in two-choice discrimination tasks (Nienborg et al. 2012). Here, we extend previous work (Haefner et al. 2013), which for a feedforward model links CPs to read-out weights and noise correlations. We include decision feedback (Nienborg & Cumming 2009), internal biases (e.g. serial dependencies, Fruend et al. 2014), and nonzero signal stimuli. We reassess the role of noise correlations and study how CPs are determined by feedback weights, by whether feedback modulate the sensory representation additively or multiplicatively, and by whether feedback depends on a mediating continuous decision variable, or on the categorical choice. These alternative feedback types reflect the nature of the neural code, e.g. whether neuronal responses represent probabilities or log-probabilities in a Bayesian inference framework, and whether feedback is best explained in terms of, e.g., beliefs or predictive coding. Their different implications for CPs can therefore be used to distinguish between these theoretical models using empirical measurements. We derive analytical CP formulas valid under general conditions in which choices are mediated by continuous decision variables, independently of the presence of feedforward or feedback influences. For this case, we identify the CP relationships with other measurable single cell responses' properties, such as their variance and choice-triggered average, and compare these relationships with the case in which CPs are due to post-decision feedback based on the categorical choice. In the former, but not the latter case, nonzero signal stimuli and internal biases lead to a characteristic increase of CPs as a function of the departure from equifrequent behavioral choices. Next, we examine some simple representative cases: pure feedforward and pure feedback, and models with internal biases. Finally, we consider a two-stage model with a feedforward phase followed by post-decision feedback.

II-18. Distinct neural dynamics in two frontal areas contribute to multi-scale temporal control of decision

Masayoshi Murakami[1] MASAYOSHI.MURAKAMI@NEURO.FCHAMPALIMAUD.ORG
Hanan Shteingart[2] CHANANSH@GMAIL.COM
Yonatan Loewenstein[2] YONATAN.LOEWENSTEIN@MAIL.HUJI.AC.IL
Zachary Mainen[1] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]The Hebrew University of Jerusalem

The ability to wait for delayed gratification is an essential feature of adaptive behavior, but little is known about the underlying neural mechanisms.Here, we sought to determine how two different frontal areas—medial prefrontal cortex (mPFC) and the secondary motor cortex (M2)—contribute to control of waiting. To do so, we trained rats to perform a task in which the animal continuously chose whether to wait for a randomly-delayed tone to obtain a large reward or to give up and obtain a smaller reward. We found that waiting time varied substantially between trials and was well-approximated by a two-stage stochastic model: (1) individual times are drawn stochastically from an exponential distribution and (2) the mean of the distribution varies over a longer time-scale depending on the history of rewards. To identify the brain regions associated with these two distinct processes, we used multi-tetrode single-unit recordings. Remarkably, we found that the fast stochastic component was robustly encoded in M2 but not in mPFC. In contrast, the slow history-dependent component was correlated with neural activity in both areas, but with different temporal characteristics. While individual mPFC neurons maintained the slow signal in persistent activity present throughout a trial and in the inter-trial interval, M2 neurons showed only transient correlations with this signal. These observations suggest a model in which the mPFC maintains a slowly-varying, experience-dependent waiting signal and M2 translates this signal into precise waiting times, while injecting trial-to-trial variability. Consistent with these findings, measurements of firing rate autocorrelations showed neurons in mPFC exhibited much slower neural dynamics than those in M2 (741 vs. 373 ms). These results help to elucidate how, through different characteristic time scales of processing, different areas within the cortical hierarchy contribute to the multi-scale dynamics of behavioral adaptation.

II-19. Confidence in memories as statistical confidence

Paul Masset PMASSET@CSHL.EDU
Adam Kepecs KEPECS@CSHL.EDU

Cold Spring Harbor Laboratory

The capacity to learn and recall previously encountered objects or situations is central to adaptive behavior at timescales beyond those of sensory systems. The ability to know when to trust recalled memories and when to doubt them allows us to make plans about current and future actions. However the accuracy of such confidence reports in memories has been questioned in the literature. We specifically consider results from experiments of word recall revealing that confidence is positively correlated with recall accuracy for correct identifications but is negatively correlated for false identification of distractors. Here we present a model that accounts for this paradox. We show that the presence of negative correlations is expected when categorizing the data according to properties only available to the experimenter. We present a statistical explanation and an implementation of the model. Our analysis shows that confidence reports in a memory task have all the expected signatures of confidence based on the statistical definition of confidence: the probability of a choice being correct given the evidence available to the subject. The observed confidence reports across the population suggest that individual subjects share a common discriminability measure across an abstract ensemble of stimuli that cannot be mapped on an intuitive metric. Our model can explain confidence reports in memory without the need of heuristics or dual process models. More generally, we highlight that when analyzing the representation of a hidden variable that is available to the subject and thought to represent an external variable that is known or controlled by the experimenter, well-behaved transformations between these variables can lead to changes in the relationships

between variables such as the negative correlation between confidence and discriminability for errors. This has strong implications on the analysis of both behavioral and neural data.

II-20. Cognitive effort and the opportunity cost of time: a behavioral examination

Ross Otto ROTTO@NYU.EDU
Nathaniel Daw NDAW@PRINCETON.EDU

New York University

In many classes of tasks, an organism's behavior reflects a fundamental tradeoff between cognitive effort and accuracy. While the question when and why an individual decides to expend—or withhold—cognitive effort has received recent theoretical attention, these issues necessitate a computationally informed experimental approach. Here we sought to 1) quantify expenditure of cognitive effort, and 2) manipulate the costs and benefits of cognitive effort to demonstrate directly how people are sensitive to this effort-accuracy tradeoff. To do this, we extend a popular theory of opportunity costs and response vigor, grounded in reinforcement learning and optimal foraging theory (Niv et al. 2007). This account formalizes a trade-off between two costs: the harder work assumed necessary to emit faster actions and the opportunity cost inherent in acting more slowly. Because in many settings this cost equals the average reward rate, this framework predicts speedier behavior in richer environments. We extend this framework from physical to cognitive effort using established tasks, for which 1) the amount of cognitive effort demanded from the task varies from trial to trial and 2) expenditure of cognitive effort affects accuracy. In one experiment, 32 subjects completed a calibrated 2AFC perceptual discrimination task, while available rewards varied randomly from trial to trial. As hypothesized, subjects' response speeds and accuracies tracked experienced average reward rate: when the opportunity cost of time was high, subjects responded more quickly and less accurately on difficult discriminations. In fits of a drift diffusion model, these changes were accounted for by a reduced decision threshold. In a second experiment, 50 subjects completed a Simon response conflict task. On response-incongruent trials—for which correct responses demand cognitive effort—we again observed reward-rate-dependent speeding accompanied by a reduction in accuracy. Thus, across both studies, expenditure of cognitive effort tracked the opportunity cost of time.

II-21. Object and spatial layout crosstalk improves scene recognition accuracy

Drew Linsley LINSLEYD@BC.EDU
Christopher Madan MADANC@BC.EDU
Sean MacEvoy SEAN.MACEVOY.1@BC.EDU

Boston College

Scene recognition is a core function of the visual system that supports navigation and effective interaction with the local environment. Scene recognition mechanisms operate on both the identities of objects in scenes and scenes' intrinsic global features, such as spatial layout. A unified judgment of scene identity emerges from a combination of these resources. We demonstrated that these resources initially combine during perception, manifesting as a systematic bias of a scene's perceived spatial layout towards the layout typically associated with its objects (Linsley & MacEvoy, 2014). For instance, the spatial layout of a room containing an oven is perceived as more 'kitchen-like' than if the oven were absent. We proposed that this early combination aids scene recognition, reinforcing perception of a scene's ambiguous spatial layout with spatial information associated with its objects. In the present study we used artificial neural networks to test this theory. Two models were trained to classify scenes based on their objects and spatial layouts. During training, one model identified co-occurring object and spatial layout feature groups and filtered across each ('crosstalk model'), whereas the other model kept these features

separate ('independent model'). Only the crosstalk model encoded scenes as more spatially similar to their category average when their objects were visible than when perceptual masks obscured them. This mirrored the spatial layout bias previously observed behaviorally and with fMRI of parahippocampal cortex (PHC), indicating that filtering co-occurring features reproduces the bias. The crosstalk model also produced scene categorization decisions that were more similar to humans than the independent model, both when scenes were intact and when their objects were masked. Finally, we observed that the crosstalk model was more accurate than the independent model at scene categorization. This work supports the theory that perceptual combination of information about scenes' objects and spatial layout benefits recognition.

II-22. Biologically realistic deep supervised learning

Jordan Guerguiev[1] JORDAN.GUERGUIEV@MAIL.UTORONTO.CA
Timothy P Lillicrap[2] TIMOTHYLILLICRAP@GOOGLE.COM
Blake Richards[3] BLAKE.RICHARDS@UTORONTO.CA

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

Supervised learning refers to learning in neural networks under the guidance of an outside teaching signal. When supervised learning occurs in multi-layer neural networks (deep networks), the results can be very powerful. Deep networks can achieve human-level performance on a range of tasks, and deep learning leads to emergent representations that resemble those in the neocortex. However, the most successful learning algorithms for supervised learning in deep neural networks invoke mechanisms that are not biologically realistic. For example, the most common algorithm for supervised deep learning, error backpropagation, requires that a neuron in a network knows the synaptic connectivity of all the neurons downstream of it. To date, a biologically plausible mechanism for supervised learning in deep networks has not been proposed. Here, we demonstrate a biologically realistic form of deep supervised learning that merges the difference target propagation algorithm proposed by Yoshua Bengio's group, and the dendritic prediction algorithm described by Robert Urbanczik and Walter Senn. Inspired by studies of neuronal circuits in the neocortex, our model utilizes stochastic spiking neurons with three functional compartments: a soma where spiking occurs, a basal dendrite which receives feedforward sensory information, and an apical dendrite which receives feedback from upper layers in the network. In addition to generating a biologically plausible algorithm for deep supervised learning, we generate specific experimental predictions about how apical dendritic inputs should shape plasticity in the neocortex. If experimental evidence for a biologically realistic deep supervised learning algorithm could be found, it would constitute a major advance in our understanding of learning in the brain, and would help to further the unification of theory and biology in modern neuroscience.

II-23. Unsupervised learning in synaptic sampling machines

Emre Neftci[1] ENEFTCI@UCI.EDU
Bruno Pedroni BPEDRONI@ENG.UCSD.EDU
Siddharth Joshi[2] SIJOSHI@UCSD.EDU
Maruan Al-Shedivat ALSHEDIVAT.MARUAN@GMAIL.COM
Gert Cauwenberghs GERT@UCSD.EDU

[1]University of California, Irvine
[2]student

Recent studies have shown that synaptic unreliability is a robust and sufficient mechanism for inducing the stochasticity observed in cortex. Here, we introduce the Synaptic Sampling Machine (SSM), a spike-based stochastic neural network model that uses synaptic unreliability as a means to stochasticity for sampling. Synaptic unreliability in the SSM plays the dual role of a mechanism that naturally implements the probabilistic computations required for stochastic neural networks, and a regularizer during learning akin to DropConnect. Similar to the original formulation of Boltzmann machines, the SSM can be viewed as a stochastic counterpart of Hopfield networks, but where stochasticity is induced by a random mask over the connections. The SSM is trained to learn generative models with a synaptic plasticity rule implementing an event-driven form of contrastive divergence. We demonstrate this by learning a model of MNIST hand-written digit dataset, and by testing it in recognition and inference tasks. We find that SSMs outperform restricted Boltzmann machines (4.4% error rate vs. 5%), they are more robust to over-fitting, and tend to learn sparser representations. SSMs are remarkably robust to weight pruning: removal of more than 80% of the weakest connections followed by cursory re-learning causes only a negligible performance loss on the MNIST task (4.8% error rate). SSMs can thus offer substantial improvements in terms of performance and complexity over existing methods for unsupervised learning in spiking neural networks. These results further suggest that synaptic unreliability may play a critical role for probabilistic inference in the brain and that SSMs can provide very power-efficient learning machines in brain-inspired (neuromorphic) computing architectures.

II-24. Derivation of adaptive and self-calibrating neural networks for dimensionality reduction

Yuansi Chen[1] YUANSI.CHEN@BERKELEY.EDU
Cengiz Pehlevan[2] CPEHLEVAN@SIMONSFOUNDATION.ORG
Dmitri Chklovskii[2] DCHKLOVSKII@SIMONSFOUNDATION.ORG

[1]University of California, Berkeley
[2]Simons Foundation

Early sensory processing reduces the dimensionality of stimuli as evidenced by a high ratio of receptors to downstream neurons (e.g. in the human retina [#photoreceptors]/[#ganglion cells] ~ 100). To model such dimensionality reduction (DR) we have previously derived a network with local learning rules from a similarity matching principle. In this network, as well as in many others, the number of output dimensions is set a priori by the number of output neurons and cannot be changed. Because the number of informative dimensions (as opposed to noise) varies widely across sensory stimuli, DR networks must automatically adapt the number of output dimensions to the stimulus. Here, we derive such DR networks from the objective functions that combine the previously proposed Multidimensional Scaling (MDS) cost with novel regularizers that penalize the nuclear norm (a convex relaxation of rank) of the output covariance. Such objective functions are solved by soft-thresholding of the eigenvalues of input covariance either by a constant or by a fraction of mean covariance eigenvalue. From these objective functions we derive online algorithms which map onto neural networks with different learning rules. We compare our learning rule predictions with the results of LTP/LTD experiments on synaptic plasticity both in adult and developing animals. Therefore, our work provides a systematic theoretical framework of neural dimensionality and can be used to interpret and guide neural plasticity experiments.

II-25. Learning engages both high- and low-covariance modes of neural population activity

Matthew Golub[1] MGOLUB@CMU.EDU
Patrick Sadtler PSADTLER14@GMAIL.COM
Kristin Quick KRISTINMQUICK@GMAIL.COM
Stephen Ryu[2] SEOULMAN@STANFORD.EDU
Elizabeth Tyler-Kabara[3] ELIZABETH.TYLER-KABARA@CHP.EDU
Aaron Batista[3] APP@PITT.EDU
Steven Chase[1] SCHASE@CMU.EDU
Byron Yu[1] BYRONYU@CMU.EDU

[1]Carnegie Mellon University
[2]Palo Alto Medical Foundation
[3]University of Pittsburgh

Learning requires changes in how we generate neural activity. We and others have found that a small number of modes of activity (i.e., dimensions) can be used to explain the patterns generated by a simultaneously recorded population of neurons. Here we asked whether the modes of population activity that show the largest task-related modulation are the modes that play the largest role during learning. We leveraged multi-electrode recordings in a Rhesus monkey and a brain-computer interface (BCI) paradigm to investigate the neural strategies underlying learning. By applying abrupt perturbations to the BCI decoder, we could ask whether the learning that followed was confined to the high-covariance modes of population activity. Interestingly, we found that the low-covariance modes play a substantial role during learning. Further, we found that the monkey learned to activate the correct modes of activity at the correct times during the task, regardless of whether those modes showed high- or low-covariance prior to learning. Although the perturbations changed the magnitude of influence that each mode of activity had on driving the BCI cursor, we did not find evidence of compensatory dynamic range rescaling (i.e., to restore the influence a mode had prior to the perturbation). Taken together, these findings begin to elucidate the neural strategies subserving the brain's flexibility in acquiring novel and abstract skills.

II-26. Context dependent evidence integration in recurrent neural networks

Sepp Kollmorgen[1] SKOLLMOR@INI.PHYS.ETHZ.CH
Valerio Mante[2,1] VALERIO@INI.UZH.CH

[1]ETH Zurich
[2]University of Zurich

Many behaviors are context dependent, meaning that the same stimulus elicits different reactions in different situations. Neural circuits that support context dependent computation have been demonstrated in various biological systems. Corresponding modelling work typically involves solutions where context is represented and maintained through an external mechanism that is largely separate from the one performing the actual computation. Here we explore how maintenance of context and context-dependent computation can be achieved jointly by a single circuit and how the resulting dynamics differ from solutions where these two functions are performed separately. We trained two different nonlinear recurrent neural network models to perform context dependent decision task, requiring integration of a contextually relevant input while ignoring a second, irrelevant, input. The first network model (steady-context-network) receives a constant context input. The second model (initial context-network) has access to context information only through its initial unit activations (on the first timestep of each trial) and context has to be maintained by the same dynamics that also enable integration of evidence. Both trained models successfully integrate the relevant inputs over long time scales. The networks implement integration as gradual movement along one-dimensional attractors. While the steady-context-network produces two different attractors, one for each context, the intial-context-network generates a single, connected, attractor with different regions associated with different contexts and different local dynamics. For the steady context network, pertubations of

network activity decay only onto the attractor enabled by the active context signal, whereas pertubations of activity in the initial context network decay onto both regions of the common attractor. We demonstrated how context dependent integration without a steady context signal can be learned by a recurrent circuit. Since there is no evidence for separate, dedicated populations representing context in frontal areas, this is a good candidate for the solution implemented by those circuits.

II-27. Matching tutors and students: Effective strategies for information transfer between circuits

Tiberiu Tesileanu[1] TTESILEANU@GMAIL.COM
Bence Olveczky[2] OLVECZKY@FAS.HARVARD.EDU
Vijay Balasubramanian[3] VIJAY@PHYSICS.UPENN.EDU

[1]City University of New York
[2]Harvard University
[3]University of Pennsylvania

Neural circuits that learn information about the world can, over longer timescales, transfer that knowledge to downstream circuits. For example, hippocampus learns from immediate experience and likely consolidates the knowledge into long-term memories elsewhere. Similarly, motor cortex provides essential input to sub-cortical circuits during skill learning, but later becomes dispensable for executing certain skills. A paradigmatic example of such information transfer occurs in songbirds, where the anterior forebrain pathway (AFP) drives short-term improvements in song that are later consolidated in pre-motor area RA. Here, we show how to match instructive signals from tutor circuits to synaptic plasticity rules used by student circuits in order to achieve effective two-stage learning. For example, consider learning a sequential pattern where a timebase generated by one area is transformed into activity patterns by synaptic connectivity with a 'student' area. Learning is based on instructive inputs from a third, 'tutor', area. If potentiation (depression) at the timebase-student synapses is driven by recent tutor input above (below) a threshold, we argue that a good instructive signal would be proportional to the instantaneous difference between current post-synaptic response and target response. In contrast, if the timing of the tutor input relative to the fixed timebase determines the sign of synaptic modifications, we find that a good instructive signal accumulates the errors in student output as the motor program progresses. For plasticity rules that interpolate between these two, we show that an effective tutor signal integrates the student error over a certain timescale, which we relate to the structure of synaptic plasticity. Mismatch between student and tutor leads to slower and less accurate learning, and in some cases can completely abolish learning. For concreteness, we implemented these ideas in biologically-plausible models of the birdsong circuit. The models reproduce qualitative firing statistics of RA neurons in juveniles and adults.

II-28. Spine-size fluctuations support stable cell assembly learning in recurrent circuit models

James Humble[1] JAMES.M.D.HUMBLE@GMAIL.COM
Haruo Kasai[2] HKASAI@M.U-TOKYO.AC.JP
Taro Toyoizumi TARO.TOYOIZUMI@BRAIN.RIKEN.JP

[1]RIKEN Brain Science Institute
[2]The University of Tokyo

Cortical circuits rewire in an experience-dependent way. A major biological mechanism underlying this is Hebbian plasticity. In models of recurrently connected networks, ongoing Hebbian plasticity is often unstable in nature because of its positive feedback, e.g., a tightly coupled and coherently active group of neurons tends to drive other neurons well and expand the group (Kunkel et al., 2011). This typically fuses multiple memory patterns and results

in a deficiency in learning/memory performance. The biological mechanism that stabilizes Hebbian plasticity is unknown. Here we combine experimentally observed fluctuations of spine sizes (Yasumatsu et al., 2008) with spike-timing-dependent plasticity in recurrently connected neural networks. We show that an appropriate level of spine fluctuations is sufficient to stabilize memory patterns without fusing, and maintains a physiological volume distribution of spines in the presence of ongoing Hebbian plasticity. In addition to stabilizing Hebbian plasticity, we posit that abnormal spine fluctuations impair learning/memory performance. Our theory explains how high spine turnover rates, experimentally observed in several animal models for autism (Isshiki et al., 2014), cause slow learning and impairs memory performance.

II-29. Regularization-free synthesis of stable, information-optimal plasticity rules in recurrent networks

Sensen Liu LIUS@ESE.WUSTL.EDU
Gautam Kumar GAUTAM.KUMAR@WUSTL.EDU
ShiNung Ching SHINUNG@WUSTL.EDU

Washington University in St. Louis

Information optimization has proven to be a powerful paradigm to study the connection between the dynamics of neural plasticity and an overt functional objective. However, with a few notable exceptions, the synthesis of information-optimal plasticity rules has been limited to feedforward networks. In this scenario, the obtained rules are generically Hebbian, leading to network instability. Consequently, the information optimization problem is typically regularized by penalizing large neuronal firing rates, leading to a stabilizing anti-Hebbian component in the synthesized rules. However, regularization imposes, in essence, a prior constraint such that the emergent rule is not wholly a consequence of information optimization. The manifestation of information optimization in recurrent networks, with or without regularization, remains unresolved. We report the emergence of provably anti-Hebbian plasticity for maximization of mutual information in recurrent, discrete spiking networks. In contrast to prior recurrent instantiations, our synthesized rules do not involve regularization and, moreover, can be well-approximated via local interactions. More specifically, we: (i) derive the plasticity rule that maximizes the mutual information between the present and history in a recurrent, stochastic spiking network, (ii) prove that the recurrent dynamics impose a 'self-regularizer' so that the emergent rule is stable without prior constraints on firing rates, (iii) derive a recursive estimator, equivalent to adding synaptic dynamics to each connection, that enables approximation of the optimal rule via purely local pairwise computations, and (iv) show that both the optimal rule and its local approximation yield a network with a globally asymptotically stable branching ratio equal to one, so that the network exhibits emergent self-organizing criticality.

II-30. Long-term synaptic statistical learning: aiming for a target postsynaptic response

Rui Ponte Costa[1] RUI.COSTA@CNCB.OX.AC.UK
Zahid Padamsey[1] ZAHID.PADAMSEY@PHARM.OX.AC.UK
James D'amour[2] JAMES.DAMOUR@MED.NYU.EDU
Robert Froemke[2] ROBERT.FROEMKE@MED.NYU.EDU
Nigel Emptage[1] NIGEL.EMPTAGE@PHARM.OX.AC.UK
Tim P Vogels[1] TIM.VOGELS@CNCB.OX.AC.UK

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

Changes in synaptic strength are believed to be the neuronal basis of learning. A multitude of experimental protocols have been used to induce such long-term synaptic plasticity. However, the exact amplitude of change, the duration of the effect and its locus of expression (pre- or postsynaptic) is often variable. The underlying causes for these variabilities have remained unclear. Here we introduce a framework in which long-term plasticity aims to optimize synaptic transmission statistics towards a presumed target strength. Consequently, the exact pre- and postsynaptic states at the time of plasticity induction determine the observed ratio of pre/post modifications. Using this framework we can explain the locus of synaptic changes observed in individual hippocampal and neocortical potentiation and depression experiments. Furthermore, our framework predicts the duration and magnitude of total weight modifications based on the initial synaptic state, and explains the expression locus of plasticity for postsynaptically silent synapses. Finally, we demonstrate that synergistic pre/post plasticity enables neural networks to learn near-optimal associations between selective neurons (e.g. to spatial or sensory features). By showing that synaptic response statistics are optimized towards a functional target, our work may lay to rest the long standing pre/post debate, and help to explain the high degree of variability in weight modification typically observed in experiments.

II-31. STDP is a reflection of spike cross-correlation: a derivation from similarity matching

Dmitri Chklovskii[1] DCHKLOVSKII@SIMONSFOUNDATION.ORG
Tao Hu[2] FOXWAVE@TAMU.EDU
Cengiz Pehlevan[1] CPEHLEVAN@SIMONSFOUNDATION.ORG

[1]Simons Foundation
[2]Texas A&M

In simulations of spiking neural networks, where synaptic weights evolve according to the Spike-Timing Dependent Plasticity (STDP) rule, the STDP kernel is typically specified a priori and applied uniformly to all neurons. Here, we propose an alternative view of neural plasticity, in which the STDP kernel is itself learned and reflects spike cross-correlation between corresponding neurons. Our approach is based on the recently developed similarity matching framework which allows one to derive both neuronal activity dynamics and synaptic weight plasticity from a principled objective function [1]. Whereas, previously, similarity matching was applied to a series of static inputs, here, we formulate similarity matching for temporal sequences. To this end we form data matrices out of time series by concatenating columns of lagged vectors resulting in a Hankel matrix structure. We formulate an optimization problem by requiring that the output of a neuron preserves the similarity of its centered inputs under the constraint on the rank of the output data matrix. This problem is solved by a series of double-exponential kernels timed according to the output spike train. Furthermore, we derive a synaptic weight update which depends on the relative timing of pre- and postsynaptic spikes as the projection of the centered spike cross-correlation function on the double-exponential kernel. If, as commonly observed experimentally, the centered spike cross-correlation function changes sign at zero we obtain a classical STDP kernel. For other forms of the cross-correlation functions we predict that the STDP kernel will be different. Whereas both different STDP kernels and different cross-correlation functions have been observed for different synapses the two functions have not been measured for the same synapse. Our theory can be tested by measuring the two, previously thought to be unrelated, functions for the same synapse and checking for correspondence.

II-32. Learning to surf the wave: unsupervised learning of cortical delay responses with propagating waves

David Barrett DAVID.BARRETT@ENG.CAM.AC.UK
Mate Lengyel M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

Various cortical areas, most prominently the prefrontal cortex and the hippocampus, display diverse, stimulus-dependent neural responses during the delay period of tasks with a working memory component (Siegel et al.

2011). The timescales of these responses are much longer than single neuron timescales. However, the principles that determine how cortical circuits learn to make use of such representations are not fully understood. We adopt a two-step approach to this problem. First, we assume that the brain learns an internal (generative) model of temporal sequences in its environment, such as the contingency between conditioned (CS) and unconditioned stimuli (US) in trace conditioning. Theoretical analyses of behavioral experiments suggest that animals do learn full generative models of conditioned and unconditioned stimuli using internal latent causes (Courville et al. 2003, 2006). Next, we construct and train a spiking neural network to perform inference and make predictions under such an internal model, with separate populations of neurons representing observed and latent variables. Neurons representing latent variables maintain a memory of a CS by triggering a traveling wave in the latent variable population, with a phase that is determined by the CS identity. The network learns to “surf the wave" and generate a response after the correct delay, by maximizing a lower bound on the likelihood of the model using variational inference. This allows a spiking network to learn long delays, more than a hundred times longer than single neuron timescales, and longer than in previous unsupervised learning models (Rezende and Gerstner, 2014; Brea et al. 2013). We show that the latent representation is consistent with activity of neurons in the prefrontal cortex recorded during a trace-conditioning task (Siegel et al. 2011). Our work provides a template for probabilistic inference and learning for complex conditioning paradigms, delay timing tasks and spatial-temporal signals.

II-33. Ramping up to an explanation of accumbens dopamine signals

Kevin Lloyd KLLOYD@GATSBY.UCL.AC.UK
Peter Dayan DAYAN@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

A cornerstone of current theorizing about dopamine’s computational role is the idea that the phasic activity of dopamine neurons represents a temporal difference (TD) prediction error. While substantial evidence supports this mapping, recent reports of ramp-like increases in accumbens dopamine concentration when animals are about to act, or are about to reach rewards, pose an important challenge to this TD hypothesis. This is because, under a TD account, the implied activity underlying such ramps is persistently predictable by preceding stimuli and as such, should be largely predicted away. Nevertheless, we suggest that dopamine ramps are largely reconcilable with standard theory, and offer three, non-mutually exclusive accounts of these phenomena. Firstly, at a computational level, we propose, along with Berke, that ramping may arise as a form of state prediction. In average-reward analyses, average reward rate has been suggested to (i) be a comparison point for the TD error, (ii) control instrumental vigour, and (iii) be represented by tonic dopamine levels. We determine that a suitable counterpart in the discounted case is a state-dependent quantity which, carried by dopamine, would manifest as a ’quasi-tonic’ signal with similar properties to those observed experimentally. Secondly, at an algorithmic level, we suggest that ramping observed just before execution of an instrumental action for reward may be caused by uncertainty about when the action will occur. TD errors occasioned by resolution of such uncertainty, such as may occur just prior to action execution, may explain these signals. Thirdly, at an implementational level, we observe that ramping may arise if dopamine has a direct influence on the timecourse of choice, such as setting the gain of an accumulative decision-making process. Even assuming purely noise-driven fluctuations in dopamine levels, resulting correlated dynamics entail an average dopamine signal which appears to ramp up to the time of decision.

II-34. A bias-variance trade-off that governs individual variability in learning and decision-making.

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

Humans often make variable decisions when presented repeatedly with the same stimuli. This trial-to-trial decision noise has typically been considered adaptive only when it promotes exploration of new, potentially useful information sources. Here we present a theory that posits that even when no new information is obtained, decision noise can reflect adaptive updating of uncertain models of the environment. A key prediction is that uncertainty about statistical structure governs a fundamental bias-variance trade-off in decision-making between adaptability (the ability to adapt to new statistical structures; i.e., inverse bias) and precision (the ability to make reliable choices in a stable environment with known structure; i.e., inverse variance). We investigate this theory in the context of change-point detection in volatile environments that challenge subjects to distinguish variability coming from a stable state from volatile changes in the underlying state itself. We recently showed that subjects can learn a hazard rate parameter that governs the rate of volatile changes, but with an average bias towards an intermediate value. The degree of learning varied considerably across subjects, with some showing more bias and others showing more learning. Here we show that subjects who learned more also tended to make noisier choices, even after accounting for the learned hazard rates. The data can be explained by a biologically plausible particle-filter algorithm that operates in model space by recursively: (1) updating the weight given to each parameter under consideration, and (2) randomly replacing the least weighted parameters with those from a prior distribution. The algorithm effectively captures the bias-variance trade-off evident in the behavioral data in terms of the precision of the prior distribution. This model represents a useful starting point for considering how the brain governs the inherent trade-off between adaptability and precision that is reflected in individual differences in human decision-making and learning behavior.

II-35. Plasticity of adaptation stabilizes neural activity and induces temporal decorrelation

Carlos Brito CARLOS.STEIN@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH
Ecole Polytechnique Federale de Lausanne

Spike-triggered adaptation regulates the activity level in single neurons and has been associated with an efficient encoding of information. The tuning of adaptation to input properties raises the question of whether adaptation could be learned. We show how spike-timing dependent plasticity (STDP) of adaptation allows for the self-organization of a neuron's activity. Similar to synaptic plasticity models, the amplitude of the adaptation current is modified at each spike, modulated by a trace of previous spikes, with a short time scale for potentiation and a long time scale for depression. Stability is reached when the auto-correlation of the spiking activity is equal at these two time scales. When injecting strong excitatory input currents, adaptation became up-regulated, and the output activity transitions from repetitive to irregular spiking, with a balance between excitation and the effective adaptation current. For inputs with power-law statistics, multiple adaptation currents with distinct time-scales were considered, leading to the learning of near power-law adaptation and consequent temporal whitening of the output activity. These findings indicate how instrinsic plasticity may be implemented in cortical neurons, allowing for stable activity regimes and an efficient temporal coding that is flexible to diverse input statistics.

II-36. Hierarchical selection, reward-dependent metaplasticity, and choice under uncertainty

Shiva Farashahi[1] SHIVA.GHAANI@GMAIL.COM
Katherine Rowe[1] KATHERINE.A.ROWE.16@DARTMOUTH.EDU
Zohra Aslami[1] ZOHRA.V.ASLAMI.18@DARTMOUTH.EDU
Daeyeol Lee[2] DAEYEOL.LEE@YALE.EDU
Alireza Soltani[1] ALIREZA.SOLTANI@DARTMOUTH.EDU

[1]Dartmouth College
[2]Yale University

Most real-world decisions require evaluating myriad of alternative objects that provide reward information based on varying combinations of their features, while that information can unpredictably change over time. This means that the brain has to not only learn reward values associated with relevant objects (e.g. red square) and multiple features of those objects (e.g. color, shape, etc.), but also use the proper model of the environment to determine how this information should be combined for making decisions. Heterogeneous neural activity (related to multiple choices and their outcomes) observed throughout the brain could provide a neural substrate for representation and evaluation of alternative models of the environment, but the specific mechanisms are unknown. Here we used a combination of computational and experimental approaches to reveal neural mechanisms underlying proper selection and update of the decision maker's internal models under uncertainty. We designed an experimental paradigm in which the subject chooses between pairs of visual targets (defined with multiple visual features and feature values) and receives feedback on every trial (Fig.1A). The reward on a given choice is determined by multiple features, allowing decision making based on alternative models (feature-based or object-based), each of which yields different performance. Experimental results showed that human subjects are able to adopt the better performing model of the environment based on reward feedback (Fig.1B). In addition, we tested three alternative scenarios for how reward feedback enables correct model adoption and update. The results from fitting subjects' behavior and testing of computational models' robustness revealed a plausible mechanism: hierarchical selections of the most informative sources of information in individual models, followed by a competition between selected sources across models to determine the choice, and finally update via reward-dependent metaplasticity. These results suggest the contribution of heterogeneous neural activity and reward-dependent metaplasticity to choice under uncertainty.

II-37. The dopamine signal in decision making tasks with temporal uncertainty

Stefania Sarno[1] STEFANIA.SARNO@UAM.ES
Victor de Lafuente[2] LAFUENTE@UNAM.MX
Ranulfo Romo[2] RROMO@IFC.UNAM.MX
Nestor Parga[1] NESTOR.PARGA@UAM.ES

[1]Universidad Autonoma de Madrid
[2]Universidad Nacional Autonoma de Mexico

Animals live in uncertain environments where they have to make decisions based on noisy sensory information and unknown timing of the relevant events. To maximize possible rewards obtained from the environment they learn how to evaluate and select their actions. Although reinforcement learning (RL) could be an appropriate framework to deal with this issue, its use in decision-making tasks with uncertain temporal and sensory information has not been investigated. To obtain further insight on this problem we used a detection task in which the time of a possible stimulation is random (de Lafuente and Romo, 2011). This task defines a non-trivial problem in which RL methods can be developed and tested. Here we reanalyze midbrain neurons recordings and propose a RL model containing a Bayesian inference module that estimates joint beliefs about stimulation and timing of the task events. Using these beliefs an actor-critic module learns to evaluate and select an action. The main results

are: 1)1. The dopamine signal and the reward prediction error exhibit similar phasic responses and modulated tonic activity that depend on the trial type. As in the data, at the go cue the model produces a phasic response reflecting the uncertainty about the future decision. 2)2. In hit and false alarm trials a phasic response occurs when an event is detected. This produces a large peak in hit trials. Instead, events responsible for false alarms (Carnevale et al, 2015) are distributed over the possible stimulation window, generating a weak modulation of the tonic activity. No effect appears in trials with stimulus-absent choices.3)3. Before the go cue, both data and model present a decreasing tonic activity. In correct rejection trials it results from the variable duration of the trial. In hit trials it comes mainly from the finite resolution in the estimation of time intervals.

II-38. Parsing sensorimotor dysfunction from reward-related abnormalities in depression using an inverse optimal control model

He Huang[1] CRANE081@GMAIL.COM
Katia Harle[2] KATIA.HARLE@GMAIL.COM
Javier Movellan[2] JRMOVELLAN@GMAIL.COM
Martin Paulus[1] MPAULUS@LAUREATEINSTITUTE.ORG

[1]Laureate Institute for Brain Research
[2]University of California, San Diego

Depression is associated with impairments in sensorimotor function and reward processing. While animal research suggests these two processes are mediated by distinct neural systems, it has been difficult in human experiments to distinguish their respective influence in goal-directed behavior. Differentiating these process dysfunctions in depression is important because it can help (a) delineate the molecular basis and neural circuitry of different cognitive dysfunctions, (b) quantify and possibly predict clinical outcomes, and (c) develop targeted interventions such as computer-based cognitive retraining. Here we propose to use inverse optimal control modeling, a computational approach that incorporates sensorimotor system in determining the reward function that the individual is optimizing, to disambiguate depression-related performance deficits into three components: sensorimotor speed, goal state, and motivation (the amount of effort one is willing to spend to achieve the goal). 66 subjects with none to severe depression symptoms ($0 \leq BDI \leq 39$) completed two tasks. In Task 1, they pushed a joystick as quickly as possible once they observed motion onset of a virtual car. In Task 2, subjects were instructed to move a virtual car as quickly as possible and stop it as close as possible to a stop-sign. We estimated sensorimotor speed from Task 1 and recovered the underlying reward function (includes goal state and motivation) that best explain subject's behavior in Task 2. Results show that, relative to healthy controls, depressed individuals: 1) have slower sensorimotor speed; 2) formulate different goals (aiming further away from the instructed goal). Severely depressed individuals also showed significantly lower motivation to achieve their goals relative to non-depressed and mildly depressed individuals. In summary, the inverse optimal control framework can disambiguate the dysfunctions of goal-directed behavior in depression and quantify the degree to which an individual with depression shows deficits in sensorimotor deficits, reward-related goal setting, and motivation.

II-39. A minimal neural mechanism for explorative behaviors

Ran Darshan[1] RAN.DARSHAN@MAIL.HUJI.AC.IL
Bill Wood[2] BILL.E.WOOD@GMAIL.COM
Susan Peters[3] SPETERS@DUKE.EDU
Arthur Leblois[2] ARTHUR.LEBLOIS@PARISDESCARTES.FR
David Hansel[2] DAVID.HANSEL@UNIV-PARIS5.FR

[1]The Hebrew University of Jerusaelm
[2]Universite Paris Descartes
[3]Duke University

Motor exploration is an essential component of sensorimotor learning. While motor exploration is expressed in adults as movement variability around a stereotyped motor pattern, young individuals produce explorative disorganized behaviors referred to as ‘motor babbling’. A common neural circuit responsible for motor exploration at all stages of development has been recently identified in songbirds. However, the cellular and network mechanisms involved in the generation of such exploratory motor output remain unknown. Here we show that motor exploration during vocal babbling shares common features among humans and various bird species, in contrast to the very different vocalizations they produce after learning. We then characterize theoretically the fundamental minimal requirements on the architecture and the dynamics in neuronal circuits capable of autonomously generating motor exploration. We show that the circuit should comprise a premotor and a motor network, connected in a feedforward topographic manner, both strongly recurrent and both operating in the balance excitation-inhibition regime. The premotor network generates spatially uncorrelated, highly irregular, neuronal activity. The spatial correlations necessary to drive irregular motor behavior emerge from the interplay of the balance of excitation and inhibition in the motor network and the topography in the feedforward projections it receives. Our work gives theoretical foundations to understand how vocal exploration is generated. It predicts that spatiotemporal properties of the neuronal activities differ along the variability-generating circuit in songbirds: weak spatial correlations in the premotor nucleus, LMAN, and correlated activity in the downstream nucleus, RA. Simultaneous recordings of multiple neurons in these structures during singing in finches confirm these predictions. In a more general perspective, our work pave the way for a mechanistic understanding for how the brain autonomously generate noise in the single neuron level and can make use of it for behavior.

II-40. A hierarchy of time scales supports chunking strategy during sequence learning

Samuel Muscinelli SAMUEL.MUSCINELLI@EPFL.CH
Johanni Brea JOHANNI.BREA@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Behaviorally relevant sequences are highly complex and span multiple time scales. The neuronal mechanisms that allow learning of such complex temporal structures are not well understood. As a specific case, we focus on motor sequence learning, in which subjects have to learn to generate fixed sequences of simple movements. Analyzing the latencies between adjacent movements, several studies [1, 2] have proposed that learned sequences are produced as connected chunks and not as a uniform chain. Here we propose an approach to sequence learning that exploits chunking and that is implementable in a neural network. We develop a multi-layer, rate-based, hierarchical model that features multiple, pre-wired effective time scales. Exploiting this structure, a target sequence can be learned using a biologically plausible, Hebbian learning rule. Due to co-activation, a unit with slow time scale develops one-to-many connections towards faster units and thus becomes associated to a chunk. This mechanism can be applied hierarchically to form higher-order chunks. The model can qualitatively reproduce the behavioral results of sequence learning experiments, i.e. latency patterns and error distributions. More generally, it can learn, after a few repetitions, complex sequences that have a broad range of time scales. The model requires the existence of a hierarchy of time scales, for which there is a growing evidence, both in human and in primate cortices [3, 4]. Using our approach, once a chunk-related unit is formed, it can be immediately reused to learn a different sequence in which the same chunk is present. This gives advantages both in terms of consumption of neural resources and in terms of learning time.

II-41. The motor cortex and supplementary motor area exhibit different population-level dynamics

Mark Churchland MC3502@COLUMBIA.EDU
Antonio Lara AHL2143@CUMC.COLUMBIA.EDU
John Cunningham JPC2181@COLUMBIA.EDU

Columbia University

Different cortical areas presumably perform different computations. Consider the supplementary motor area (SMA) and motor cortex (M1). SMA appears to be involved in movement timing. M1 participates in generating descending commands. SMA and M1 thus presumably (1) process different information and (2) perform different computations via different network dynamics. The first expectation has been addressed by examining single-neuron responses. To address the second expectation, we analyzed SMA and M1 population responses recorded during the same reach task. We employed a novel 'hypothesis guided dimensionality reduction' approach which translates a hypothesis into a cost function. The minimum-cost projection reveals the presence or absence of hypothesized structure. Our hypothesis was motivated by a central motif found in prior M1 simulations and results: a large overall translation of the neural state followed by rotational dynamics. Our cost function thus sought projections where some dimensions exhibit a condition-independent response (e.g., similar for left vs. right reaches) while other dimensions capture trajectories that follow linear dynamics (not necessarily rotational). The M1 population response displayed the hypothesized structure, consistent with prior results. A large condition-independent translation of the neural state preceded movement onset. Subsequently, the neural state evolved according to dynamics dominated by rotations. The SMA population response did not display the hypothesized structure: no clear rotations or other linear dynamics were present. This finding supports the idea that SMA and M1 perform different computations. The specificity of the central M1 motif (absent in upstream SMA and downstream muscle populations) supports the hypothesis that the motif is related to pattern generation rather than generic aspects of the data. Yet despite different dynamics, some signals were shared. SMA and M1 exhibited a nearly identical condition-independent translation prior to movement initiation. Subsequently, aspects of the oscillatory patterns produced by rotational M1 dynamics also appeared in the muscles.

II-42. Closed loop optogenetic control of neural circuits: Tracking dynamic trajectories of neural activity

Michael Bolus[1,2] MBOLUS@BELLSOUTH.NET
Adam Willats[3,2] AWILLATS3@GATECH.EDU
Clarissa Whitmire[3,2] CLARISSA.WHITMIRE@GATECH.EDU
Zak Costello[3] ZAK.COSTELLO@GATECH.EDU
Magnus Egerstedt[3] MAGNUS.EGERSTEDT@ECE.GATECH.EDU
Christopher Rozell[3] CROZELL@GATECH.EDU
Garrett Stanley[3,2] GARRETT.STANLEY@BME.GATECH.EDU

[1]W.H. Coulter Dept. BME, GT
[2]Emory University
[3]Georgia Institute of Technology

Previously we demonstrated using closed loop optogenetic control to clamp the firing rate of single neurons to static targets. Clamping firing rates to fixed levels has many potential applications as a systems-level analog of Hodgkin and Huxley's voltage clamp. However, time-varying activities in the brain, such as mode-switching and brain state oscillations, are thought to be critical in mediating behavior and perception, which motivates developing a tool to replicate and manipulate these signals. Here we develop a framework for tracking dynamic trajectories, using optogenetic control of single unit thalamic firing activity in-vivo. Specifically, we utilize the thalamocortical circuit in the rat vibrissa pathway as a model system, closing the loop around the spiking activity of single units expressing channelrhodopsin-2 (ChR2) in somatosensory thalamus. Integral to the framework is the implementation of an observer designed to estimate the latent firing rate from spiking activity, providing feedback for control. We take a first order approach to observing firing rate using a fixed bandwidth exponential filter as well as a more sophisticated adaptive approach. Experimentally and computationally we find controller performance is highly sensitive to observer design. When the time constant of the exponential filter (tau) is tuned according to the target frequency, control of some dynamic targets can be achieved but at the cost of increasingly noisy estimates of firing rate. To ameliorate this tradeoff in controller performance, we have taken an adaptive point process filter (aPPF) approach, whereby the filter used to estimate firing rate is modulated according to spiking activity itself. Using an aPPF improves not only the ability to track dynamic firing rate trajectories, but also the fidelity of the estimate. Taken together, these developments represent a fundamental building block for control of neural activity which we are extending to larger-scale problems of multi-unit-, population-, and systems-level control.

II-43. Efficient state-space modularization for planning: theory, behavioral and neural signatures

Daniel McNamee D.MCNAMEE@ENG.CAM.AC.UK
Daniel Wolpert WOLPERT@ENG.CAM.AC.UK
Mate Lengyel M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

Even in state-spaces of modest size, planning is plagued by the ‘curse of dimensionality’. Hierarchically organized modular representations have long been suggested to underlie the capacity of the nervous system to efficiently and flexibly plan in high-dimensional environments. In such a modular representation, planning can first operate at a global level across modules acquiring a high-level ‘rough picture’ of the trajectory to the goal and, subsequently, locally within each module to ‘fill in the details’. Having the right modularization can thus greatly aid planning. However, the principles underlying efficient modularization remain obscure, making it difficult to identify the behavioral and neural signatures of modular representations. In particular, previous approaches computed optimal state-space decompositions based on optimal policies and optimal value functions thus requiring, a priori, knowledge of the environment solutions. Here, we compute a modularization optimized for planning directly from the transition structure of the environment without assuming any knowledge of optimal behavior. We propose a normative theory of efficient state-space representations which partitions an environment into modules, by minimizing the average (information theoretic) description length of path planning within the environment thereby optimally trading off the complexity of the global and local planning processes. We show that such optimal representations provide a unifying account for a diverse range of hitherto unrelated phenomena at multiple levels of representation and behavior: the strong correlation between hippocampal activity and ‘degree centrality’ in spatial cognition, the appearance of ‘start’ and ‘stop’ signals in sequential decision-making, ‘task-bracketing’ in goal-directed control, and route compression in spatial cognition.

II-44. Timing during transitions in Bengalese finch song: implications for motor sequencing

Todd Troyer[1] TODD.TROYER@UTSA.EDU
Michael Brainard[2] MSB@PHY.UCSF.EDU
Kristofer Bouchard[3,2] KRISTOFER.BOUCHARD@GMAIL.COM

[1]Biology Dept. and Neurosciences Institute
[2]University of California, San Francisco
[3]Lawrence Berkeley National Laboratory

Nearly all behaviors consist of sequences of more basic actions arranged in time. Bengalese finches (BFs) sing sequences of 40-200 ms vocalizations known as syllables. Songs include ‘branching syllables’ where transitions to several possible subsequent syllables appear to be probabilistic. Here we examine the relationships between transition timing (gaps) and transition frequency statistics to gain insight into the neural mechanisms for producing of complex sequential behaviors. One hypothesis is that after branch syllables, activation builds within distinct neural populations coding for the transition to subsequent syllables. Transitions are triggered when the first population reaches a threshold level of activation. Populations receiving greater input should reach threshold faster and win the competition more often, suggesting that a race mechanism would lead to negative correlations between transition probability and gap durations, as previously reported and observed here. However, we show that race models do not produce the observed correlation: to be observed, low probability transitions must outcompete high probability transitions, leading to short observed latencies for both outcomes. Taking advantage of BF song branching structure, we investigated whether the periods before and after syllable selection might account for variations in gap duration. We find that 63% of gap variation can be predicted from the average gap duration for the other transitions converging on a target syllable. This suggests that variations in latency between syllable selection and syllable initiation accounts for the bulk of variation in gap duration. To examine possible contributions to this latency, we determined that the frequency of the target syllable explains 14% of gap variance, and this increases to 23% when considering transition frequencies averaged across converging transitions. We hypothesize that the latency to syllable initiation decreases with increased practice, but the number of possible neural pathways leading to the excitation of that syllable also play a role.

II-45. Spatial averaging behavior in reach decisions reflects motor rather than visual averaging

Vassilios Christopoulos[1] VCHRISTO@CALTECH.EDU
Vince Enachescu[2] ENACHESC@USC.EDU
Paul Schrater[3] SCHRATER@UMN.EDU
Stefan Schaal[2] SSCHAAL@USC.EDU

[1]California Institute of Technology
[2]University of Southern California
[3]University of Minnesota

Classic psychological theories view decision and action as two separate cognitive processes - i.e., action planning begins only after a decision is made. Recent studies challenge this theory suggesting that decisions emerge via a competition between internal representations of potential actions. According to these studies, the brain generates in parallel multiple actions that compete for selection and accumulates value information to bias the competition until one of the actions is selected. In large part this theory is based on reaching experiments showing that when people are presented with multiple competing targets, reaches often launch to a spatially averaged location between the targets. Although this behavior can be interpreted as evidence of ‘motor averaging’ (i.e., simultaneous preparation of multiple competing actions), it has been argued that it reflects ‘visual averaging’ (i.e, single action towards an averaged location). In the current study, we aim to dissociate these two competing hypotheses. We designed a ‘reach-before-you-know’ experiment in which subjects had to perform reaches to two potential targets presented simultaneously in both hemifields. Critically, the actual goal location was unknown prior to movement onset. Control trials with single targets presented either in the left or the right space interleaved with the two-target trials. We modeled the task within a neurodynamical framework and showed that if decisions emerge through a competition of multiple actions, the reaction time should vary with the target probability and the approach direction of the reaches. Consistent with the model predictions, subjects had faster responses in two-target trials in which one of the targets was assigned with higher probability than in trials with equiprobable targets. Additionally, the reaction time was longer for reaches aimed to an intermediary location. Our findings provide direct evidence that the brain plans multiple actions prior to selecting one of them in reach decisions.

II-46. Short-term motor learning through sensorimotor temporal gating in a central serotonergic circuit

Takashi Kawashima[1,2] KAWASHIMAT@JANELIA.HHMI.ORG
Chao-Tsung Yang[1,2] YANGC@JANELIA.HHMI.ORG
Brett Mensh[1,2] MENSHB@JANELIA.HHMI.ORG
Misha Ahrens[1,2] AHRENSM@JANELIA.HHMI.ORG

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

Learning the relationships between motor action and sensory feedback is critical for adapting behavior to changes in the body and the environment. Although neural circuits mediating long-term forms of motor learning have been studied in depth, the neural substrate of short-term motor learning - the adjustment of movements over timescales of seconds or minutes - is unknown. We use whole-brain neuron-resolution calcium imaging in zebrafish to identify an essential role for serotonergic neurons in the dorsal raphe nucleus (DRN) for short-term motor learning. We find that DRN neurons encode the visual feedback that accompanies every swim bout. By de-synchronizing the visual stimulus from motor output, we find that these DRN neurons encode the visual feedback only when it is coincident with a swim bout This temporal gating of sensory feedback by motor action may serve to distinguish self-generated sensory input from sensory input generated by external sources. Further, when the fish is trained under higher feedback gain to adapt their motor output, DRN activity accumulates and turns into persistent firing. This firing underlies a short-term motor memory, so that the adapted behavior persists for over 10 seconds even in the absence of behavior and sensory feedback. These results reveal a novel neural mechanism for short-term motor learning, and extend the known repertoire of the functions of the serotonergic system.

II-47. Dynamic modulation of feed-forward circuit function and its dysregulation in fragile X syndrome

Vitaly Klyachko[1,2] KLYACHKO@SALK.EDU
Sarah Wahlstrom-Helgren[3] WAHLSTROMHELGREN@WUSTL.EDU

[1]Dept Cell Biology and Physiology
[2]Dept Biomedical Engineering
[3]Washington University

Feed-forward inhibitory (FFI) circuit is a canonical unitary circuit found throughout the brain and essential for many fundamental computations. It is widely believed that FFI circuit function depends on the fine-tuned balance of excitatory and inhibitory (E/I) circuit components that are time-locked with each other. The E/I balance within the FFI circuits is rapidly and dynamically modulated by the ongoing activity due to the short-term plasticity (STP) of both excitatory and inhibitory components. How dynamic changes in the E/I balance influences fundamental FFI circuit operations during ongoing activity remains poorly understood. Moreover, E/I circuit imbalance has been implicated in many neurodevelopmental disorders associated with intellectual disability, including Fragile X syndrome (FXS), the leading genetic cause of intellectual disability and autism. How fundamental operations performed by FFI circuits are affected in this disorder remains largely unexplored. Here, we examined how synaptic dynamics during spike bursts modulate the functions of two unitary hippocampal FFI circuits, both in healthy conditions and in the FXS mouse model.

II-48. Training excitatory-inhibitory recurrent networks for cognitive tasks: A simple, flexible framework

Francis Song FRANCIS.SONG@NYU.EDU
Robert Yang GY441@NYU.EDU
Xiao-Jing Wang XJWANG@NYU.EDU

New York University

The ability to simultaneously record from large numbers of neurons in behaving animals has ushered in a new era for the study of the neural circuit mechanisms underlying cognitive functions. One promising approach to uncovering the dynamical and computational principles governing population responses is to analyze model recurrent neural networks (RNNs) that have been optimized to perform the same tasks as behaving animals. Because the optimization of network parameters specifies the desired output but not the manner in which to achieve this output, "trained" networks serve as a source of mechanistic hypotheses and a testing ground for data analyses that link neural computation to behavior. Complete access to the activity and connectivity of the circuit, and the ability to manipulate them arbitrarily, make trained networks a convenient proxy for biological circuits and a valuable platform for theoretical investigation. However, existing RNNs lack basic biological features such as the distinction between excitatory and inhibitory units (Dale's principle), which are essential if RNNs are to provide insights into the operation of biological circuits. Moreover, trained networks can achieve the same behavioral performance but differ substantially in their structure and dynamics, highlighting the need for a simple and flexible framework for the exploratory training of RNNs. Here, we describe a framework for gradient descent-based training of excitatory-inhibitory RNNs that can incorporate a variety of biological knowledge. We provide an implementation based on the machine learning library Theano, whose automatic differentiation capabilities facilitate modifications and extensions. We validate this framework by applying it to well-known experimental paradigms such as perceptual decision-making, context-dependent integration, multisensory integration, parametric working memory, and motor sequence generation. Our results demonstrate the wide range of neural activity patterns and behavior that can be modeled, and suggest a unified setting in which diverse cognitive computations and mechanisms can be studied.

II-49. Fluctuating regimes and learning in excitatory-inhibitory rate networks

Francesca Mastrogiuseppe FRANCESCA.MASTROGIUSEPPE@ENS.FR
Srdjan Ostojic SRDJAN.OSTOJIC@ENS.FR

Ecole Normale Superieure

Recurrent networks of non-linear rate units can display a large variety of dynamical regimes depending on the structure of their synaptic connectivity. Classical studies have shown that randomly coupled networks can exhibit a transition from a fixed point to chaotic activity, where firing rates fluctuate in time and across different units [Sompolinsky et al, 1988]. More recent works have highlighted the potential computational capacities of this novel dynamical regime [Sussillo and Abbott, 2009; Laje and Buonomano, 2013]. Most previous studies focused on highly simplified random networks, where excitation and inhibition are not segregated. In these models, furthermore, firing rates fluctuate symmetrically around zero in the positive and negative direction, so that the transition to chaos is characterized solely in terms of second order statistics. To help bridge the gap between classical models and more realistic spiking networks, here we investigate the dynamics of a rate network that includes additional biological constraints. In this perspective, we adopt a sparse, two-populations connectivity matrix which obeys Dale's law. We select positively defined activation functions, and we introduce noisy inputs into each unit to mimic spiking noise. Extending the dynamical mean field theory, we show that network dynamics can be effectively described through two coupled equations for the mean activity and the auto-correlation function. As the synaptic coupling is pushed above the critical value, we find that the onset of chaotic fluctuations strongly perturbs the first order statistics, driving an increase in the mean firing rate. Moreover, we show that above the transition two different fluctuating regimes can be distinguished: for moderate synaptic coupling, recurrent inhibition is sufficient to stabilize fluctuations; for strong coupling, firing rates are stabilized solely by the upper bound imposed on activity.

Finally, we investigate the learning properties of that class of networks, and we analyze the limitations introduced by noise.

II-50. Functional requirements for homeostatic inhibitory plasticity in recurrent networks

Owen Mackwood[1,2] OWEN.MACKWOOD@BCCN-BERLIN.DE
Henning Sprekeler[1] H.SPREKELER@TU-BERLIN.DE

[1]Technische Universitaet Berlin
[2]Bernstein Center for Computational Neuroscience

Neuronal networks in the brain display a wide variety of plasticity mechanisms, some of which have been interpreted as homeostatic regulators of neural activity. This includes synaptic mechanisms such as activity-dependent scaling and GABAergic plasticity, as well as plasticity of glutamatergic synapses onto inhibitory interneurons. While the first two have been adopted in computational models, the last has attracted less attention. Here, we investigate how excitatory synapses onto interneurons can function as a homeostatic regulator of excitatory activity. To this end, we simulate locally connected recurrent networks with a variety of excitatory-to-inhibitory plasticity rules, all of which have a fixed point where the presynaptic rate for each cell equals a homeostatic target. We show that for a Hebbian rule, excitatory cells compete to control inhibition, driving most of the network into a quiescent state. The competition arises from the fanout in recurrent connectivity that produces de-localized inhibition when local inhibition is asked for. This problem is only overcome when a) the post-synaptic rate dependence in the rule is removed, b) the presynaptic rate is replaced by a local population average, and c) the region over which it is averaged is sufficiently large compared to the connectivity range. Such a rule is conceptualized as plasticity mediated by a diffusing agent released by active excitatory cells. The diffusion-mediated rule successfully controls the local population rate, while leaving local variability uncompensated, in contrast to homeostatic plasticity on GABAergic synapses which compensates each excitatory cell in a fine-grained manner. These results predict that manipulating the activity of small groups of principal cells should produce homeostatic changes to glutamatergic synapses onto inhibitory interneurons in neighbouring principle cells. Furthermore, any plasticity observed at these synapses that is not mediated by a diffusive agent is unlikely to homeostatically regulate excitatory firing rates. [Supported by BMBF (FKZ 01GQ1201)]

II-51. Modeling the dynamics of large-scale cortical networks with laminar structure

Jorge Mejias[1] JORGE.F.MEJIAS@GMAIL.COM
John D Murray[2] JOHN.MURRAY@YALE.EDU
Henry Kennedy[3] HENRY.KENNEDY@INSERM.FR
Xiao-Jing Wang[1] XJWANG@NYU.EDU

[1]New York University
[2]Yale University
[3]Universite de Lyon

Interactions between visual cortical occur in both feedforward and feedback directions along the visual hierarchy. It is known that these inter-areal interactions differ in their anatomical properties: feedforward and feedback projections tend to originate from supragranular and infragranular layers, respectively. Physiologically, several recent works have identified a clear spectral profile as well: feedforward (bottom up) interactions seem to be associated with oscillations in the gamma band (30-70Hz), while feedback (top down) interactions relate to lower frequencies, in the high alpha or low beta range (8-20 Hz). In this work, we developed an extended large-scale computational model of monkey cortex (Chaudhuri et al., Neuron, 2015) endowed with a laminar structure of cortical areas, to

investigate the dynamical mechanism underlying frequency-specific interactions in the visual system. The model spans multiple scales, and at each level (local circuit, laminar network, inter-areal interactions, and large-scale cortical network) it is anatomically constrained and then tested against electrophysiological observations, which provides novel and valuable insight about their circuit mechanisms. At the large-scale network level, the model is built upon state-of-the-art anatomical connectivity data from the macaque brain, which provides information on both the strength of the directed inter-areal connections and the laminar origin of such connections. This allows the model to quantitatively predict the emergent frequency-dependent functional connectivity and its relationship to the underlying structural connectivity. It provides a mechanistic explanation for the emergence of functional hierarchies among visual cortical areas, under the condition that feedback projections predominantly target deep layers, a model prediction that can be tested experimentally. Our work highlights the importance of multi-scale approaches, with anatomical and physiological constrains at each step, in the construction of large-scale brain models.

II-52. Pattern generation in simple inhibition-dominated networks

Katherine Morrison[1] KATHERINE.MORRISON@UNCO.EDU
Anda Degeratu[2] ANDA.DEGERATU@MATH.UNI-FREIBURG.DE
Vladimir Itskov[3] VLADIMIR.ITSKOV@MATH.PSU.EDU
Carina Curto[3] CCURTO@PSU.EDU

[1]University of Northern Colorado
[2]University of Freiburg
[3]Pennsylvania State University

Many networks in the brain exhibit internally-generated patterns of activity – that is, emergent dynamics that are shaped by intrinsic properties of the network rather than inherited from an external fluctuating input. For example, spontaneous sequences have been observed in both cortical and hippocampal networks (Luczak et al., 2007; Stark et al., 2015), and patterned motor activity arises from central pattern generator circuits (CPGs) (Marder and Bucher, 2001). A common feature of all these networks is an abundance of inhibition, which has led to the idea that cortical circuits may function similarly to CPGs (Yuste et al., 2005). The mechanisms underlying pattern generation, however, remain unclear. Possible explanations range from single cells, such as “pacemaker" or other intrinsically rhythmic neurons (White et al., 1998; Whittington et al., 2000), to network-level properties like the structure of connectivity. In this work, we narrow our focus to the role of connectivity alone by studying a model with simple perceptron-like neurons but complex connectivity. Specifically, we introduce the Combinatorial Threshold Linear Network (CTLN) model, a simple inhibition-dominated network model whose dynamics are controlled solely by the structure of an underlying directed graph. By varying only the graph, we observe a rich variety of emergent patterns including: multistability, fast sequences, slower "cell assembly" sequences, and complex rhythms. These patterns are reminiscent of population activity in cortex, hippocampus, and various CPGs. We also prove new mathematical results on the CTLN model that allow us to predict many features of the dynamics by examining properties of the underlying graph. We show examples illustrating how these theorems enable us to engineer complex networks with prescribed dynamic patterns.

II-53. The structure of correlated variability in balanced cortical circuits

Robert Rosenbaum[1] ROBERT.ROSENBAUM@ND.EDU
Matthew Smith[2] SMITHMA@PITT.EDU
Jonathan Rubin[2] JONRUBIN@PITT.EDU
Brent Doiron[2] BDOIRON@PITT.EDU

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

Correlated variability in cortical circuits impacts population coding and provides a window to synaptic connectivity structure. Linking circuit structure to correlation structure is therefore important for studying sensory coding and for interpreting cortical recordings. We combine computer simulations, mathematical analysis and in vivo recordings to study the link between spatial correlation structure and the spatial structure of synaptic connectivity in neural circuits with balanced excitation and inhibition. We show that when recurrent, intra-laminar projections are narrower than external feedforward projections, spiking is weakly correlated between neurons at all distances. Broader recurrent projections, however, produce a distinctive spatial correlation structure in which nearby neuron pairs are positively correlated, pairs at intermediate distances are negatively correlated and the average correlation between pairs is extremely small. This correlation structure is revealed in recordings from primate primary visual cortex, but only after subtracting low-dimensional latent variability. Our findings provide critical intuition for how synaptic connectivity structure shapes correlations structure in cortex.

II-54. Balance out of control: robust stabilization of recurrent circuits via inhibitory plasticity

Guillaume Hennequin[1] GJEH2@CAM.AC.UK
Tim P. Vogels[2] TIM.VOGELS@CNCB.OX.AC.UK

[1]University of Cambridge
[2]University of Oxford

During limb movements, neurons of the primary motor cortex (M1) exhibit large, temporally complex, and short-lived activity transients. To account for these, we have recently proposed a new type of balanced architectures, in which strong and intricate recurrent excitation (E) is stabilized by detailed feedback inhibition (I) [1]. While we were able to construct such networks using algorithms from control theory and optimization, exactly how stabilizing feedback can be learned through realistic forms of inhibitory synaptic plasticity (ISP) remains unclear. Here we resolved this question in balanced networks with linear(ized) stochastic rate dynamics. First, we numerically confirmed a simple intuition: that Hebbian ISP [2], known to establish a detailed E/I balance in feedforward circuits, successfully stabilizes continuous growth in recurrent excitation. However, for strong enough excitation, ISP eventually fails. To understand – and ultimately circumvent – this limitation, we formulated optimal ISP as H2-norm minimization, a well-known control-theoretic approach to robust stability optimization. Through direct analytical comparisons with the optimal solution, we could show that Hebbian ISP always increases the degree of network stability, but becomes progressively less effective as excitatory connectivity builds up. Indeed, in that regime, knowledge of pre/post-synaptic activity correlations (the essence of Hebbian ISP) must be complemented by information about the network's input sensitivity ("which inputs evoke the strongest collective responses?"). Surprisingly, this information can be collected locally via a spontaneous diffusion process running backwards through the synapses. We thus augmented Hebbian ISP with retrograde messengers, and obtained greatly enhanced stabilization performance over a much broader range of connectivity strengths. Our results provide the foundations for understanding how the brain reaches global objectives (here, stability) through local synaptic modifications, and suggest a new functional role for retrograde synaptic transmission. [1] Hennequin et al., Neuron (2014) [2] Vogels et al., Science (2011)

II-55. Discovering spatiotemporal structure in epileptic activity through discrete latent variable modeling

John Szymanski[1] JRS2236@COLUMBIA.EDU
Daniel McNamee[2] D.MCNAMEE@ENG.CAM.AC.UK
Michael Wenzel MW2946@COLUMBIA.EDU
Rafael Yuste RMY5@COLUMBIA.EDU

[1]Columbia University
[2]University of Cambridge

The generation and propagation of epileptic neural activity involves complex patterns of multineuronal activity, and the spatiotemporal dynamics of this process are largely uncharacterized in intact brains. In this analysis of calcium imaging of acute pharmacologically induced epileptic seizures in the intact mouse brain, we use an unsupervised approach to discover a characteristic sequence of discrete states that underlies the propagation of seizures through a field of view comprising hundreds of neurons in mouse cortex. Previous analyses of multineuronal cortical responses to sensory stimuli have revealed that cortical population activity can progress through sequences of discrete states, and that transitions between these states occur at stochastic intervals from stimulus onset and from each other. Exploratory analysis of dimensionally reduced population activity data from our seizure model revealed similarly abrupt transitions between states within seizures, as well a diversity of roles of parvalbumin (PV) interneurons that depends on their local density. To better characterize these states, we fit a hidden Markov model (HMM) to the data, identifying a sequence of states that was largely stable between seizures and also evolved as seizures grew in magnitude over the course of the experiment. The sequence includes a pre-ictal state, a state surrounding the initial peak of cellular activity, and one or two states during the end of the seizure. These temporal state sequences served as the basis for further spatial analysis of network activity within seizures, revealing spatially structured changes in population coupling and pairwise correlations between cells underlying the propagation of seizures.

II-56. Irregular spiking and decline in response variability emerge in recurrent networks at criticality

Yahya Karimipanah YAHYAKP@GMAIL.COM
Zhengyu Ma ZHENGYUMA@WUSTL.EDU
Ralf Wessel RW@PHYSICS.WUSTL.EDU

Washington University in St. Louis

Irregular spiking is a widespread phenomenon in neuronal activities in vivo. In addition, the firing rate variability decreases after the onset of external stimuli. Although these are known as two universal features of cortical activity, a universal explanation underlying such phenomena is still missing. Independently, the collective cortical activity is coordinated and exhibits scale free dynamics. This scale invariance is known to be a signature of a critical point; a point right at the transition between a phase of short-lasting activity and of a long-lasting chaotic regime. In parallel, it has been shown, both theoretically and experimentally, that such criticality brings about certain functional advantages such as maximum information transfer and dynamic range. However, despite the strong evidence for criticality hypothesis, it is still very little known on how it can be leveraged to facilitate neural coding. As the decline in response variability is regarded as an essential mechanism to enhance coding efficiency, we asked whether this feature could be addressed within the context of criticality. Using a simple binary probabilistic model, we show that irregular spiking and decline in response variability, both arise as emergent properties of a recurrent network poised at criticality. Our results provide us with a unified explanation for the ubiquity of these two features, without a need to exploit any further mechanism. We believe this finding is even capable of making a paradigm shift in studying criticality, as not only does it make a bridge between critical dynamics and neural coding, but it also provides us with new measures that can be used to test the criticality hypothesis itself.

II-57. Exogenous modulation of ongoing oscillations in spiking neural networks

Jeremie Lefebvre[1,2] JEREMIE.LEFEBVRE@HOTMAIL.COM
Axel Hutt[3] AXEL.HUTT@INRIA.FR
Christoph Herrmann[4] CHRISTOPH.HERRMANN@UNI-OLDENBURG.DE
Micah Murray[5] MICAH.MURRAY@CHUV.CH

[1]Toronto Western Research Institute
[2]University Health Network
[3]INRIA Nancy
[4]Carl von Ossietzky University
[5]Centre Hospitalier Universitaire Vaudois

Rhythmic neural activity is believed to play a central role in neural computation. Oscillatory brain activity has been associated with myriad functions such as homeostasis, attention, and cognition as well as neurological and psychiatric disorders, including Parkinson's disease, schizophrenia, and depression. Numerous studies have shown that that non-invasive stimulation, such as Repetitive Transcranial Magnetic Stimulation (rTMS) and Transcranial Alternating Direct Current Stimulation (TACS), provide the means of modulating large-scale oscillatory brain dynamics by perturbing and/or entraining both resting state and task activity. These stimulation-induced perturbations of neural oscillations have been shown to alter cognitive performance and perception, effects that are further known to depend on brain state prior and during stimulation. Yet, the surge of interest in these approaches is compromised by the existence of complex interference patterns between exogenous and endogenous cyclic dynamics. Indeed, the intrinsic non-linearity of neural systems suggests that external fluctuations shape ongoing rhythmic neural activity in intricate ways, beyond what is expected by entrainment and resonance alone. To explore these important questions, we used a generic computational model of excitatory and inhibitory spiking neurons, to explore various combinations of stimulation frequencies and amplitudes, revealing those in which resonance and/or rhythmic entrainment (i.e. Arnold tongues) can be observed. Furthermore, building on recent advances in non-linear dynamics, we report novel regimes of oscillatory acceleration that characterize network responses to high-frequency exogenous forcing. Our results provide new computational perspectives on the response of synchronous non-linear neural systems, in which a plurality of intrinsic mechanisms can be recruited to modulate emergent oscillatory states.

II-58. An onset scenario for febrile seizures: Enhanced neural synchronization at a critical temperature

Janina Hesse[1,2] JANINA.HESSE@BCCN-BERLIN.DE
Nikolaus Maier[3] NIKOLAUS.MAIER@CHARITE.DE
Dietmar Schmitz[3] DIETMAR.SCHMITZ@CHARITE.DE
Jan-Hendrik Schleimer[1] JH.SCHLEIMER@HU-BERLIN.DE
Susanne Schreiber[1] S.SCHREIBER@HU-BERLIN.DE

[1]Humboldt-Universitaet zu Berlin
[2]BCCN Berlin
[3]Charite-Universitatsmedizin Berlin

About 2 to 5% of young children experience seizures during high fever, so-called febrile seizures. While changes in pH seem to play a role in the induction of these convulsions, our modeling results suggest that the increase in temperature preceding a seizure could also directly contribute to seizure induction. We investigated the influence of temperature on type I model neurons and identified a critical temperature range where even mild temperature increases ($\approx 1\hat{\ }\circ C$) alter single-cell dynamics such that network synchronization increases strongly. This critical temperature range lies close to a codimension-two bifurcation: the saddle-node-loop bifurcation (SNL). At the SNL point, the dynamics of firing onset switch from a saddle-node bifurcation (typical for type I neurons) to a homoclinic

orbit bifurcation and, importantly, the phase response curve (PRC) changes from the canonical, symmetric shape to an asymmetric one, strongly favoring synchronization. Such a drastic change in synchronization caused by small increases in temperature constitutes a possible mechanism for the induction of febrile seizures. Our model analysis provides several predictions including the temperature dependence of PRCs and spike shapes, which we test experimentally in recordings of hippocampal pyramidal cells. We further analyze whether the proposed mechanism can explain genetic predispositions for strong febrile seizures: prevalence of febrile seizures should increase if a genetic mutation brings a neuron closer to the SNL point. Exploring established models of febrile-seizure-related ion-channel mutations, we find that this is indeed the case. Mutations lower the critical temperature of the SNL point, where synchronization and excitability explode. We conclude that the temperature-induced changes in neuronal dynamics could indeed contribute to the induction of febrile seizures. Because also other parameters can shift neuronal dynamics towards the SNL point, the relevance of the proposed mechanism could extend to other systems marked by a sudden increase in synchronization, such as epileptic seizures.

II-59. Dynamic information routing at the edge of synchrony

Agostina Palmigiano[1] AGOS@NLD.DS.MPG.DE
Theo Geisel[2] GEISEL@NLD.DS.MPG.DE
Fred Wolf[2] FRED@NLD.DS.MPG.DE
Demian Battaglia[3] DEMIAN.BATTAGLIA@UNIV-AMU.FR

[1]Max Planck Institute for Dynamics
[2]Max Planck Institute for Dynamics and Self-Organization
[3]Aix-Marseille Universite

Behavior and cognition require precisely targeted switching between communication pathways to selectively process, from a variety of competing stimuli, those relevant for the current task. Currently proposed models of neural routing largely rely on the separation of a local processing circuit and an embedded feedforward sender network. Here, working within the communication-though-coherence paradigm, we show that in models of networks coupled by long range excitation and undergoing transient synchrony episodes in the gamma band, the same mechanisms generating neuronal oscillations give rise to the transient phase locking used for communication. More specifically, we show that in networks of conductance based neurons with highly heterogeneous parameters a robust regime characterized by transient synchrony emerges. When local circuits are interconnected by long range excitation, this short-lived events self organize to give rise to frequency tracking and transient phase locking, as seen in vivo. During these temporally restricted episodes of coordinated phase-locked activity, information is regulated by the level of synchronization and follows the hierarchy imposed by the phase relation. Indicators of information transfer such as transfer entropy and mutual information reveal uni directional information transmission from a phase leading population to the lagging ones. A wide range of complex random input signals to a source area can be selectively gated and read out in a distant target area depending on the pattern of phase differences between the circuits activity. We study simple motifs of two or three populations, this later case emulating experimental paradigms in which selective attention is directed to one out of two simultaneously presented stimuli. We further show that weak external biases can naturally be used to induce switching between different routing configurations by altering the phase relation and revealing a vast repertoire of functional motifs.

II-60. Functional dichotomy of spatially modulated entorhinal island and ocean cells

Chen Sun[1] CSUN@MIT.EDU
Takashi Kitamura[1] KITAMURA@MIT.EDU
Susumu Tonegawa[1,2] TONEGAWA@MIT.EDU

[1]Massachusetts Institute of Technology
[2]RIKEN Brain Science Institute

The importance of entorhinal-hippocampal circuits in the mammalian brain for an animal's spatial and episodic experience is known, but the neural basis for these different spatial computations is unclear. Medial entorhinal cortex layer II contains Island and Ocean cells that project via separate pathways into hippocampus. To date, few studies have investigated the relationship between anatomy and spatial coding in entorhinal cortex, thus there is little known about how these separate populations of Ocean cells and Island cells differ in their coding of space, if at all. We performed cell type specific Ca2+ imaging in freely exploring mice using cellular markers and a miniature head-mounted fluorescence microscope, the first time in entorhinal cortex. We found that both DG projecting Oceans and CA1 projecting Islands contain grid cells and other spatially modulated cells. However, Ocean cells, unlike Island cells, rapidly forms a distinct representation of a novel context and drives context-specific fear memory. On the other hand, Island cells are significantly more speed-modulated than Ocean cells. Together, the spatially modulated cells found in these two excitatory medial EC layer II inputs to the hippocampus may subserve different downstream spatial functions: Ocean cells may facilitate contextual representation in downstream circuits whereas speed modulated Island cells may contribute more to spatial path integration.

II-61. Novel probabilistic computations unifies diverse grid cell responses

Allen Cheung A.CHEUNG@UQ.EDU.AU

The University of Queensland

Mounting evidence shows mammalian brains are probabilistic computers [1-3], but the specific cells involved remain elusive. Parallel research suggests that grid cells of the mammalian hippocampal formation are fundamental to spatial cognition but their diverse response properties still defy explanation. No plausible mechanism exists which explains stable grids in darkness for 20 minutes or longer, despite being one of the first results ever published on grid cells [4]. Similarly, no current explanation can tie together grid fragmentation [5] and grid rescaling [6, 7], which show very different forms of flexibility in grid responses when the environment is varied. Other properties such as attractor dynamics [8] and grid anisotropy [9] seem to be at odds with one another. Modelling efforts have largely ignored the breadth of response patterns, while frustratingly also fail to account for the disastrous effects of sensory noise during spatial learning and recall, especially in darkness [10-12]. Here, electrophysiological and theoretical evidence is presented which shows that grid cell responses are accurately predicted by probabilistic learning computations. Diverse response properties of probabilistic grid cells are statistically indistinguishable from rat grid cells across key manipulations. A simple coherent set of probabilistic computations explains stable grid fields in darkness, partial grid rescaling in resized arenas, low-dimensional attractor grid cell dynamics, and grid fragmentation in hairpin mazes. The same computations also reconcile oscillatory dynamics at the single cell level [13-15] with attractor dynamics at the cell ensemble level [8]. Additionally, a clear functional role for boundary cells [16-18] is proposed for spatial learning. These findings provide a parsimonious and unified explanation of grid cell function, and implicate grid cells as an accessible neuronal population involved in probabilistic computations.

II-62. A model for spatially periodic firing based on interacting excitatory and inhibitory plasticity

Simon Weber WEBER@TU-BERLIN.DE
Henning Sprekeler H.SPREKELER@TU-BERLIN.DE

Technische Universitaet Berlin

Neurons in the hippocampal formation exhibit a variety of spatially tuned firing patterns. The mechanisms by which these different patterns emerge are not fully resolved, although competing computational models exist for several of them. Here we present a new model that can generate all observed spatial firing patterns by a single mechanism. The model consists of a feedforward network with a single output neuron. Its essential ingredients are i) spatially modulated excitatory and inhibitory inputs and ii) interacting excitatory and inhibitory Hebbian plasticity. The inhibitory plasticity homeostatically controls the output firing rate by balancing excitation and inhibition. We show in simulations and by a mathematical analysis that the output neuron rapidly develops periodic firing patterns along a spatial dimension if inhibitory inputs are more broadly tuned than excitatory inputs along this dimension. More generally, depending on the relative spatial auto-correlation length of the excitatory and inhibitory inputs, the model exhibits firing patterns that are similar to those of place cells, grid cells or band cells. For inputs with combined spatial and head direction tuning, the same mechanism leads to output firing patterns reminiscent of head direction cells and conjunctive cells (neurons that fire like grid cells in space but only at a particular head direction). A linear stability analysis of the homogeneous steady state accurately predicts the spatial periodicity obtained from simulations. The model combines the robust pattern formation of attractor models, with the spatial (rather than neural) structure formation of models based on synaptic plasticity. In contrast to attractor models, our model predicts that the spatial scale of grid cells should be robust to global modifications in inhibitory synaptic strength. (Funded by the German Federal Ministry for Education and Research, FKZ 01GQ1201)

II-63. Longitudinal visualization of spontaneous activity across early cortical development

Gordon Smith[1] GORDON.SMITH@MPFI.ORG
Bettina Hein[2] HEIN@FIAS.UNI-FRANKFURT.DE
David Whitney[1] DAVID.WHITNEY@MPFI.ORG
Philipp Huelsdunk[2] HUELSDUNK@FIAS.UNI-FRANKFURT.DE
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

Numerous models of cortical map formation have attributed a critical role to activity dependent interactions within developing cortical circuits. It is plausible that correlated activity in the early cortex could establish the networks of spatially distributed but functionally co-tuned neurons present in the mature cortex, but there is so far no experimental evidence supporting this hypothesis. Here, we perform the first longitudinal study of correlation structures in ferret visual cortex across development, starting at an age when layer 2/3 cells are starting to receive feed-forward input and extending to a stage when circuits reach full maturity. By measuring the correlation structure of both spontaneous and stimulus-evoked cortical activity, we are able to relate correlations to tuning properties as they emerge across time. In order to reveal the correlation structure across both cellular and columnar scales we perform two-photon and wide-field epifluorescence imaging. We first determined the influence of anesthesia on spontaneous activity using a novel head-restrained preparation, finding that spontaneous activity in the awake cortex is more pronounced than under anesthesia. However the spatial correlation structure is highly consistent across states, suggesting that this structure reflects features of a stable circuit architecture. Furthermore, the correlation structure at the cellular level is highly similar to that obtained with wide-field imaging, suggesting a cellular basis for large scale patterns of correlated activity. Finally, we show that spontaneous correlations are higher between locations which develop similar orientation preferences in the mature cortex and are lower for locations developing orthogonal tuning. Notably, this structure is already apparent in the highly immature cortex as early as 10 days prior to eye opening, and is not evident from stimulus-evoked cortical activity through closed-eyelids. Together, these results suggest that early spontaneous activity is suitable to fulfill the instructive role proposed by models of cortical development.

II-64. Populations of retinal ganglion cells are robustly in a ‘low temperature’ state

Mark Ioffe MIOFFE@PRINCETON.EDU
Gasper Tkacik GTKACIK@IST.AC.AT
William Bialek WBIALEK@PRINCETON.EDU
Michael Berry BERRY@PRINCETON.EDU

Princeton University

Previous results have suggested that the maximum entropy models inferred from neural data are near a phase transition [Tkacik 2015]. Specifically, models that constrain time-averaged statistics (such as firing rates and pairwise correlations), while maximizing the entropy, are mathematically identical to the Ising model. This model is one of the simplest models that can have a phase transition, marked by a discontinuity in one of the derivatives of the free energy, such as the specific heat. By fitting an Ising model to the data, and then introducing a parameter equivalent to the temperature, one can explore whether there is a phase transition in the vicinity of the operating point of the system (T=1). Here, we establish that this phase transition is robustly present in multiple experimental conditions with statistically significant differences in neural population statistics. We then manipulate the measured covariance matrix to test which properties of the correlation structure give rise to these properties. By scaling the overall covariance matrix, we find a step-like shift in the specific heat at relatively low covariance level, where the nature of the model neural population jumps from essentially independent to structured. This indicates that real networks are distant, in this parameter space, from independence. By setting some pairwise correlations to zero, we find also that these effects depend crucially on a web of many, weak correlations. Thus, populations of retinal ganglion cells have a web of correlation that is well within the range needed to put the system in a structured, ‘low temperature’ state. Because the requirements for this state are fairly generic, neural populations in many other brain areas may also share this structure. We are actively exploring the hypothesis that in this low temperature state, sets of activity patterns naturally cluster into population codewords.

II-65. A biophysical model of evoked local field potential and EEG incorporating layer-dependent synaptic activities

Jingjing Luo[1] JINGJING.LUO@READING.AC.UK
Michael Bruyns-Haylett[1] M.BRUYNS-HAYLETT@READING.AC.UK
Jason Berwick[2] J.BERWICK@SHEFFIELD.AC.UK
Aneurin Kennerley[2] A.J.KENNERLEY@SHEF.AC.UK
Luke Boorman[2] L.BOORMAN@SHEFFIELD.AC.UK
Samuel Harris[2] SAM.HARRIS@SHEFFIELD.AC.UK
Daniel Coca[2] D.COCA@SHEFFIELD.AC.UK
Stephen A Billings[2] S.BILLINGS@SHEFFIELD.AC.UK
Ying Zheng[1] YING.ZHENG@READING.AC.UK

[1]University of Reading
[2]The University of Sheffield

We present a dynamic model of the evoked local field potential (LFP) and EEG at the somatosensory cortex of anesthetised rats. The model is a neural population model based on the finding that synaptic excitation and inhibition co-vary with inhibition lagging excitation, and that LFP and EEG reflect mainly synaptic activities of local pyramidal neural populations [Zheng et al 2012]. The model thus decomposes LFP and EEG into excitatory and inhibitory synaptic activities of neural populations within the local cortical column. In order to understand the interaction between neural excitation and inhibition at the population level, we pharmacologically manipulated inhibitory synaptic activity by applying GABAA receptor antagonist bicuculline in the barrel cortex. It showed that P1 in EEG and LFP of supergranular layer (LFP1) was not affected by the intervention while N1 in LFP of deeper layers increased and prolonged significantly [Bruyns-Haylett et al 2015] (Fig. 1). This implies that P1 in EEG and LFP1 is independent of inhibition; instead, it is a combination of current source from granular synaptic excitation and current sink from supergranular synaptic excitation [J. J. Luo et al 2014]. The model incorporates layer-dependent synaptic activities, thus decomposing EEG and LFP into neural excitation and inhibition of different cortical layers. The current model, integrating neural activities at both granular and supergranular layers, is capable of fitting evoked EEG and LFP throughout the time course of bicuculline intervention. The estimated inhibitory components in both granular and supergranular layers were minimum when the bicuculline effect was maximum, and returned as the bicuculline effect weakened. However, estimated excitatory components remained throughout (Fig. 2). This study demonstrates the decomposition of P1 of EEG into synaptic excitation of granular and supergranular layers thus has important implications on understanding and interpreting the non-invasive EEG signals.

II-66. Representation of space and choice in mouse parietal cortex during virtual navigation

Michael Krumin M.KRUMIN@UCL.AC.UK
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

Posterior parietal cortex (PPC) in primates is believed to be involved in cognitive operations such as coordinate transformations and decision making. In rat PPC, some experiments revealed decision signals, but others revealed signals related to navigation. Does rodent PPC carry signals related to decision or to navigation, or a mixture of these and other signals? To address these questions we used 2-photon imaging to record populations of neurons in PPC of mice that performed a contrast discrimination task in a virtual T-maze. The mice traversed a corridor and reported the position of a grating by either turning right or turning left at the end of it. The activity of PPC neurons was strongly modulated by the mouse's virtual position along the main corridor (z) and virtual head direction (θ). Many of the cells had localized 'position-heading fields' in these coordinates, and this $z - \theta$ model gave a prediction of each neuron's activity with explained variance of up to 89% ($21 \pm 16\%$ for all active cells). These position-heading fields could occur at any location in $z - \theta$ space, and indeed the neurons recorded in a single session typically tiled the T-maze. In some cells, the strength of the position-heading field was additionally modulated by the contrast and position of the visual grating (on either wall). However, the final choice made by the animal did not seem to affect the responses of PPC cells beyond what could be simply predicted from their position-heading field. These results suggest that mouse PPC can play a key role in visually guided behavior. Although this region might play other roles in other task conditions, the large fraction of variance that could be predicted by spatial variables alone suggests that our mice have adopted a strategy in which PPC is involved in spatial computations, rather than choice representation.

II-67. Distributed neural prediction of prey motion in amphibians

William Mowrey[1,2] MOWREYW@JANELIA.HHMI.ORG
Anthony Leonardo[1,2] LEONARDOA@JANELIA.HHMI.ORG

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

Rapid ballistic movements are used to track objects, capture prey, and escape from predators. These behaviors are often highly accurate, despite significant sensory and motor delays. Recent work has shown that amphibian tongue projections compensate for phototransduction and head movement delays by extrapolating prey motion (Borghuis and Leonardo, in press). This extrapolation may be accomplished in the retina, as activity in amphibian fast-OFF retinal ganglion cells (RGCs) has been shown to predict object motion (Berry et al., 1999). An accurate tracking estimate can be extracted from fast-OFF population activity by a population vector average (PVA), though this tracking estimate is imperfect and features characteristic size- and speed-dependent errors (Leonardo and Meister, 2013). This stands in contrast to optimal linear decoding of retinal activity, which tracks moving objects with remarkable accuracy (Marre et al., 2015). A characteristic error of fast-OFF/PVA tracking is a shape-dependent bias: fast-OFF cells are strongly driven by the leading edge of dark prey-like objects, such that population activity becomes biased ahead of center for objects extended in the direction of motion. We tested for this bias by presenting artificial prey of different shapes to amphibians on a touchscreen, and found good agreement between the fast-OFF/PVA model and tongue projection accuracy. This supports the notion that PVA decoding of a specific RGC subtype underlies prey tracking in amphibians, rather than optimal linear decoding across all RGC types. We further tested predictions of speed-dependent error in prey capture. In contrast to the fast-OFF/PVA model, toads projected with similar accuracy to prey over a 4-fold range of speeds. Together, our data support a PVA-like readout of activity in the early visual system for tongue projection behavior, but suggest that additional predictive computation in the brain corrects for speed-dependent bias observed in retinal fast-OFF population activity.

II-68. Beyond kinematic and EMG tuning: object-related activity in M1 single neurons and populations during grasping

Rex N Tien[1] RNT9@PITT.EDU
Sagi Perel[2] SAGI.PEREL@GMAIL.COM
Andrew Schwartz[1] ABS21@PITT.EDU

[1]University of Pittsburgh
[2]Google

Neurons in primate primary motor cortex (M1) display modulated activity during the grasping and manipulation of objects. Previous studies have sought to explain these modulations by regressing neural activity against movement features such as kinematics (endpoint position and velocity and arm, hand and finger joint angles and velocities) or electromyographic (EMG) activity. In this study, we present evidence that single neurons and populations in M1 encode information about the specific object being grasped beyond encoding the movements and muscle activities required to execute the grasp. Two rhesus macaques grasped 6 different objects while we simultaneously recorded arm and hand kinematics and EMGs and up to 9 neurons in M1. The variation in M1 firing rates (FRs) associated with movement features was removed by regressing FRs against time-lagged kinematics and EMGs. The presence of object-related structure in the residuals from these regressions would suggest that the FRs encoded the identity grasped object beyond encoding the specific kinematics and EMG activities needed to grasp that object. Indeed, ANOVAs with factor object revealed that nearly all of the neurons encoded object identity in their residual FRs (RFRs). Treating the neural data from several sessions as a pseudo-population revealed yet more about the nature of this object-related activity. Classifiers built from RFRs could identify which object was grasped at levels significantly above chance. Additionally, pairs of objects that elicited the most similar sets of kinematics and EMGs elicited the most disparate sets of RFRs, suggesting that this activity may serve

to separate the neural representations of movements made to similar objects. Evaluation of different regression models showed that this object encoding may be complex; the best fits were for models in which neural tuning parameters to individual movement features were allowed to change depending on the identity of the object.

II-69. Parallel visual search in the archer fish

Mor Ben-Tov[1] MORBENTOV@GMAIL.COM
Opher Donchin[1] DONCHIN@BGU.AC.IL
Ohad Ben-Shahar[1] BEN-SHAHAR@CS.BGU.AC.IL
Ronen Segev[1,2] RONENSGV@BGU.AC.IL

[1]Ben-Gurion University of the Negev
[2]University of California, Berkeley

Archer fish are known for their remarkable ability to shoot down insects situated on vegetation above water. Motivated by the capacity of the archer fish to detect such small targets rapidly and accurately, we explored how they perform visual search. Specifically we were interested in exploring the parallel mode of visual search, aka pop-out in visual search. Pop-out enables fast detection times which are independent of the number of distracting objects present. Due to the importance of such mode of visual search, an extensive research was conducted to understand which visual attributes elicit pop-out, and how the brain performs visual search. Behaviorally, pop-out has been observed until recently only in mammals, where its neural correlates are found in the primary visual cortex in contextually modulated neurons that encode aspects of saliency. In a recent study, we showed that the archer fish also exhibits this important visual search mode when presented with moving targets. In addition, using single cell recordings, we found contextually modulated neurons in the optic tectum, the main visual area in the archer fish brain. Motivated by these results, we now explored which visual attributes elicit pop-out in the archer fish. In mammals, it is known that color, size, orientation and motion are undoubted features that elicit parallel search in visual search. Using a set of behavioral experiments we tested the effect of these features and found they elicit pop-out in the archer fish also. Taken together, our results indicate that there may by universality in the performance and computation of pop-out across vertebrates. In addition, they raise the question of what are the most informative attributes that can elicit pop out visual search in a natural environment.

II-70. Attentional modulation of interareal and interlaminar connectivity in macaque V1 and V4

Michael Boyd[1] M.BOYD@NCL.AC.UK
Jochem van Kempen[2] JOCHEM.VAN-KEMPEN@NCL.AC.UK
Michael Savage[2] MICHAEL.SAVAGE1@NCL.AC.UK
Miguel Dasilva[3] MIGUEL.DASILVA@MANCHESTER.AC.UK
Alexander Thiele[2] ALEX.THIELE@NCL.AC.UK

[1]Institute of Neuroscience
[2]Newcastle University
[3]The University of Manchester

Attention prioritises processing of task relevant information in the presence of irrelevant or distracting stimuli. Attention modulates neuronal spiking activity, local field potentials and coherence between neuronal groups. These analyses have provided insight into the effects of attention within single cortical layers and areas and sometimes between areas. We lack knowledge how attention modulates the interactions and information exchange between groups of neurons in different cortical layers within and between areas. Male rhesus macaques performed a covert visuospatial attention task whilst we recorded simultaneously from V1 and V4 using multichannel laminar electrodes spanning infragranular, granular and supragranular layers within cortical columns. Recording channels

in V1 and V4 had overlapping receptive fields. Using current source density analysis channels were aligned to the granular input layer. We used field-field, spike-field, and spike-spike coherence to investigate the interactions within and between areas. Attention increased spiking activity throughout V1 and V4, with the exception of the V1 granular layers. V1 neurons showed reductions in low frequency spike-spike coherence with attention, consistent with a reduction in noise correlations. Attention modulates both interareal and interlaminar coherence in a layer dependent manner. In V1, low frequency (4-13Hz) coherence was increased with attention between infragranular contacts as well as infragranular- granular and infragranular-supragranular contacts. Field coherence between supragranular layers, as well as between supragranular-granular and supragranular-infragranular in the low gamma range was reduced by attention. Attention increased V1-V4 coherence in beta, gamma and high gamma frequency bands across layers, but a specific low gamma band coherence 'bump' only occurred between V1 (all layers)-V4 (granular) layers. Attention significantly increased V4 (supragranular)-V1 (all layers) coherence in the low frequency range (4-13Hz). These data demonstrate that attention has specific effects on intra-areal and inter-areal communication consistent with predicted feedforward and feedback projection patterns. Supported by the Wellcome Trust

II-71. Synergy in motion trajectory encoding of direction-selective ganglion cells in the salamander retina

Norma Kuhn NORMA.KUEHN@MED.UNI-GOETTINGEN.DE
Tim Gollisch TIM.GOLLISCH@MED.UNI-GOETTINGEN.DE

University Medical Center Gottingen

The detection of motion direction and velocity in visual scenes is important in everyday tasks, e.g., for avoiding cars when crossing the street or for assessing the optic flow when navigating through an environment. We study the encoding of motion in visual scenes on the level of the retinal ganglion cells. There, direction-selective ganglion cells (DSGCs) are known to encode certain directions of a drifting pattern. They provide important information about the optic flow to higher brain areas. In the salamander, we found three subtypes of DSGCs whose preferred directions are separated by $120\hat{}\circ$, similar to the system of On-type DSGCs observed in the mouse. To probe motion encoding with more complex patterns than simple drifting gratings, we projected textures with random trajectories onto the isolated salamander retina and recorded up to 30 DSGCs simultaneously with multielectrode arrays for population analysis. We used a commonly deployed linear decoder to derive stimulus predictions from individual DSGCs and population responses. We estimated the mutual information between stimulus and linear prediction and found that the joint activity of DSGC populations with different preferred directions conveys more information than the sum of its individual contributions. This synergy is caused by the fact that motion encoding by individual DSGCs for complex stimuli is strongly confounded by non-motion-direction-related activity, leading to a non-monotonic nonlinearity in the motion encoding. The concerted firing of DSGCs with different preferred directions can counteract the ambiguity of this non-monotonic stimulus dependence. Our findings suggest that in complex visual scenes, DSGCs do not exclusively encode motion direction but are challenged by a vast spectrum of motion velocities and spatial contrasts. This leads to a strongly nonlinear direction encoding whose unfavorable effects for linear decoding are partly compensated by the concerted signalling of DSGCs with different preferred directions, causing synergy.

II-72. Perceptual distortion measured with a gain control model of LGN response

Alexander Berardino[1] AGB313@NYU.EDU
Valero Laparra[1,2] VALERO@NYU.EDU
Johannes Balle[1,3] JOHANNES.BALLE@NYU.EDU
Eero Simoncelli[1,3] EERO.SIMONCELLI@NYU.EDU

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

Models for perceptual image distortion are generally designed to capture properties of human vision, such as the spatial frequency dependence of contrast sensitivity, and the masking effects of superimposed oriented patterns. However, the widely-used Structural Similarity (SSIM) index (Wang et al., 2004), is based on a unique construction devised to disregard changes in local luminance or contrast, while emphasizing changes in local structure. Since these properties are evident in the responses of early visual neurons, we wondered whether a more explicit model of physiological responses might provide a more suitable substrate for constructing a distortion measure. We built a functional model of early spatial visual processing (LGN-GC) that incorporates known properties of retina and LGN. Similar to the model of Mante et al., 2008, the model includes bandpass linear filtering, rectification, and local luminance and contrast gain controls. The distortion between an original and corrupted image is determined by passing each through the model, and measuring the Euclidean distance between the two response vectors. The model parameters (filter sizes and amplitudes) were fit to a database of human perceptual quality judgments (TID2008 - Ponomarenko et al., 2009). We find that the fitted parameters are consistent with measured physiological properties of LGN neurons in the macaque monkey. Moreover, LGN-GC outperforms SSIM, and performs comparably to multi-scale SSIM (MS-SSIM) at predicting perceptual distortions (despite its restriction to a single spatial scale), explaining more of the (cross-validated) variance in the human data (SSIM: 60%, MS-SSIM: 64%, LGN-GC: 67%). Finally, we performed a direct comparison of LGN-GC to SSIM by examining stimuli optimized to differentiate them (known as MAD competition; Wang et al., 2008). This comparison provides evidence that distance as measured by our physiologically-inspired model corresponds more closely with human perception than SSIM.

II-73. Perceptual evaluation of artificial visual recognition systems using geodesics

Olivier J Henaff[1,2] OJH221@NYU.EDU
Robbe LT Goris[2,1] ROBBE.GORIS@NYU.EDU
Eero Simoncelli[2,1] EERO.SIMONCELLI@NYU.EDU

[1]Howard Hughes Medical Institute
[2]New York University

Recognition is a demanding visual task because image-domain renderings of any given object vary widely. Successful recognition requires representations that distinguish different items, while exhibiting invariance (or at least ‘tolerance’) to naturally occurring variations (DiCarlo & Cox 2007). A new class of artificial systems, deep neural networks, performs real-world visual recognition with surprising accuracy. However, the invariances of these representations have not been compared with those of the human visual system. Here, we propose a new synthesis methodology for comparing the invariances of machine representations to those of human vision. Given two reference images (typically, differing by some transformation), we synthesize a sequence of images that follow a minimal-length (geodesic) path in the response space of the representation. We hypothesize that if the human visual system has the same invariances, then this sequence should also represent a minimal-length perceptual path. Candidate representations can thus be compared by assessing which one produces a sequence that is shortest in perceptual terms. We apply this paradigm to the simple test case of a pair of translated images, and find that a current state-of-the-art model for object recognition generates unnatural-looking geodesic sequences, with easily discriminable successive frames. On the other hand, replacing the max-pooling operation in this network with a more physiologically-plausible quadratic pooling leads to smoother and more natural sequences, in which successive frames are difficult to discriminate. This can be quantified more precisely with formal psychophysical measurements, but even informal viewing demonstrates that the latter sequence follows a perceptually 'shorter' path, and thus that the associated model better captures this invariance of human vision. Our method is general, and can be used to study transformations between arbitrary pairs of images, thus offering a thorough and principled means of comparing representations as models of biological vision.

II-74. Rats' proficiency in shape discrimination is accounted by the complexity of their visual strategy

Vladimir Djurdjevic[1] DJURDJEVIC@SISSA.IT
Daniele Bertolini[1] DBERTOLINI@SISSA.IT
Alessio Ansuini[1] ANSUINI@SISSA.IT
Jakob H Macke[2] JAKOB.MACKE@CAESAR.DE
Davide Zoccolan[1] ZOCCOLAN@SISSA.IT

[1]SISSA (International School for Advanced Studies)
[2]Neural Systems Analysis, caesar, Bonn

In recent years, rodents have emerged as promising models for the investigation of visual processing, including high-level functions, such as visual object recognition. Several studies have shown that rats, in particular, are able to recognize visual objects in spite of substantial variation in their appearance, by relying on a multifeatural shape processing strategy. At the same time, the complexity of this strategy has been found to be highly variable across subjects and largely dependent upon the specific shape of the objects the animals had to discriminate. In this study, we aimed at understanding the functional implications of these earlier findings, by training 6 rats to discriminate a 'reference' object from eleven 'distractor' objects that spanned a wide spectrum of image-level similarity with the reference. Rat performance varied widely across subjects, with some animals succeeding only in the simpler comparisons, and some others achieving above chance performances also with distractors that were very similar to the reference. To explain this phenomenon, we presented the rats with random variations of the reference object, requiring them to classify these noisy stimuli as either the reference or a distractor. We then processed rat responses using a classification image method that produced saliency maps showing what shape features were associated with decisions of individual animals. We found that the complexity of the recognition strategy varied substantially across subjects—some rats relied on a few diagnostic features only, while others built rich perceptual templates made of many distinct features. More importantly, strategy complexity correlated well with rat ability to correctly discriminate the harder distractors—the best performing rats were those displaying the richer perceptual strategies. Overall, these findings explain the variable proficiency of rats in complex visual discrimination tasks, and show how rodents can be capable of surprisingly refined processing of shape information.

II-75. A deep convolutional energy model of V4 responses to natural movies

Michael Oliver MICHAELOLIVER@BERKELEY.EDU
Jack Gallant GALLANT@BERKELEY.EDU

University of California, Berkeley

Area V4 is an important intermediate stage of visual processing. However, it has been very difficult to characterize and model the tuning properties of V4 neurons. For example, no current models of V4 neurons can accurately predict responses to natural images. This is in stark contrast to models of V1 and MT, where responses can be predicted well. V4 neurons have large, nonlinear receptive fields, and this makes it difficult to estimate their tuning

properties using conventional methods: modeling V4 amounts to solving a high-dimensional non-linear regression problem with limited data. To effectively attack this problem, we first sought to collect as much data as possible by chronically implanting electrode arrays in area V4 of two macaque monkeys. Neurons were recorded while the awake animals viewed clips of large, full color natural movies. The chronic recordings were stable enough that neurons could often be held for several days. This allowed us to collect responses to hundreds of thousands (up to over 1 million) distinct movie frames, for hundreds of different V4 neurons. We then used several different neural network architectures to fit the data obtained from each neuron. The training signals for each fit neural network were the stimulus movie and the response from one neuron. The most successful neural network architecture that we tested was one that reflected insights from the Adelson-Bergen energy model, the scattering transform and deep convolutional neural networks. We call this the deep convolutional energy model. This model is simple and interpretable, and it predicts V4 responses significantly better than previous models. Deep convolutional energy models fit to V4 neurons approach the prediction performance of the best current models of V1 and MT neurons. Interpretation of the fit models provides important insights about the representation of visual information in area V4.

II-76. Predicting V2 activity from V1 population activity

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

[1]Champalimaud Centre for the Unknown
[2]Carnegie Mellon University
[3]Albert Einstein College of Medicine

While the brain is composed of functionally distinct cortical areas, little is known about the way in which interconnected populations of neurons communicate relevant information between each other. In particular, it is unclear how, on a moment-by-moment basis, response variability propagates from one area to the next. To tackle this question, we simultaneously recorded from populations of neurons in the output layers of primary visual cortex (V1) and in the input layers of area V2 in anesthetized monkeys. We characterized how the trial-to-trial covariability structure of the V1 population influenced activity in the V2 population. We used V1 population activity to predict V2 activity on a moment-by-moment basis and observed that the prediction accuracy was comparable to that found between V1 populations, suggesting that there is a measurable interaction between V1 and V2. Furthermore, we found that not all modes of activity in V1 were equally effective in driving activity in V2, and only a small subset had a measurable influence on the prediction. Finally, we asked whether the most important modes of activity in V1, the ones that explained most of its shared variability, were also the most important in predicting V2. We found that this was not the case, as these had a significantly lower predictive power when compared to the equivalent optimal predictor. These results suggest that the V1-V2 interaction can be described by V2 "reading out" particular modes of the V1 population activity.

II-77. Predictive computations in the primary visual cortex

Jan Homann JHOMANN@PRINCETON.EDU
David W Tank DWTANK@PRINCETON.EDU
Michael Berry BERRY@PRINCETON.EDU

Princeton University

Predictions about the future are important for an animal in order to interact with its environment. Therefore, predictive computation might be a core operation carried out by neocortical microcircuits. We explored whether the primary visual cortex can perform such computations by presenting repeated temporal sequences of static images with occasional unpredictable disruptions. Simultaneous recordings of 150-250 neurons were performed using two-photon Ca++ imaging of layer 2/3 neurons labeled with GCaMP6f in awake mice, who were head-fixed but free to run on a styrofoam ball. In our visual stimuli, each spatial frame consisted of either an oriented grating or a random superposition of Gabor filters. We found that most of the neurons ($\sim 98\%$) showed a strong reduction in activity over a few repeats of the temporal sequence. When we presented a frame that violated the temporal sequence, these neurons responded transiently. In contrast, a small fraction ($\sim 2\%$) had activity that ramped up over several repeats, before reaching a steady, sequence-modulated response. This partitioning of the neural population into 'transient' and 'sustained' responses was observed for all temporal sequences tested. At the same time, the identity of which neurons were transient versus sustained depended on the temporal sequence. These features—adaptation to a repeated temporal sequence and a transient response to a sequence violation—are hallmarks of predictive coding. After a few repeats, the temporal sequence becomes predictable and can be efficiently represented by a small subset of the neural population. The unpredictable frame then elicits an 'error' signal because it encodes a potentially important novelty. In order to explore whether neural novelty signals could be useful to the animal, we performed behavioral experiments with matched visual stimuli that demonstrated that mice could easily learn to lick in response to a violation of an ongoing temporal sequence.

II-78. The impact of sensory uncertainty on maximally informative adaptive dynamics in neural populations

Wei-Mien Mendy Hsu[1,2] WHSU@SALK.EDU
David B Kastner[3] DBKASTNER@GMAIL.COM
Stephen Baccus[4] BACCUS@STANFORD.EDU
Tatyana O Sharpee[2] SHARPEE@SALK.EDU

[1]University of California, San Diego
[2]Salk Institute for Biological Studies
[3]University of California, San Francisco
[4]Stanford University

Sensory neural populations are optimized to transmit information about the sensory environment. This optimization has occurred through the course of evolution in the creation of different cell types, and also occurs dynamically to adapt neural responses to the current stimulus environment. Recent experimental and theoretical work has shown that when encoding a particular stimulus feature, the existence of multiple neuronal types with different thresholds increases information transmission when sensory noise drops below a certain level. This prediction across an evolutionary timescale simultaneously explains the existence of adapting and sensitizing Off retinal ganglion cells (RGCs), which have high and low thresholds for spiking, respectively, as well as the absence of comparable types among On that have higher effective noise level. However, the difference in thresholds between adapting and sensitizing cells is systematically lower than the one that would yield maximal information in an environment of stationary contrast. Yet, to achieve the optimal threshold for the current environment, ganglion cells must dynamically measure the contrast. Here we show that smaller differences in thresholds are optimal in the case where the stimulus contrast is not known but is estimated from sensory inputs. Further, we find that sensory uncertainty increases both the average firing rate and information transmission, but only in the regime of low firing rates. At high firing rates, sensory uncertainty increases the average firing rate but decreases information. Our findings reveal the relationship between sensory uncertainty and maximal information transmission in neural populations, and provide additional arguments for sparse coding in the brain.

II-79. Visualizing parallel information processing in the fruit fly retina

Yiyin Zhou YIYIN@EE.COLUMBIA.EDU
Konstantinos Psychas KP2547@COLUMBIA.EDU
Nikul Ukani NIKUL@EE.COLUMBIA.EDU
Aurel Lazar AUREL@EE.COLUMBIA.EDU

Columbia University

Individual photoreceptors in the Drosophila Retina have been extensively characterized while the overall processing of the visual field by the entire retina has not been addressed. Recent studies suggest, however, that the retina acts as a visual information preprocessor rather than merely a sampler. What is the role of this front-end in preparing the visual stimuli for later stage processing tasks such as motion detection? Towards characterizing its role as an information preprocessor, we describe algorithms for the first full-scale circuit emulation of the Retina. This enables an intuitive evaluation of retinal information processing by way of visualization of its inputs and outputs as visual scenes. This has allowed us to inspect several processing properties of the retina including 1) circuit implementation of logarithmic compression, 2) dependence of response property on mean luminance, 3) noise reduction under the neural superposition rule, 4) motion blur, and 5) visual acuity. These properties strongly influence the processing in the subsequent stages of fly early visual system. Fully compliant with the Neurokernel API, the model can be readily interfaced with models of other neuropils of the fly brain of the Neurokernel platform. The model, accompanied by open-source code, has been made publicly available in the form of a Neurokernel Request for Comments (RFC, https://neurokernel.github.io/docs). We introduced RFCs, as a successful model of knowledge creation in computer networks and operating systems, to neuroscience by publishing research papers/results on the Neurokernel website that are backed up by executable code running on GPU clusters. The open-source retina implementation presented here is highly flexible for computational experiments that complement well electrophysiological recordings from single photoreceptors and may generate hypotheses for further experiments. We welcome interested parties to experiment with our retina implementation, suggest extensions and improvements.

II-80. Inhibition underlying V1 simple cell receptive fields shapes the timing of spiking output

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

University of Pennsylvania

Thalamic afferents carrying sensory information are excitatory to their targets in cortical layer 4 (L4), which are engaged in local recurrent networks of excitatory and inhibitory neurons and ultimately shape the local response to sensory input. In the visual system, the interplay of thalamic and local excitatory and inhibitory inputs shapes neurons' receptive fields (RFs) and functional properties. In this system, excitation and inhibition are often thought to play opposing roles in the RF (Hirsch et al., 1998).We estimate synaptic excitatory and inhibitory conductances underlying the spatiotemporal receptive fields of simple cells of primary visual cortex (V1). We use optimally oriented bars flashed in different positions within the RF, while varying the membrane potential with current injection. Contrary to predictions that local excitation and inhibition act in opposition, we find that the excitatory input arising from an optimally oriented, contrast sign-matched bar flashed in a subregion of a simple cell RF evokes both an excitatory and an inhibitory conductance. This inhibition is comparable in magnitude and delayed relative to excitation. In contrast, the input from a sign-mismatched stimulus evokes pure inhibition, thus resulting in a larger spatial footprint of inhibition compared to excitation. Surprisingly, inhibition from a sign-matched ('excitatory') stimulus is often greater than that from a sign-mismatched ('inhibitory') one. We further show that delayed inhibition plays a functional role in reducing the 'window of opportunity' (Pinto et al., 2000) for sensory inputs to trigger spikes and effectively increasing the precision of the cells' output. Our findings are consistent with theoretical and experimental work in this and other systems (Cardin et al., 2010; Wilent & Contreras, 2005; Higley & Contreras, 2006; Wehr & Zador, 2003). Our results show that excitatory drive from a preferred stimulus evokes inhibition that dynamically shapes the timing of a V1 simple cell's spiking output.

II-81. Learning sparse representations of visual stimuli from natural movies

Bernal Jimenez[1] BERNALJG@BERKELEY.EDU
Jesse Livezey[1,2] JESSE.LIVEZEY@BERKELEY.EDU
Michael DeWeese[1] DEWEESE@BERKELEY.EDU

[1]University of California, Berkeley
[2]Redwood Center for Theoretical Neuroscience

Understanding what algorithms biological neurons might be implementing in V1 is crucial for understanding what algorithms the cortex might be implementing to learn in general. The simplest models of simple cells in V1 describe them as being responsive to edge patterns with varying sizes and orientations; these patterns correspond to their receptive field. A simple model for encoding visual information in V1 uses sets of spiking rates as a sparse representation of the stimuli such that the superposition of these neurons' receptive fields resemble the original image. Algorithms modelling V1 using this image representation should learn similar receptive fields to those in V1 simple cells when trained on natural visual input. Sparse coding algorithms such as locally competitive algorithms (Rozell, et. al. 2006) and SAILNet (Zylberberg, et. al.; 2011) are able to replicate the receptive fields found in V1 simple cells by enforcing sparsity and competition between neurons. In contrast with locally competitive algorithms, the Sparse and Independent Local Network (SAILNet) presents a more biophysically realistic model in which the nodes learn exclusively from information available at that node. To make SAILNet yet more biophysically plausible, the network was trained using natural movies as input instead of static images. The learning rules for SAILNet were extended to depend on neuronal spike timing rather than spiking rates. The new learning rules changed the amount of competition between neurons depending on the distance between spikes in time. When these changes were implemented in SAILNet, the network was still able to create sparse representations by learning receptive fields similar to those in V1 simple cells and obtained network statistics similar to those of the original SAILNet. This work shows that sparse coding models can learn similar biologically realistic receptive fields using continuous input, bringing them closer to learning from natural visual input.

II-82. Optimal estimation of motion-in-depth from stereo natural-image movies

Johannes Burge JBURGE@SAS.UPENN.EDU

University of Pennsylvania

As animals navigate through the environment, they must estimate the speed at which they approach objects and objects approach them. Accurate estimates of motion-in-depth depend on accurate estimates of the speeds at which images slip across the two retinas. Here, we use a Bayesian encoding-decoding framework that specifies how to encode and process task-relevant information in natural stimuli to construct neurons that are selective for the speed of motion-in-depth. First, we first generate a large, groundtruth-labeled training set of binocular natural image movies for a range of different 3D trajectories and speeds (1deg, 250ms, 1000mov/spd). Next, using Accuracy Maximization Analysis, we learn a small population of model simple cells having space-time receptive fields (RFs) optimized for the estimation of motion through depth. Two distinct binocular RF types emerge; each corresponds to a distinct mechanism for processing a different cue to motion-in-depth. The first RF type supports processing of changing disparity over time, and have low ocular dominance indices (strongly binocular). The second RF type supports processing of interocular velocity differences and have high ocular dominance indices (i.e. strongly monocular). However, although the RFs encode the information most useful for 3D motion estimation, each is not itself selective for 3D motion. The population response to natural stimuli specifies the non-linear combination rules for constructing units that are selective for 3D motion. Optimal decoding yields accurate, precise estimates of motion in depth. Recent neurophysiological work has demonstrated that many MT neurons are selective for (tuned to) the speed and trajectory of motion-in-depth. Psychophysical work has demonstrated that humans have process each aforementioned cue to motion-in-depth. The receptive fields and combination rules reported here constitute normative, testable predictions for the computations underlying these abilities with natural stimuli in real neural systems.

II-83. Cortical responses to natural and model-matched synthetic sounds reveal hierarchical computation

Samuel Norman-Haignere SVNH@MIT.EDU
Josh McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

A central goal of computational neuroscience is to construct models that explain neural responses to natural sensory signals. One approach is to test whether neural responses can be predicted from stimulus features extracted by a model. Unfortunately, there is no guarantee that features which predict cortical responses will be those that drive them, because distinct features are often correlated across natural stimuli. Here we advocate an alternative approach for model testing: the synthesis of 'model-matched' stimuli designed to evoke the same response as a natural signal in a model. If the model provides a good description of the neural response, then responses to natural and model-matched synthetic signals should be similar. We used this approach to test a standard model of auditory cortex: the spectrotemporal modulation filter bank. With fMRI, we measured voxel responses to natural and synthetic sounds matched in the marginal statistics (histogram) of each filter's response. Critically, the synthetic sounds were otherwise unconstrained, and were perceptually distinct from the natural sounds they were matched to (presumably because they lack higher-order statistical dependencies present in natural sounds). Despite these perceptual differences, natural and model-matched synthetic sounds produced nearly equivalent voxel responses in primary auditory cortex, suggesting that the spectrotemporal model accounts for much of the neural response there. In contrast, voxel responses in non-primary regions differed markedly to the natural and synthetic sounds, with many voxels producing little to no response for model-matched synthetic stimuli. This functional difference was much less pronounced when spectrotemporal power was used to predict neural responses: modulation statistics were effective predictors in non-primary regions, presumably because they are correlated with higher-order features to which non-primary regions are tuned. Our approach reveals higher-order selectivity in non-primary regions of human auditory cortex, and illustrates the use of model-matched stimuli in testing theories of cortical computation.

II-84. Learning mid-level codes for natural sounds

Wiktor Mlynarski MLYNAR@MIT.EDU
Josh McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

Auditory perception depends critically on abstract and behaviorally meaningful representations of natural auditory scenes. These representations are implemented by cascades of neuronal processing stages in which neurons at each stage recode outputs of preceding units. Explanations of auditory coding strategies must thus involve understanding how low-level acoustic patterns are combined into more complex structures. While models exist in the visual domain to explain how phase invariance is achieved by V1 complex cells, and how curvature representations emerge in V2, little is known about analogous grouping principles for mid-level auditory representations. We propose a hierarchical, generative model of natural sounds that learns combinations of spectrotemporal features from natural stimulus statistics. In the first layer the model forms a sparse, convolutional code of spectrograms. Features learned on speech and environmental sounds resemble spectrotemporal receptive fields (STRFs) of mid-brain and cortical neurons, consistent with previous findings. To generalize from specific STRF activation patterns, the second layer encodes patterns of time-varying magnitude (i.e. variance) of multiple first layer coefficients. Because it forms a code of a non-stationary distribution of STRF activations, it is partially invariant to their specific values. Moreover, because second-layer features are sensitive to STRF combinations, the representation they support is more selective to complex acoustic patterns. The second layer substantially improved the model's performance on a denoising task, implying a closer match to the natural stimulus distribution. Quantitative hypotheses emerge from the model regarding selectivity of auditory neurons characterized by multidimensional STRFs and sensitivity to increasingly more abstract structure. The model also predicts that the auditory system constructs representations progressively more invariant to noise, consistent with recent experimental findings. Our results suggest that mid-level auditory representations may be derived from high-order stimulus dependencies present in the natural environment.

II-85. Behavioral and neural tuning for acoustic communication signals in Drosophila

Jan Clemens CLEMENSJAN@GMAIL.COM
David Deutsch DDEUTSCH@PRINCETON.EDU
Isabel D'Alessandro ISABELD@PRINCETON.EDU
Mala Murthy MMURTHY@PRINCETON.EDU

Princeton University

Acoustic communication in Drosophila constitutes an ideal model system for understanding the neural computations underlying pattern recognition. During courtship, male flies chase females and produce song while females base their mating decision on temporal song features. Both sexes rapidly change their speed when hearing song, thereby providing a highly-resolved behavioral readout of song responses: Females slow down in response to conspecific male song, allowing the male to interact with her; males speed-up in search of a courtship target. To comprehensively describe the behavioral tuning for temporal features of song, we use a novel automated, single-fly playback assay and find that females are highly selective for particular song features. That is, only few artificial stimuli are attractive to females (induce slowing), while most song stimuli are neutral or even aversive (induce speeding-up)—the existence of aversive song stimuli in Drosophila has to our knowledge not been reported before. A model consisting of a Gabor-filter and an integrator can explain this behavioral tuning, indicating that relatively simple computations underlie song processing in Drosophila. However, the relation between stimulus features and motor outputs is highly dynamical, e.g. some stimuli induce transient slowing followed by speeding-up. In addition, aversive or neutral stimuli can suppress responses to subsequent attractive songs, demonstrating that flies integrate song across stimulus presentations. To probe the neural code for song in auditory neurons in the brain we perform Calcium imaging. Generally, song feature tuning of a subset of 4th order auditory neurons resembles the behavioral tuning. However, neuronal responses lack the rich dynamics and the sexual dimorphism observed in the behavior. These neurons thus act as feature encoders, while most of the behavioral dynamics must be generated in downstream decision circuits. Our behavioral experiments thus provide a powerful constraint for models of the neural computations underlying acoustic communication in Drosophila.

II-86. Integration across stimulus dimensions in auditory cortex.

David Sloas[1] SLOAS@BU.EDU
Hongbo Xue[1] HONGBOX@BU.EDU
Ran Zhuo[1] RAZH9605@BU.EDU
Anna Chambers[2,3] ANNA_CHAMBERS@MEEI.HARVARD.EDU
Eric Kolaczyk[1] KOLACZYK@BU.EDU
Daniel Polley[2,3] DANIEL_POLLEY@MEEI.HARVARD.EDU
Kamal Sen[1] KAMALSEN@BU.EDU

[1]Boston University
[2]Eaton Peabody Laboratories
[3]Massachusetts Eye and Ear

Although sensory cortex is thought to be important for perception related to complex objects, its role in representing complex stimuli remains unknown. Complex objects are often rich in information along multiple stimulus dimensions. For example a complex auditory object can be characterized by parameters such as frequency, bandwidth, amplitude modulation, intensity and location. The position of cortex in the sensory hierarchy suggests that cortical neurons may integrate across these dimensions to form a more gestalt representation of auditory objects. Yet, studies of cortical neurons typically only explore a single or few dimensions due to the difficulty of searching for optimal stimuli in a high dimensional stimulus space. Recently, we demonstrated the power of using evolutionary algorithms (EA's) to explore multi-dimensional stimulus spaces for auditory cortical neurons by developing an EA that uses real-time single unit spike feedback in awake mice to rapidly converge on optimal stimuli in a five-dimensional stimulus space1. However, it remains unknown whether multi-dimensional characterizations are necessary for cortical neurons or if it is sufficient to characterize responses to a single dimension at a time, i.e., one-dimensional tuning curves. The answer to this depends on whether cortical neurons are sensitive to interactions between stimulus dimensions. Here, we apply a statistical non-linear regression method, the Generalized Additive Model (GAM), to address this question. We find that auditory cortical neurons are sensitive to interactions across stimulus dimensions, and these interactions are highly diverse across the population, indicating significant integration across multiple stimulus dimensions in auditory cortex. These results motivate the use of multi-dimensional stimuli in auditory cortex, and the novel paradigm combining the EA and the GAM has the potential to reveal new facets of cortical integration in the auditory and other sensory cortices. We also illustrate the use of this paradigm in comparing cortical vs. pre-cortical processing.

II-87. Nonlinear Bayesian cue integration explains the dynamics of vocal learning

Baohua Zhou BZHOU7@EMORY.EDU
Sam Sober SAMUEL.J.SOBER@EMORY.EDU
Ilya Nemenman ILYA.NEMENMAN@EMORY.EDU

Emory University

The acoustics of vocal production in songbirds is tightly regulated during both development and adulthood as birds progressively refine their song using sensory feedback to match an acoustic target. Here we perturb this sensory feedback using headphones to shift the pitch (fundamental frequency) of the song. When the pitch is shifted upwards (downwards), birds eventually learn to compensate and sing lower (higher), bringing the experienced pitch closer to the target. Paradoxically, the speed and amplitude of this sensorimotor learning decrease for larger introduced errors, so that the animal responds rapidly to a small sensory perturbation, while seemingly ignoring and never correcting a much bigger one. Similar results are observed for other species and behaviors. We develop a mathematical model based on nonlinear Bayesian integration of two sensory modalities (one perturbed by the headphones and the other not) that quantitatively explains all of these observations. Furthermore, the model makes predictions about the structure of the probability distribution of the pitches sung by birds during the pitch shift experiments, which are confirmed experimentally.

II-88. Oxytocin enables maternal behavior by balancing cortical inhibition

Bianca Marlin BIANCAJONESMARLIN@GMAIL.COM
Mariela Mitre MARIELA.MITRE@MED.NYU.EDU
Ioana Carcea IOANA.CARCEA@MED.NYU.EDU
Jennifer Schiavo JENNIFER.SCHIAVO@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 essential for maternal behavior (Insel and Young, 2001; Churchland and Winkielman, 2012; Dulac et al., 2014). Here we describe how increased synchronous activity in primary auditory cortex correlates with the improvement of an auditory-dependent maternal behavior in mice, and how oxytocin enhances responses to pup calls in left primary auditory cortex. Naive virgin females initially do not retrieve pups. We found that oxytocin (pharmacologically-applied or optogenetically-released) accelerated the time to first retrieval in virgin females. Expression of retrieval behavior required left but not right auditory cortex, and was accelerated by systemic or local oxytocin. We made new antibodies specific to the mouse oxytocin receptor, and found preferential expression in left auditory cortex (Marlin et al., 2015). Electron microscopy revealed oxytocin receptors at synapses, including inhibitory terminals on excitatory neurons. We next performed in vivo whole-cell recordings to determine the transformation of pup call responses from the virgin state to the maternal state. We made current- and voltage-clamp recordings to measure spiking and synaptic responses. We used awake behaving tetrode recordings and awake in-vivo calcium imaging to track changes in auditory cortical neurons during the presentation of pup call stimuli, before and after co-housing. 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.

II-89. Regulation of sensory feature selectivity by the thalamic reticular nucleus

Lukas Schmitt LUKAS.SCHMITT@NYUMC.ORG
Michael Halassa MICHAEL.HALASSA@NYUMC.ORG

New York University

The thalamic reticular nucleus (TRN) is a critical gate for sensory information flowing from thalamus to cortex. We, and others, have demonstrated that the TRN is composed of individual subnetworks, each capable of broadly controlling sensory flow in a modality specific manner. Given the precise topographic relationship between the TRN and sensory thalamus, several investigators have suggested that the TRN may exert more fine-grained control over thalamic transmission, including the selection of particular sensory features. This idea is consistent with recent findings in the visual TRN, where neurons appear to exhibit complex sensory representations. However, the generality of such feature specific representations across sensory modalities and their impact on thalamic output have not been investigated. Here, we address these issues by recording auditory responses from the auditory TRN (audTRN) and thalamus (medial genicular body; MGB) in the awake, head-fixed mouse. In agreement with previous studies, we found that MGB neurons exhibit spectrotemporal receptive fields (STRF) characterized by small numbers of nearby excitatory peaks with adjacent suppressive bands. Silencing the audTRN greatly diminished the suppressive bands and resulted in broader spectral and temporal tuning of MGB neurons. In addition, this manipulation produced a number of non-linear effects including shifts in the main response peak and the appearance of multiple peaks. This diversity of changes was mirrored in the complex STRFs of audTRN

neurons, which show multiple peak responses widely distributed in frequency space. Interestingly, these peaks showed characteristic spectral and temporal organization suggesting that TRN neurons may respond preferentially to temporally dynamic sound stimuli. Together, these results raise the possibility that the TRN is involved in sensory feature detection, and that it may significantly influence preprocessing of sensory input occurring at the level of the thalamus.

II-90. Compressed sensing allows robust and accurate sensory signal recovery

Rishabh Raj RRAJ@STOWERS.ORG
Dar Dahlen DAR.DAHLEN@GMAIL.COM
Ron Yu CRY@STOWERS.ORG

Stowers Institute for Medical Research

The brain extracts salient features from vast numbers of possible sensory stimuli in the environment. The neural code that allows the brain to robustly encode high dimensional stimulus and effectively represent salient features remains unknown. In many sensory systems, representation of sensory stimuli is sparse and sparse coding emerges as an effective coding scheme. However, it is not clear how population activities of neurons are decoded to accurately represent the stimuli. In our study, we find that sparse transformation of neuronal responses permits a decoding scheme to recover signal from neuronal response patterns that is akin to Compressed Sensing. In the mouse olfactory system, odor identities are first represented by the combinatorial activation of the glomeruli in the olfactory bulb. The responses are sparsened in the cortex. Based on olfactory glomerular responses to large panels of odorants, we find that sparsening allows odor identity and intensity information to be recovered precisely and robustly from the responses of small numbers of randomly selected olfactory glomeruli and in the presence of high levels of noise. We show that a biologically realistic network can perform the decoding effectively. Our results suggest that Compressed Sensing allows recovery of sensory input from incomplete and inaccurate responses patterns of the neuronal population to provide a robust code of neural information.

II-91. Invariant odor coding through a primacy code

Christopher Wilson[1] CHRISTOPHER.WILSON@NYUMC.ORG
Alexei Koulakov[2] KOULAKOV@CSHL.EDU
Dmitry Rinberg[1] DMITRY.RINBERG@NYUMC.ORG

[1]New York University
[2]Cold Spring Harbor Laboratory

Odors retain stable perceptual features across concentrations. This perceptual stability occurs despite changes in the odor's representation in the sensory input layer of the olfactory system. We propose that the olfactory system uses a temporal ranking coding scheme to achieve concentration-tolerant representations. Under this coding scheme, called 'primacy coding', only a small number of sensor neurons which are activated earliest are responsible for forming odors' perceptual identities. Here, we test two predictions of this model: that the earliest odor-evoked activity should be both sufficient and necessary to form olfactory identities. To test this, we developed an optogenetic masking behavioral paradigm to test the time scale on which olfactory sensory information is integrated. Using this technique, we demonstrate that mice can preform a simple olfactory discrimination task using odor-evoked activity occurring the first 100 ms following the initial inhalation of odorant. We also demonstrate that a masking stimulus which disrupts initial odor-evoked activity while preserving some late-occurring activity impairs animals' behavioral performance. These results provide evidence that this novel coding scheme is a viable candidate for forming olfactory identity percepts from unstable sensory input.

II-92. Functionally and structurally distinct subnetworks in the sensory cortex

Jason Wittenbach[1,2] WITTENBACHJ@JANELIA.HHMI.ORG
Simon Peron[1,2] PERONS@JANELIA.HHMI.ORG
Karel Svoboda[1,2] SVOBODAK@JANELIA.HHMI.ORG
Jeremy Freeman[1,2] FREEMANJ11@JANELIA.HHMI.ORG

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

Neurons within cortical circuits can be assigned to different functional groups on the basis of their activity patterns. Little is known about the circuit mechanisms that give rise to these distinct functional groups. In layer (L) 2/3 of the mouse primary vibrissal somatosensory cortex (vS1), 'whisking' neurons encode whisker position and 'touch' cells encode whisker-object contact. We have recently shown that targeted ablations of touch neurons dramatically reduce activity in the remaining population; in contrast, ablations of whisking neurons have little effect. Three mechanisms could account for this difference: the time scale of stimulus input (which differs for touch and whisking); the specificity of stimulus input to a subnetwork of targeted neurons; the degree of strong recurrent coupling within a subnetwork. To understand the contributions of these different mechanisms we analyzed a computational model of the L2/3 microcircuit. The model shows that all three factors contribute, but recurrent coupling most directly affects susceptibility to ablations. With sufficient recurrent coupling, ablation effects are found across a range of time scales, though effects are enhanced for faster inputs. While a minimal level of input specificity is required, it otherwise has little effect. Alongside our experimental findings, these results predict strong recurrent coupling in the L2/3 touch network but not the whisking network.

II-93. Evolving amygdala neuronal networks shape the unpleasantness of pain

Gregory Corder GCORDER@STANFORD.EDU
Biafra Ahanonu BAHANONU@STANFORD.EDU
Benjamin Grewe GREWE@STANFORD.EDU
Mark Schnitzer MSCHNITZER@GMAIL.COM
Gregory Scherrer GS25@STANFORD.EDU

Stanford University

Pain is a multidimensional experience comprising sensory and emotional modules, but it remains unclear how nociceptive sensory signals are imbued with negative affect. The basolateral amygdala (BLA), classically involved in learned fear and anxiety, assigns an emotional valence to highly salient, external stimuli. To understand how the BLA encodes the experience of qualitatively different noxious stimuli, and how this activity changes during the development of pathological neuropathic pain, we imaged BLA ensemble dynamics in response to numerous sensory stimuli, increasing from neutral to noxious intensities, over 42 days after a peripheral nerve injury. Large-scale imaging of BLA projection neurons (expressing the calcium indicator GCaMP6m) was performed in freely behaving mice implanted with a small microendoscope. In normal mice, prior to nerve injury, principle component representational maps of extracted, time-locked neuronal and behavioral nociceptive responses revealed small (5-8% of all imaged neurons) and unique ($>$10% overlap between modality-specific representations) patterns of activity for thermal and mechanical, and innocuous and noxious stimuli. Strikingly, after the establishment of neuropathic pain the ensemble responses evoked by prior innocuous stimuli transformed such that the network representation was more similar to activity patterns evoked by frankly noxious stimuli. Thus, to establish the function of this neuronal network to the affective component of pain we chemogenetically sileced neuronal activity by virally-expressing the inhibitory DREADD(hM4) receptor in pain-active BLA neurons of FosCreERT2 mice. Silencing the BLA pain network in mice with neuropathic pain did not alter reflexive responses to noxious stimuli. By contrast, the amount of time spent attending to the injured paw following noxious stimulation, as well as motivated escape away from the noxious stimulus, was dramatically reduced. Together, our results point to a

possible neural substrate for chronic pain and show for the first time how neural ensemble coding may shape the experience of pain unpleasantness.

II-94. On the spike train variability characterized by variance-to-mean power relationship

Shinsuke Koyama SKOYAMA@ISM.AC.JP
The Institute of Statistical Mathematics

We propose a statistical framework for modeling the non-Poisson variability of spike trains observed in a wide range of brain regions. Central to our approach is the assumption that the variance and the mean of ISIs are related by a power function characterized by two parameters: the scale factor and exponent. This single assumption allows the variability of spike trains to have an arbitrary scale and various dependencies on the firing rate in the spike count statistics, as well as in the interval statistics, depending on the two parameters of the power function. On the basis of this statistical assumption, we show that the power laws with various exponents emerges in a stochastic leaky integrate-and-fire model and in a conductance-based neuron model with excitatory and inhibitory synaptic inputs, depending on the input regimes and the ratio between excitation and inhibition. We also discuss based on this result that the conventional assumption of proportional relationship between the spike count mean and variance could lead to the wrong conclusion regarding the variability of neural responses. Finally, we propose a statistical model for spike trains that exhibits the variance-to-mean power relationship, and a maximum likelihood method is developed for inferring the parameters from rate-modulated spike trains.

II-95. Margin learning in spiking neurons

Rafael Brune BRUNE@EM.MPG.DE
Robert Guetig GUETIG@EM.MPG.DE
Max Planck Institute - Experimental Medicine

Most neurons receive inputs from thousands of afferents. Correspondingly, a pattern of presynaptic activity can be thought of as a point in a high dimensional space whose coordinates represent the activity of each individual afferent. To realize a specific function, e.g. recognize grandmother's face, a neuron is required to respond exclusively to input patterns that lie within a specific region of this input space. The ability of neurons to realize difficult classification tasks through simple decision boundaries that separate complex input categories is tightly linked to the high dimensionality of their input spaces. However, this advantage of high dimensional neural representations comes at a price: Learning is difficult in high dimensional spaces. In particular, a neuron's ability to generalize from a limited number of training examples is impaired by overfitting when the number of synaptic efficacies is large. Constituting a major breakthrough in machine learning, learning in high-dimensional spaces has been greatly improved by margin techniques that maximize the minimal distance between available training examples and the learned decision boundary. However, naive applications of margin learning to spiking neurons are limited to binary responses were the margin between no-output-spike versus at-least-one-output-spike can be captured by the distance between a neuron's maximal postsynaptic potential and its firing threshold. Here we introduce a margin measure for spiking neurons that can be applied to arbitrary numbers of output spikes, e.g. the margin between three-output-spikes versus four-output-spikes. We show that a family of gradient-based learning rules that operate on these margins strongly improves the learning capabilities of spiking neurons. We successfully apply this margin learning to a neural model of phoneme recognition. This work transfers powerful margin-based learning concepts from machine to neurobiological learning and enables biologically plausible models of spiking neurons to learn from limited examples in high dimensional inputs spaces.

II-96. All-optical causal measurement of single-neuron effective connectivity in behaving mice

Selmaan Chettih[1] CHETTIH@FAS.HARVARD.EDU
Christopher D Harvey[2] CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

[1]Harvard University
[2]Harvard Medical School

The spiking output of a neuron can be viewed both as representing information about the organism and world, and as playing a causal role in generating the dynamics of neural activity. Measuring monosynaptic connectivity between neurons is the dominant method to assess a neuron's contribution to dynamics. However this measurement is rarely combined with functional characterization due to experimental difficulty, and furthermore there are substantial difficulties predicting dynamics from connectivity. We therefore developed a complementary method to directly measure the causal influence of one neuron's spiking on its neighbors. Our approach uses two-photon excitation to stimulate single neurons expressing the opsin C1V1 while monitoring supra-threshold activity in hundreds of nearby cells with GCaMP6 imaging in the superficial cortex of awake, head-fixed mice. This integrates our measurement into a common experimental preparation for monitoring spontaneous and behavior-evoked activity. We built a two-photon microscope with two scan heads integrated into a virtual-reality system for head-fixed mice performing navigation and other cognitively demanding tasks, permitting simultaneous resonant-scanning GCaMP imaging and linear galvanometer-targeted photostimulation. Previous studies suggested that the magnitude and resolution of stimulation would be limiting factors. We thus introduced a sub-cellular localization sequence to C1V1 restricting its presence to the cell body, reducing unintended off-target stimulation and increasing the fraction of total current available to cell body stimulation. Rapidly scanning the cross-sectional area of single neurons using high-numerical aperture excitation, we were able to induce activity in targeted neurons with substantially enhanced magnitude and resolution compared to published results. We performed experiments repeatedly stimulating $>$100 cells one-by-one while imaging activity in hundreds of neurons. We found significant responses to single-neuron stimulation across distances of tens to hundreds of microns, at the population level and for a sparse set of individual pairs, establishing the feasibility of our approach.

II-97. Apical dendrite as a canonical correlation analyzer

Tatsuya Haga TATSUYA.HAGA@RIKEN.JP
Tomoki Fukai TFUKAI@RIKEN.JP

RIKEN Brain Science Institute

Pyramidal neurons have apical dendrites that are divided into proximal and distal parts. It has been experimentally found that synaptic inputs to distal dendrites are not directly integrated into the somatic membrane potential, but they interacts with the somatic activity via calcium spikes in the apical dendrite which amplifies and potentiates coincident inputs across compartments (Larkum, 2013). To explore the role of this dendritic computation, we construct a neuron model that consists of a somatic compartment (including proximal dendrites) and a distal dendritic compartment, in which synchronous synaptic inputs to the two compartments generate calcium spikes. Synapses on both compartments are modified by independent Bienenstock-Cooper-Munro (BCM) theory and additional potentiation by calcium spikes. In contrast to principal component analysis in a single-compartmental Hebbian learning model (Oja, 1982), this two-compartmental model can perform canonical correlation analysis (CCA) (Hotelling, 1936) between distal dendritic and somatic inputs. Hebbian learning generally extracts major inputs from the largest population, whereas CCA-like learning of our model modifies synaptic weights to maximize correlation between activities of two compartments. Therefore, our model can analyze inputs to one compartment by the instruction inputs to the other compartment, and they are selectively amplified by calcium spikes. This property helps the formation of the conjunctive representation of multiple features, and the extraction of weak sensory features based on top-down signals or internally generated activity patterns. Furthermore, we combine our neuron model with a recurrent network that generates spontaneous firing sequences to realize the learning

of receptive fields in distal dendrites based on pre-embedded patterns in somatic network. This model explains ‘preplay’ sequences of place cells in hippocampus (Dragoi and Tonegawa, 2011).

II-98. Supervised learning sets benchmark for robust spike rate inference from calcium imaging signals

Matthias Bethge[1] MBETHGE@UNI-TUEBINGEN.DE
Lucas Theis[1] LUCAS.THEIS@BETHGELAB.ORG
Philipp Berens[1] PHILIPP.BERENS@BETHGELAB.ORG
Emmanouil Froudarakis[2] EFROUD@CNS.BCM.EDU
Jacob Reimer[2] REIMER@BCM.EDU
Miroslav Roman-Roson[1] MIROSLAV.ROMAN-ROSON@CIN.UNI-TUEBINGEN.DE
Thomas Baden[1] THOMAS.BADEN@UNI-TUEBINGEN.DE
Thomas Euler[1] THOMAS.EULER@CIN.UNI-TUEBINGEN.DE
Andreas Tolias[2] ASTOLIAS@BCM.EDU

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

A fundamental challenge in calcium imaging has been to infer spike rates of neurons from the measured noisy calcium fluorescence traces. We collected a large benchmark dataset (>100.000 spikes, 73 neurons) recorded from varying neural tissue (V1 and retina) using different calcium indicators (OGB-1 and GCaMP6s). We introduce a new algorithm based on supervised learning in flexible probabilistic models and systematically compare it against a range of spike inference algorithms published previously. We show that our new supervised algorithm outperforms all previously published techniques. Importantly, it even performs better than other algorithms when applied to entirely new datasets for which no simultaneously recorded data is available. Future data acquired in new experimental conditions can easily be used to further improve its spike prediction accuracy and generalization performance. Finally, we show that comparing algorithms on artificial data is not informative about performance on real data, suggesting that benchmark datasets such as the one we provide may greatly facilitate future algorithmic developments.

II-99. Bayesian latent state space models of neural activity

Scott Linderman[1] SLINDERMAN@SEAS.HARVARD.EDU
Aaron Tucker AARONTUCKER@COLLEGE.HARVARD.EDU
Matthew Johnson[2] MATTJJ@CSAIL.MIT.EDU

[1]Harvard University
[2]Harvard Medical School

Latent state space models such as linear dynamical systems and hidden Markov models are extraordinarily powerful tools for gaining insight into the latent structure underlying neural activity. By beginning with simple hypotheses about the latent states of neural populations and incorporating additional beliefs about the nature of this state and its dynamics, we can compose a sequence of increasingly sophisticated models and evaluate them in a statistically rigorous manner. However, inferring the latent states and parameters of these models is particularly challenging when presented with discrete spike counts, since the observation likelihoods are not conjugate with latent Gaussian structure. Thus, we often resort to model-specific approximate inference algorithms which preclude rapid model iteration and typically provide only point estimates of the model parameters. As a result, it is difficult compare models in a way that is robust to the approximation and the particular estimates of the model parameters. Here, we develop a unified framework for composing latent state space models and performing efficient Bayesian inference by leveraging a data augmentation strategy to handle the discrete spike count observations.

This framework is easily extensible, as we demonstrate by constructing an array of latent state space models with a variety of discrete spike count distributions and fitting them to a simultaneously recorded population of hippocampal place cells. Our Bayesian approach yields a posterior distribution over latent states and parameters, which enables robust prediction and principled model comparison. Moreover, we show that our method is very efficient, performing faster than alternative point-estimation approaches in our experiments.

II-100. Towards ground truth in ultra-dense neural recording

Brian Allen[1] BDALLEN@MIT.EDU
Caroline Moore-Kochlacs[2] CAROMK@GMAIL.COM
Jorg Scholvin[1] SCHOLVIN@MIT.EDU
Justin Kinney[1] JKINNEY@MIT.EDU
Jacob Bernstein[1] JBERNS@MIT.EDU
Suhasa B Kodandaramaiah[1] SUHASABK@MIT.EDU
Nancy Kopell[2] NK@MATH.BU.EDU
Edward Boyden[1] ESB@MEDIA.MIT.EDU

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

We are interested in whether ultra-dense neural recording can enable high fidelity, automatic spike sorting. We recently developed probes with close-packed recording sites (e.g. 9x9 μm sites, $1\mu m$ apart) designed so that spikes from a given neuron may be sensed by many recording sites. Computational modeling work and preliminary data suggest that this feature may allow for the effective use of independent component analysis (ICA) for spike sorting , which operates robustly and optimally when there are at least as many signals (voltage readings on the probe sites) as sources (neurons), when combined with a suitable classifier. Here we perform a series of ultra-dense recordings in primary visual cortex of head-fixed mice, with simultaneous patch clamp recording. Patch clamp recording grants ground truth access to the voltage of a single neuron. We here discuss preliminary data from a 128-channel probe designed to record across multiple cortical layers, with n = 4 cells analyzed in detail, and describe potential future directions for using ground truth measurement to validate probe architectures and spike sorting algorithms.

II-101. Discovering structure in connectomes using latent space kernel embedding

Eric Jonas[1] JONAS@ERICJONAS.COM
Srinivas Turaga[2,3] TURAGAS@JANELIA.HHMI.ORG

[1]University of California, Berkeley
[2]Janelia Farm Research Campus
[3]Howard Hughes Medical Institute

It has now become possible to map the synaptic connectivity of neural circuitry at the cellular resolution using electron microscopy. In this work, we present a new class of models for the analysis of connectomic data. Many theories of neural computation propose specific patterns of neural connectivity tied to the tuning properties of neurons. We propose an extension to traditional latent space models to uncover continuous hidden structure in these connectomes, such as the neural tuning property of a neuron and the function that determines neural connectivity. Our scalable model provides the flexibility to recover structure in both directed and undirected graphs. We demonstrate our model on synthetic connectomes and on the recently published mouse retinal connectome .

II-102. Sortfree: Using threshold crossings to evaluate scientific hypotheses in population analyses.

Eric Trautmann[1] ETRAUTMANN@GMAIL.COM
Sergey Stavisky[1] SERGEY.STAVISKY@STANFORD.EDU
Matt Kaufman[2] MKAUFMAN@CSHL.EDU
K Cora Ames[3] KCA2120@COLUMBIA.EDU
Stephen Ryu[4] SEOULMAN@STANFORD.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU

[1]Stanford University
[2]Cold Spring Harbor Laboratory
[3]Columbia University
[4]Palo Alto Medical Foudnation

In this work, we aim to address a major challenge facing systems neurophysiological experiments. How can we cope with the challenge of spike sorting datasets as the number of recorded channels increases from roughly 100 up to many hundreds or thousands? For a typical experiment, it currently takes approximately 8-16 hours to hand sort spikes on 100 channels. Hand-sorting all datasets may not be necessary, however. The highest performance brain machine interfaces demonstrated to date in primates and humans use threshold crossings instead of tracking neurons (e.g: Gilja 2012). Why is it reasonable to use threshold crossings instead of carefully isolated single units? In many brain areas, the recorded dimensionality of neural activity is lower than number of neurons. If the network activity is low dimensional relative to the number of recorded units, then it is reasonable to expect that combining several single units into multiunit channels will not introduce large distortions when estimating the dynamical activity at the level of the population. To investigate this, we reprocessed and re-analyzed data collected by Ames and Shenoy 2014 and Churchland et al. 2012, substituting electrical threshold crossings for hand sorted units. In both experiments, we found that the analyses reached significance and supported the hypothesis and conclusions presented, closely recapitulating both qualitative and quantitative features in the original datasets. This approach is most suitable for analyses that rely on linear readouts of population activity. We anticipate that using threshold crossings in place of spike sorting will become increasingly important and relevant for population analyses in order to address the deluge of data created by new recording technologies, as this method is theoretically justified, empirically supported, and simple.

II-103. Bayesian targeted dimensionality reduction for neural population activity

Mikio Aoi[1] MAOI@PRINCETON.EDU
Valerio Mante[2,3] VALERIO@INI.UZH.CH
Jonathan W Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]University of Zurich
[3]ETH Zurich

A growing body of evidence indicates that neural population activity during perceptual, motor, and decision-making behavior is low dimensional relative to the number of neurons we can record in an experiment. It is also becoming clear that many brain areas exhibit ‘mixed selectivity’ and carry heterogeneous representations of sensory and behavioral information. However, classical methods for dimensionality reduction like PCA do not identify dimensions that carry information about external variables. To address this shortcoming, we introduce a Bayesian method for targeted dimensionality reduction, inspired by a method from (Mante, Sussillo, & Newsome 2013), and similar in spirit to ‘demixed PCA’ (Machens 2010). We frame dimensionality reduction as a reduced-rank regression problem, which provides a principled, model-based method for identifying a low-dimensional projection of neural activity that is informative about task and behavioral variables. The resulting model characterizes the

relationship between external variables and neural population activity in terms of a small number of weight vectors (each of which defines a ‘direction’ in the high-dimensional space of neural activity) and time-courses (each of which describes temporal modulation of neural activity by a particular external variable). The framework allows us to integrate over weights and use marginal likelihood to estimate the dimensionality of the subspace encoding each covariate. To illustrate the performance of our method, we analyze of the dimensionality of population encoding in neural data recorded from prefrontal cortex (PFC) in macaques performing a context-dependent discrimination task. We show that each external variable (color, motion, decision, context) has a multi-dimensional encoding in the neural population. We believe that this approach will have useful applications for analyzing mixed, heterogeneous neural representations in a wide variety of other brain areas.

II-104. Fast, scalable Bayesian inference for high-dimensional neural receptive fields

Mikio Aoi[1] MAOI@PRINCETON.EDU
Anqi Wu[1] WUANQI1129@GMAIL.COM
Ikuko Smith[2] IKUKOTS@GMAIL.COM
Spencer Smith[2] SLSLAB@GMAIL.COM
Jonathan W Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]University of North Carolina

We examine the problem of rapidly and efficiently estimating a neuron’s linear receptive field (RF) from responses to high-dimensional stimuli. This problem poses important statistical and computational challenges due to the need for strong regularization in high-dimensional parameter spaces. In particular, we focus on several methods for scaling up automatic smoothness determination (ASD) [Sahani & Linden 2003], an empirical Bayesian RF estimation method that uses the marginal likelihood to set hyperparameters governing shrinkage and smoothness. First, we provide an arbitrarily precise Fourier representation of the ASD prior covariance that leads to substantially better scaling with smooth filters. Second, we introduce several approximate methods that exploit Kronecker and Toeplitz structure in the stimulus covariance, an approach inspired by the method of expected log-likelihoods, but permits use of naturalistic stimuli for which the stimulus covariance is unknown. The resulting estimator allows for a substantial gains in efficiency for evaluating the marginal likelihood: classical ASD requires $O(d\hat{}3m)$ time and $O(d\hat{}2m)$ memory for an m-dimensional RF with d elements per dimension; our tricks reduce this cost to $O(mdlogd)$ time and $O(md)$ memory. We show that evidence optimization for a linear RF with 160K coefficients using 5K samples of data can be carried out on a laptop in $\approx 2s$, with minimal loss in accuracy. We demonstrate the use of these methods on neuronal responses recorded from mouse V1. The resulting methods will allow neuroscientists to consider large-scale RF estimation problems that were previously infeasible, and suggest a new approach for scaling up other neural characterization methods.

II-105. Cortex-wide cellular-resolution imaging and analysis of neural activity in head-fixed behaving mice

Ben Huang BSH@UCLA.EDU
Arash Bellafard ABELLAFARD@MEDNET.UCLA.EDU
Peyman Golshani PGOLSHANI@MEDNET.UCLA.EDU

University of California, Los Angeles

Two-photon Ca2+ imaging in awake head-fixed mice has enabled studies of behaviorally-relevant circuit dynamics at cellular resolution. Yet, one current limitation is the need to restrict imaging to a pre-selected region-of-interest. While this targeted approach has produced valuable insights into more functionally-circumscribed regions, it can be limiting for studies of higher-order cognitive functions involving distributed networks spanning multiple regions. Studies of cognitive circuits and computation thus require imaging access that spans the entire cortex within an individual behaving animal. To achieve this goal, we have developed a cortex-spanning cranial-window preparation that provides long-term cellular-resolution optical access to the entire dorsal cortical surface of adult mice. To further extend access to cortical surface buried within the midline fissure, we incorporated right-angle microprisms within our cortex-spanning window to image regions such as the mPFC (medial prefrontal cortex) and ACC (anterior cingulate cortex) that play important roles in cognitive processing. We have implanted these windows in GCaMP6-expressing transgenic mice and imaged cortical activity across the frontal, parietal and occipital regions, producing rich datasets of cortex-wide population dynamics within individual animals. To characterize the spatiotemporal structure of cortex-wide activity, we analyzed the correlation structure of the spontaneous activity and response variability during state-dependent behaviors. We implemented a simple startle paradigm to head-fixed mice that were free to run on a spherical treadmill, while we imaged calcium activity in layer-2/3 pyramidal neurons across the entire anterior-posterior axis from frontal to occipital cortex. Our initial analysis reveals that the structures of the activity correlation matrix are distinct between regions, but are stable within each region across behavioral states (stationary, running, startle). Using the surface vasculature as a map, we can register the imaged regions to anatomical coordinates, allowing us to evaluate whether the differences in correlation structure coincide with anatomically-defined regions.

III-1. On the neural basis of probabilistic inference during perceptual decision making

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

[1]University of Rochester
[2]Brown University
[3]NIH/National Eye Institue

The mathematical concept of probabilistic inference has previously been suggested as a framework for understanding perception (Lee & Mumford 2003; Yuille & Kersten 2006), but linking these ideas to the spiking activity of sensory neurons has been a challenge because little is known about the details of the brain's internal model. Here, we show how to overcome this challenge in the context of psychophysical tasks. Our insight is that training a subject on a task induces a task-dependent perturbation in the subject's internal model; and, assuming that sensory neurons' responses represent the posterior over some latent variables, this pertubation leads to predictable changes in the responses of sensory neurons independent of the details of the representation (whether parametric (Ma et al. 2006) or sampling-based (Fiser et al. 2010)). Based on this insight, we derived empirically testable predictions for an orientation discrimination task and confirmed them using population recordings from monkey V1 (Bondy & Cumming, SfN 2013). Importantly, we show that the variability and covariability of V1 responses have significant components aligned with the prior predicted by the task. To confirm this task-dependence, we then constructed other 'hypothetical' tasks and showed that the empirical moments are most aligned with the actual task—patterns that are not explainable in a pure feedforward framework. Furthermore, we propose a method to infer the internal model used by the brain for a given task from the statistical structure of sensory responses. We illustrate the capabilities of our method by applying it to synthetic data (Haefner et al. 2014) for which the ground truth is known. Given only the sensory responses, we are able to infer a rich characterization of the internal structure of the model, demonstrating our method's power for tracking changes due to learning or attention in the brain.

III-2. Marked point process filter for clusterless, and adaptive encoding-decoding of multiunit activity

Kensuke Arai[1] KENSUKE.Y.ARAI@GMAIL.COM
Daniel Liu[2] DANIEL.LIU@UCSF.EDU
Loren Frank[2] LOREN@PHY.UCSF.EDU
Uri Eden[1] TZVI@MATH.BU.EDU

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

Real-time, closed-loop experiments can uncover causal relationships between specific neural activity and behavior. An important advance in realizing this is the marked point process filtering framework which utilizes the "mark" or the waveform features of unsorted spikes, to construct a relationship between these features and behavior, which we call the encoding model. This relationship is not fixed, because learning changes coding properties of individual neurons, and electrodes can physically move during the experiment, changing waveform characteristics. We introduce a sequential, Bayesian encoding model which allows incorporation of new information on the fly to adapt the model in real time, for experiments which can be segmented into multiple epochs where encoding followed by decoding, occur. In the encoding phase, behavior and marks are observed simultaneously to construct updates to the encoding model, while in the decoding phase, only the marks are observed, and we decode the corresponding behavior. Model parameter posteriors are obtained during encoding using Gibbs sampling, with posterior of the previous epoch being incorporated as a prior to the current epoch. The priors reflect how certain we are about model parameter values, and we may relate the "width" of the prior to the notion of how quickly receptive fields and recording quality changes in a unit of time. A possible application of this framework is to the decoding of the contents of hippocampal ripples in rats exploring a maze. During physical exploration, we observe its position and marks to update the encoding model, which is employed to decode contents of ripples when rats stop moving, and switch back to updating the model once the rat starts to move again.

III-3. Critical role of spontaneous activity for performing cognitive tasks

Xiaowei Gu XWGU@ION.AC.CN
Chengyu Li TONYLICY@ION.AC.CN

Chinese Academy of Sciences

Brain is not silent even without sensory input or motor output and internally generated neural activity was widely observed. Moreover, spontaneous activity was demonstrated to be correlated with evoked response to external stimuli or behavioral performance. However, the causal role of spontaneous activity in cognitive behaviors remains unclear. In our published work, we have demonstrated that the delay-period activity of medial prefrontal cortex (mPFC) was critical for information maintenance in learning but not well-trained phase of a working memory (WM) task. However, the functional significance of neural activity in other periods is still elusive. To address these questions, in the current unpublished study we optogenetically manipulated spontaneous activity of pyramidal neurons in mouse mPFC, during the baseline period before sensory delivery. We performed optogenetic experiments in two different WM tasks: a delayed non-match-to sample (DNMS, Fig. 2A) and a delayed paired associative learning (DPAL, Fig. 1A) tasks. In both tasks, performance of mice was significantly impaired after spontaneous activity was suppressed or enhanced optogenetically during the learning phase, but not after the mice were well trained (Fig. 1B and 1C). The behavioral impairment was manifested as increased miss or false choice rates. To understand the mechanism of spontaneous activity underlying behavioral performance, we further recorded mPFC neural activity during leaning of WM tasks (Fig. 2B). Trial-by-trial variation of spontaneous neural activity in baseline period was correlated with the performance engagement in tasks. Furthermore, mean firing rates of baseline period were modulated through learning in a single day (Fig. 2C). In addition, endogenous baseline-period mPFC activity showed correlation with behavioral performance and memory retention during delay period (Fig. 2D and 2E). Thus, behavioral-context regulated mPFC spontaneous activity in baseline period is critical for

learning of WM tasks. Our findings causally revealed the importance of spontaneous neural activities for cognitive behaviors.

III-4. Distinct timescales of cortical reorganization in a long-term learning task

Xiao Zhou[1] ZHOUXIAO@CMU.EDU
Rex N Tien[2] RNT9@PITT.EDU
Steven Chase[1] SCHASE@CMU.EDU

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

Skill learning is associated with a functional reorganization of cortical neural activity. However, the link between changes in neural activity and the concurrent behavioral improvements is not well understood. This is primarily because in most tasks it is difficult to interpret the behavioral impact of particular changes in neural activity. Here we leveraged a brain-computer interface (BCI) learning paradigm to determine how long-term practice (several weeks) leads to skill acquisition in a BCI movement task. In a BCI, the experimenter provides the subject a definitive mapping between neural activity and the movement of an effector (in our case, a computer cursor). Each new BCI mapping thus provides the subject a new tool that must be mastered through continued practice. One advantage of this paradigm is that the mappings of individual neurons to cursor movement can be manipulated to test the specificity of adaptive responses. We found two distinct timescales of cortical reorganization during long-term learning of the new BCI mapping. In the first phase, rapid, coordinated changes in activity across all neurons act to quickly (within one day) reduce behavioral errors during task performance. In the second phase, long-timescale changes in the tuning of individual neurons act to gradually improve the efficiency of the movement over weeks of practice.

III-5. Do people think like computers?

Bas van Opheusden SVO213@NYU.EDU
Gianni Galbiati GVG218@NYU.EDU
Zahy Bnaya ZAHY.BNAYA@GMAIL.COM
Wei Ji Ma WEIJIMA@NYU.EDU

New York University

Decision-making tasks such as negotiating, developing military strategy and career planning require a sequence of choices between multiple alternatives at each step, causing a combinatorial explosion of the decision tree. Much of our intuition about how people explore large decision trees comes from artificial intelligence, primarily computers playing board games. Common AI agents for board games use a search algorithm guided by a heuristic function that assigns values to board positions. Do people think the same way? To find out, we collected data from human subjects playing a generalized version of tic-tac-toe. In the first experiment, subjects played against each other. In the second, subjects played against AI opponents, performed a two-alternative forced-choice task, and evaluated board positions. In the third, we tracked subjects' eye movements while they played against AI opponents. We model subjects' choices with an AI agent who uses a value function with simple weighted features and constructs a decision tree with best-first search and value-based pruning. To capture mistakes, we include noise and stochastic feature dropping in the value function, and a lapse rate. Our model predicts individual subjects' choices on single trials better than chance or control models generated by lesioning model components. The model generalizes across tasks. After fitting the model to subjects' choices in games against AI opponents, we can predict their choices in the 2AFC (accuracy: $55.7 \pm 0.7\%$) and evaluation tasks (correlation coefficient predicted/observed: 0.48 ± 0.05). Subjects' fixation time on each square correlates with how often the search algorithm visits that square (correlation coefficient: 0.546 ± 0.006). From these results, we conclude that AI-inspired models can indeed capture people's decision-making process in this game.

III-6. A flexible research platform for sensory substitution in a 3D environment

Yang Liu YOUNGLEOEL@OUTLOOK.COM
Markus Meister MEISTER@CALTECH.EDU

California Institute of Technology

The goal of our project is to help blind people to navigate. Our approach is to translate information about their surroundings into sounds. These psychophysics experiments will also reveal aspects of crossmodal perception and the formation of internal images. There are 45,000,000 blind people worldwide. Several approaches are being developed to directly restore visual function, but these are invasive and not yet available for clinical use. As a non-invasive alternative, sensory substitution methods encode scene features into stimuli of other sensory modalities, such as sound ('sonification'). Existing devices encode low-level scene features, such as the distance to an obstacle, or the intensity pattern in a static image, and have met with limited success. We propose that technologies exist today or in the near future that can acquire all the knowledge required to navigate a particular scene. For example, driverless cars already combine GPS mapping and real-time computer vision to a high-level and actionable representation of the environment. Thus the remaining challenge is to efficiently communicate this high-level knowledge to blind subjects. We have created an experimental platform to compare diverse strategies for sonification, using readily available sighted subjects. The subject wears virtual reality goggles and navigates in a 3D virtual environment, which provides visual experience as well as an acoustic sonification. After a training period we remove the visual input and the subject navigates based on acoustic signals alone. Any given method of sonification can be assessed by the subject's final performance and training time required. In ongoing studies we are evaluating a strategy in which each actionable object calls out its name. The subject experiences the sound as coming from the object's location. Obstacles to be avoided emit a generic noise signal. Naive subjects have learned to perform navigational tasks in this environment during the first day.

III-7. Optimal prediction in natural gaze behavior

SangWook Lee LEEALEX@UCHICAGO.EDU
Stephanie Palmer SEPALMER@UCHICAGO.EDU
Leslie Osborne OSBORNE@UCHICAGO.EDU

The University of Chicago

Delays in sensory processing give rise to a lag between a stimulus and the organism's reaction. This presents a particular challenge to tracking behaviors like smooth pursuit where a difference in eye and target motion creates image motion blur on the retina. One strategy that might compensate for processing delays is to extrapolate the future target position and make anticipatory eye movements. We recorded eye position in humans and monkeys engaged in tracking tasks to test for predictive information in eye movements. For humans, we created a video game task based on Atari Pong in which the subject moves a paddle to keep a ball bounding within an arena. We controlled the level of predictability, keeping collisions elastic or adding stochasticity. Subjects also watched 'movies' of the Pong game being played, or of a ball bouncing within an enclosed arena (no paddle). We employed a third task for monkeys based on the classic step-ramp pursuit paradigm in which the target jumped and translated in a pseudorandom fashion. On this double step-ramp task, target motion was predictable over a much shorter time scale than in the Pong task ($\sim 500ms$ vs $>2s$). We computed the mutual information of eye and target position as a function of time shift. We find substantial predictive information in all tasks in both humans and monkeys, but there is a striking difference between active game play and the other tracking-based tasks. In active Pong, the mutual information in the eye position is very near the bound determined by the predictive information in the target itself, $I(T(t), T(t + \delta t)$, thus all that is predictable about the target is incorporated into behavior. Eye-target information peaks at zero delay, such that prediction compensates for processing delays. In contrast, watching a movie of the Pong game or tracking

III-8. Information seeking is driven by two types of uncertainty

Ethan Bromberg-Martin[1,2] NEUROETHAN@GMAIL.COM
Michael Platt[3] MPLATT@MAIL.MED.UPENN.EDU
David Barack[1] DBARACK@GMAIL.COM

[1]Columbia University
[2]Kavli Institute for Brain Science
[3]University of Pennsylvania

Conventional theories of reinforcement learning explain how we choose actions to gain rewards, but we can also choose actions to gain information about rewards, thereby reducing our uncertainty about future events. This preference for information, referred to as information seeking in economics and observing behavior in psychology, is evolutionarily conserved, found in rats, birds, monkeys, and humans. However, a fundamental question remains unanswered: what type of uncertainty drives information seeking? Here we test two hypotheses about the uncertainty that motivates information seeking. According to entropy-reduction theory, animals are driven to gather information to reduce the entropy of the distribution of future rewards. According to variance-reduction theory, animals are driven to gather information to reduce the variance of their estimates of reward value. To test these hypotheses, we analyzed data from monkeys performing two behavioral tasks: an information choice task in which animals were able to view advance information about the outcome of an impending gamble, and a naturalistic traplining task in which animals received information about the sequence of future rewards while foraging from multiple food sources. Critically, both tasks offered animals actions with varying yields of entropy-reduction and variance-reduction, allowing us to dissociate their effects on behavior. Unexpectedly, we found that information seeking behavior in these tasks followed distinct informational motives. In the information choice task, animals exclusively tracked reduction in variance, not entropy. In the traplining task, however, animals tracked both types of uncertainty reduction. Notably, traplining allowed animals to gather information about sequences of future rewards, and they were motivated by reductions in uncertainty about these sequences. Thus, our findings suggest that information seeking is not beholden to a single form of uncertainty. Just as taste, calories, and nutrient content motivate food seeking, multiple informational features of natural environments motivate information seeking.

III-9. Attention-related BOLD modulation with and without superior colliculus inactivation

Richard Krauzlis[1] RICHARD.KRAUZLIS@NIH.GOV
Amarender Bogadhi[1] AMARENDER.BOGADHI@NIH.GOV
Anil Bollimunta[1] ANIL.BOLLIMUNTA@NIH.GOV
David Leopold[2] LEOPOLDD@MAIL.NIH.GOV

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

The superior colliculus (SC) has been demonstrated to play a causal role in the control of visual attention, but acts through mechanisms that are dissociable from the well-known signatures of attention in visual cortex (Zenon & Krauzlis 2012). To identify which brain regions might be part of the SC attention network, we have conducted fMRI in a monkey performing a spatial attention task, with and without SC inactivation. Imaging runs contained three types of blocks: Baseline (B), Foveal Attention (FA) and Peripheral Attention (PA) blocks. In B block trials, the relevant stimulus was a central fixation point that dimmed at randomized times. FA block trials were similar to B block trials but added a peripheral motion-change stimulus as an irrelevant distracter. In PA block trials, the fixation point did not dim and the peripheral motion-change was the relevant stimulus. The monkey's task was to maintain central fixation and to report the relevant stimulus change (fixation dimming in B & FA blocks, peripheral motion change in PA blocks) by releasing a lever to get a juice reward. In each anatomically defined ROI, we identified voxels with significant differences between PA and B blocks, and calculated an attention modulation index (AMI)

for each of these voxels based on the %change in BOLD for PA and FA blocks. During SC inactivation, the AMIs for most attention-related cortical areas (V1, MT/MST, LIP and FEF) remained normal, despite the presence of significant deficits in the attention task. However, attention-related modulation was eliminated in one cortical area (FST), and reduced in subcortical regions (caudate, pulvinar). The results suggest that the SC contributes to attention through a circuit involving cortical area FST and regions of the thalamus and basal ganglia, highlighting the possible role of subcortical decision-making mechanisms in the control of attention.

III-10. Selective attention suppresses responses to competing distractors in auditory cortex

Zachary Schwartz SCHWARZA@OHSU.EDU
Stephen David DAVIDS@OHSU.EDU

Oregon Health and Science University

Auditory selective attention is required for extracting behaviorally relevant information from crowded acoustic environments. Studies in humans have identified changes in neural population activity (MEG, ECoG) associated with selective attention, but single-unit correlates of attention in animals are not well characterized. Most animal studies have compared neural activity between passive listening and a single behavior condition or between behaviors with different structure, making it difficult to isolate effects of selective attention from other changes in behavioral state. To address this problem, we developed a behavior that dissociates selective attention from general task engagement. We recorded single-unit activity in primary auditory cortex of ferrets during two attention conditions and during passive listening. These data revealed that the effects of attention and task engagement on neural activity are separable. Changing selective attention produced a systematic decrease in activity evoked by distractors competing with the attended target. On the other hand, engaging in the task increased spontaneous spike rates but did not systematically affect response gain. The attention-mediated suppression of distractor responses produced enhanced neural discriminability between the distractor and target in the attended location.

III-11. Laminar organization of attentional modulation in visual area V4

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

Salk Institute for Biological Studies

In order to flexibly adapt to behavioral demands, the brain needs to rapidly modulate the operating mode of the underlying cortical circuits and thereby control the way information is routed. Tasks that control attention provide a powerful way to manipulate cortical information processing. Traditional single unit electrophysiology has provided key insights into the probable neural mechanisms underlying attention in area V4 [1], including modulation of mean firing rate and reduction in correlated variability among pairs of simultaneously recorded units [2,3]. These changes in neuronal response are thought to result from feedback signals generated in attentional control centers, which impinge on the laminar circuits of the visual cortices. To gain insight into the role of laminar circuits in this transformation one must measure the laminar profile of attentional modulation. This requires laminar electrodes to be positioned perpendicular to the cortical surface. In V4 this is a challenge because it straddles a narrow gyrus, with only a narrow strip of cortex that lays flat beneath the calvarium. Here, we replaced the native dura with an optically clear artificial dura. This allowed us to very precisely target laminar array electrodes to pass down the cortical column (Fig 1A). We obtained receptive field maps that were well aligned (Fig 1B), and clear current-source-density (CSD) maps (Fig 1C). We used CSD maps to distinguish the laminar position—granular (G), infra-granular (IG), supra-granular (SG)—of each channel. We recorded activity while monkeys performed an attention- demanding task (Fig 2). We found differential modulation of sensory gain, response variability and correlated variability across layers, with the highest modulation in G and SG layers (Fig 3). Neurons in SG layers transmit signals to higher-order areas. Our finding, that attentional modulation is strong in the superficial layers, supports the idea that attention improves signals that are transmitted to higher-order areas.

III-12. Evidence accumulation and change rate inference in dynamic environments

Adrian Radillo[1] ADRIAN.RADILLO@GMAIL.COM
Alan Veliz-Cuba[2] AVELIZCUBA1@UDAYTON.EDU
Kresimir Josic[1] JOSIC@MATH.UH.EDU
Zachary Kilpatrick[1] ZPKILPAT@MATH.UH.EDU

[1]University of Houston
[2]University of Dayton

Humans and other animals make perceptual decisions based on noisy sensory input. The underlying decision processes are often modeled using evidence accumulators. Behavioral and electrophysiological data have provided evidence for such models in several cases. Recent studies focus on ecologically realistic situations in which the correct choice or the informative features of the stimulus change over time. Importantly, optimal evidence accumulation in changing environments requires discounting prior evidence at a rate determined by environmental volatility. Experiments suggest that humans can learn the rate of these changes to make choices nearly optimally. To explain these observations we extend previous accumulator models of decision making to the case of multiple choices with asymmetric, unknown transition rates between them. We show that an ideal observer can optimally infer, on-line, these transition rates, and accumulate evidence in order to make the best decision. When the number of choices is discrete, our algorithm is simpler than those proposed previously, as it does not require knowing the probability of all possible run lengths since the last change point. Moreover, we showed in a previous study that the optimal, nonlinear model is well approximated by a linear model. This motivates a physiologically plausible neural implementation for the current problem: We show that a Hebbian learning rule can shape interactions between multiple populations representing the different choices, allowing the network to integrate inputs nearly optimally. Our work therefore links statistical principles for optimal inference with stochastic neural rate models that can adapt to the environmental volatility to make optimal decisions in a changing environment.

III-13. Using social information to infer on one's own taste in the important case of temporal discounting

Michael Moutoussis[1] M.MOUTOUSSIS@UCL.AC.UK
Ray Dolan[2] R.DOLAN@UCL.AC.UK
Peter Dayan[3] DAYAN@GATSBY.UCL.AC.UK

[1]Wellcome Trust Centre for Neuroimaging at UCL
[2]Max Planck Computational Psychiatry Centre
[3]Gatsby Computational Neuroscience Unit, UCL

In this work we provide evidence that shifting preferences upon observing the choices of others largely reflects Bayesian inference. Inference is performed on the parameters defining one's own utility function. These parameters can in turn be said to describe the decision-maker's individual tastes. Temporal discounting parameters are a paradigmatic example of such tastes, very important in computational psychiatry and economics. High values of discounting parameters are associated with several psychiatric disorders, lower IQ and poverty. Temporal discounting tasks have good psychometric properties, leading to a well-established hyperbolic model. Somewhat surprisingly, individual preferences shift in the face of observing the preferences of others even if this shifting is not itself rewarded. The computational basis of this is unclear. We propose a new model of tastes as (uncertain)

Bayesian beliefs, allowing for one's tastes to be updated through observing choices made by other, epistemically-trusted people. Such random tastes may form an important part of random-utility based choices for the individual. If uncertainty is thus reflected in choice variability, then a key signature of our account is that baseline choice variability should correlate with the magnitude of apparent preference change. We examined discounting in a novel community study of 740 young people who made choices between a smaller but immediate, versus a larger but delayed, reward. They did this both before and after learning about the preferences of a 'partner'. We found that participants displayed considerable choice variability. The degree of preference shift upon learning about the partner was correlated with baseline choice variability, lending support to our Bayesian account. Younger people were influenced by others more than older ones, and this was explained by the former being less certain about their own preferences. These findings raise the possibility of tastes being subject to Bayesian inference in other important domains.

III-14. Why are we so slow to decide between two good options?

Satohiro Tajima SATOHIRO.TAJIMA@GMAIL.COM
Jan Drugowitsch JDRUGO@GMAIL.COM
Alexandre Pouget ALEX.POUGET@GMAIL.COM

University of Geneva

For decades now, normative theories of perceptual decisions, and their implementation as drift diffusion models (DDMs), have significantly improved our understanding of human and animal behavior and the underlying neural processes (Ratcliff, 1978; Gold & Shadlen, 2001; Drugowitsch et al., 2012). While similar processes seem to govern value-based decisions (e.g., Krajbich, Armel & Rangel, 2010), we still lack the theoretical understanding of why this ought to be the case. In particular, such processes predict that decisions are solely governed by the difference in reward across choice options, independent of their absolute reward magnitudes, such that decisions are slowest if these options are similar in values. This is counter-intuitive, as one would expect decisions between two similar high-valued options, yielding certain high rewards, to be performed more quickly. Here, we show that, similar to perceptual decisions, DDMs indeed optimize the reward rate (i.e. per unit of time) for value-based decisions. Such optimal decisions require the models' decision boundaries to collapse over time, and to depend on the a-priori knowledge about reward contingencies. In particular, if rewards are high on average, decisions are still governed by reward differences, but with fast collapsing decision boundaries leading to faster choices. If the goal of the subject is to maximize the reward per trial, as opposed to per unit of time, the optimal bound collapse is slower, leading once again to long reaction times for two similar high-valued options, as has been observed experimentally. Finally, in general, diffusion models only implement the optimal strategy under specific task assumptions, and cease to be optimal once we start relaxing these assumptions, by, for example, using nonlinear utility functions. Our findings thus provide the much-needed theory for value-based decisions, explain the apparent similarity to perceptual decisions, and predict conditions under which this similarity should break down.

III-15. Lucky rhythms in orbitofrontal cortex bias gambling decisions in humans

Sridevi Sarma[1] SRIDEVI.SARMA@GMAIL.COM
Kevin Kahn[1] FROSTYM288@GMAIL.COM
Matthew Kerr[1] MKERR10@JHMI.EDU
Jorge Gonzalez-Martinez[2] GONZALJ1@CCF.ORG
Hyun Joo Park PARKH@CCF.ORG
Juan Bulacio[2] BULACIJ@CCF.ORG
Matthew A Johnson[2] JOHNSOM@CCF.ORG
Susan Thompson THOMPSS16@CCF.ORG
Jaes Jones[2] JONESJ30@CCF.ORG
Vikram Chib VCHIB@JHU.EDU
John T Gale[2] GALEJ@CCF.ORG

[1]Johns Hopkins University
[2]Cleveland Clinic

It is well established that emotions influence our decisions, yet the neural basis of this biasing effect is not well understood. Here we directly recorded local field potentials from the orbitofrontal cortex (OFC) in human subjects performing a financial decision-making task. Our results are striking: we observed an increase in gamma-band (37-60Hz) oscillatory activity that selectively biased subject's to risk more virtual money on trials with the same expected payout for all possible decisions. Additionally, these gamma rhythms were linked to the subject's emotional state through how "lucky" the subject was on recent trials. These results suggest that the OFC plays a pivotal role in both the subject's emotional state and the subject's future decisions, and bridge two popular theories of the OFC, the emotional marker theory and response inhibition.

III-16. A theory of learning dynamics in perceptual decision-making

Christopher Baldassano[1] CHRISB@PRINCETON.EDU
Andrew M Saxe[2] ASAXE@FAS.HARVARD.EDU

[1]Princeton University
[2]Harvard University

Humans, monkeys, and even rats can make strikingly optimal decisions from a noisy stream of perceptual evidence. Yet perhaps even more remarkable is the fact that this optimal performance is not innate, but emerges through learning. How can we bridge the gap between our relatively well-developed models of optimal behavioral choice, which posit fixed neural structures and mechanisms, and this evident learning? Can choice behavior be accurately described at all points during task acquisition? Here we develop a theory of two alternative forced-choice (2AFC) experiments which describes the entire trajectory from baseline performance to near optimality, and links plasticity in the underlying neural circuitry to behavioral improvements. Using a linear recurrent neural network which receives biased noisy inputs and must generate a decision by a fixed deadline, we update every synapse after each decision trial using error-driven learning. Remarkably, the learning dynamics of the network reduce exactly to an Ornstein-Uhlenbeck (OU) process with time-dependent parameters, explaining how unconstrained learning in a generic population of neurons connected via modifiable synapses can yield behavior characterized by highly successful phenomenological models. By virtue of its OU formulation, the model makes a variety of readily-testable behavioral predictions, including the shifting shape of psychometric functions over time, the effect of manipulating the SNR of the inputs, and the effect of delivering information early versus late in a trial. The model yields several neural predictions, including that individual neurons will typically have mixed selectivity, while the population response will become increasingly correlated over the course of learning. More broadly, the model can be viewed as an extension of the fundamental drift diffusion model to incorporate learning dynamics, with diverse applications to the dynamics of learning in free response paradigms, with mixed-difficulty input signals, and in choices between more than two alternatives.

III-17. Serotonin delays switching by promoting active reward seeking in a probabilistic foraging task

Eran Lottem[1] ERAN.LOTTEM@NEURO.FCHAMPALIMAUD.ORG
Pietro Vertechi[1] PIETRO.VERTECHI@NEURO.FCHAMPALIMAUD.ORG
Matthijs N Oude Lohuis[1,2] MATTHIJS.OUDELOHUIS@NEURO.FCHAMPALIMAUD.ORG
Dhruba Banerjee[1,3] DHRUBA.BANERJEE@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[1] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]University of Amsterdam
[3]University of California, Irvine

The central neuromodulator serotonin (5-HT) is involved in impulse control through behavioral inhibition. Activation of serotonergic neurons was shown to promote patience in mice trained to wait for randomly delayed rewards, and inhibition of these neurons in a similar task resulted in premature leaving. These effects may be the result of an overall reduction in movement vigor following serotonin release. Alternatively, rather than attenuating behavioral outputs, serotonin may be modulating decision-making processes to favor certain types of behavior. To tease apart these alternatives we developed a simple foraging task in which mice nose-poked for rewards at two probabilistically depleting ports. Importantly, each trial was comprised of multiple successful and failed attempts at reward, thus making it possible to distinguish between the motor- and decision-related effects of serotonin. We found that the optogenetic activation of serotonergic neurons increased the average number of reward attempts before leaving (n=8 mice). Furthermore, stimulating mice after leaving one site and before reaching the other had little effect on movement. We found that behavior in this task could be parsimoniously modeled by a drift-diffusion process, in which successive outcomes are integrated until reaching a fixed bound that signals the decision to leave. In this framework, serotonin stimulation resulted in a reduction in the weight of individual attempts, thereby increasing the time taken to reach the threshold. Taken together, these results invite reconsideration of behavioral inhibition theory of serotonin effects in light of a computational perspective on decision-making.

III-18. Evidence accumulation detected in BOLD signal using slow perceptual decision making

Paul Krueger[1,2] PAUL.M.KRUEGER@BERKELEY.EDU
Marieke van Vugt[3] M.K.VAN.VUGT@RUG.NL
Patrick Simen[4] PSIMEN@OBERLIN.EDU
Leigh Nystrom[2] NYSTROM@PRINCETON.EDU
Philip Holmes[2] PHOLMES@MATH.PRINCETON.EDU
Jonathan Cohen[2] JDC@PRINCETON.EDU

[1]University of California, Berkeley
[2]Princeton University
[3]University of Groningen
[4]Oberlin College

The process of decision making can be described by mathematical models of evidence accumulation such as the drift diffusion model (DDM). We assessed whether neural dynamics of evidence accumulation could be observed in BOLD signal recorded during a perceptual decision making task. The hemodynamic response function acts as a low-pass filter on neural activity, which presents two challenges for measuring evidence accumulation. First, the temporal resolution of fMRI is low, and perceptual decisions are typically fast. To overcome this, we slowed down perceptual decision making dramatically using theoretical predictions of the drift diffusion model. Specifically, participants were penalized more harshly for incorrect responses than they were rewarded for correct responses. Second, because the BOLD signal is convolved with the low-pass hemodynamic response, it is very difficult to disentangle different signals. To address this, we improved our ability to distinguish BOLD activity related

to detection (modeled using a boxcar) from that related to integration (modeled using a ramp) by minimizing the collinearity of our GLM regressors. This was achieved by dissecting a boxcar into its two most orthogonal components: an ‘up-ramp’ and a ‘down-ramp.’ This gave our model the most power to distinguish integration (which we modeled using the up-ramp regressor) from detection (modeled as the sum of an up-ramp and down-ramp regressor). We also used a well-matched control condition in which stimuli and responses were similar to the experimental condition, but no evidence accumulation or decision making took place. We identified areas with significant ramping activity that was greater during the decision making task than the control task. These areas included bilateral insula, cerebellum, and bilateral middle occipital gyrus.

III-19. Attention and choice across domains

Ian Krajbich KRAJBICH@GMAIL.COM
Stephanie Smith SMITH.10264@OSU.EDU

The Ohio State University

People tend to choose items that draw their attention. Previous research has demonstrated clear relationships between attention and choice in a variety of different value-based settings. However, these relationships have been studied and modeled separately in each setting. Here we aimed to bridge these gaps by studying the consistency of the decision-making process across choice domains. We investigated both how the link between attention and choice varies within an individual across tasks and how the strength of that link correlates with other measures of attention. We employed four distinct choice tasks, including a binary food-choice task, a risky food-choice task, a risky monetary-choice task, and a social monetary-allocation task. After finishing these four choice tasks, subjects completed a final non-choice task to measure their attentional gradient. We modeled the data from these tasks by combining a utility-function approach from economics with both standard logistic regressions and a multi-attribute extension of the attentional drift diffusion model. We find that while looking time generally does not correlate with value/utility, it does predict choices. Across tasks we find a remarkably consistent pattern of results where a 0.5s looking-time advantage for one side translates to an increase in choice probability of $\sim 25\%$, and at indifference the probability that the last fixation is to the chosen item is $\sim 70\%$. At the individual level we find that subjects with strong attentional influences in one task are likely to have strong attentional influences in the other tasks. We also see a significant correlation between a subject’s average attentional influence (across tasks) and their attentional gradient. Across subjects then, we find clear relationships between the sharpness of attentional gradient, attentional weighting, and subsequent choice. Together, these connections provide support for a common attention-based decision-making process responsible for choice across domains.

III-20. What can we learn from choice probabilities in neural networks with feedback?

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

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

Feedback constitutes an important attribute of information processing in the brain. Here we use a simple model system to identify how neural feedback shapes neural variability and correlations with choice in time. For analytical tractability, we investigate a linear dynamical system with additive gaussian noise. The architecture of a recurrent network influences its information content, as well as the dynamics of choice correlations, (CCs, the correlation of neural activity with behavioral reports). We relate CC directly to the matrix of recurrent feedback, which determines the dynamics and structure of the network. Appropriate recurrence allows the network to integrate evidence over time. We derive an analytical solution for the optimal matrix of recurrent feedback weights, and compute how it affects the noise covariances in the network (Fig 1a). Our calculations reveal that if output noise is fed back into the network but the output firing rate is constrained, then the optimal feedforward gain is strongly enhanced while the feedback gain is reduced (Fig 1b). We also show how the vector of CCs for all neurons may change over time, depending on the structure of the recurrent feedback matrix. These relationships allows one to infer the timescale and relevant patterns of feedback from measured CCs. The predictive coding paradigm also uses feedback as a key component. We find that internal models (priors on the stimulus) alter choice probabilities for neurons encoding expectation or residual errors (Fig 1c). We show analytically and confirm by simulations how the parameters of the internal model are reflected in choice probabilities when a network implements predictive coding.

III-21. Deep networks reveal the structure of motor control in sensorimotor cortex during speech production

Jesse Livezey[1,2] JESSE.LIVEZEY@BERKELEY.EDU
Gopala Anumanchipalli[3] GOPALAKRISHNA.ANUMANCHIPALLI@UCSF.EDU
Brian Cheung[1] BCHEUNG@BERKELEY.EDU
Mr. Prabhat[4] PRABHAT@LBL.GOV
Michael DeWeese[1] DEWEESE@BERKELEY.EDU
Edward Chang[3] EDWARD.CHANG@UCSF.EDU
Kristofer Bouchard[4] KEBOUCHARD@LBL.GOV

[1]University of California, Berkeley
[2]Redwood Center for Theoretical Neuroscience
[3]University of California, San Francisco
[4]Lawrence Berkeley National Laboratory

Vocal articulation is a complex task requiring the orchestration of several parts of the vocal tract. Understanding the organization of cortical signals that control the vocal articulators during the production of simple speech segments (consonants, vowels) is key to our understanding of how the brain produces the complex sequences that compose spoken language. As with many ethological behaviors, speech production can be decomposed into features that are sequentially blended to generate complex sequences. How cortical activity is organized during the production of such complex behaviors is poorly understood. To explore how the human brain produces speech, cortical activity from the surface of speech sensorimotor cortex was recorded using high-density electrocorticography (ECoG) in neurosurgical patients during the production of a diverse set of American English consonant-vowel (CV) syllables. In contrast to previous efforts, which used linear methods on trial-averaged data at specific time points, we trained deep networks (DNNs) to classify speech from single-trial spatiotemporal patterns of activity. We found that the predictions made by the DNNs exhibited a rich, hierarchical organization, simultaneously structured by consonant features and vowels. Comparison to linear networks indicated that the non-linear processing of DNNs was critical to revealing this structure. Additionally, DNNs exceeded state-of-the-art performance classifying speech from human sensorimotor cortex, with the best performance of 38.7% for syllable classification (23x over chance). Interestingly, we found that classifying CV's resulted in improved accuracy (relative to chance) compared to classifying either of the individual constituents (C's, V's). This demonstrates that, during speech production, vSMC neural activity is structured on longer time-scales than individual phoneme durations. Together, these results demonstrate the power of DNNs for extracting the structure of neural activity underlying of complex motor tasks, such as speech, from noisy single-trial signals. Therefore, DNNs may be a powerful tool for understanding brain computations in general.

III-22. Interactive control of diverse complex characters with artificial neural networks

Igor Mordatch[1] IGOR.MORDATCH@GMAIL.COM
Kendall Lowrey[2] LOWREY@CS.WASHINGTON.EDU
Galen Andrew[2] GALEN@CS.WASHINGTON.EDU
Pieter Abbeel[1] PABBEEL@CS.BERKELEY.EDU
Emo Todorov[2] TODOROV@CS.WASHINGTON.EDU

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

Interactive real-time controllers that are capable of generating complex, stable and realistic movements have many potential applications to serve as computational models in neuroscience and biomechanics. We present a method for training artificial recurrent neural networks to act as near-optimal feedback controllers. It is able to generate behaviours for a range of virtual dynamical systems and tasks – swimming, flying, biped and quadruped walking with different body morphologies. It does not require motion capture or task-specific features or state machines. The controller is simply an artificial neural network, having a large number of feed-forward units that learn elaborate state-action mappings, and a small number of recurrent units that implement memory states beyond the physical system state. The action generated by the network is defined as predicted future velocity. Thus the network is not learning a traditional control policy, but rather the dynamics under an implicit policy (in essence making predictions about the future states). Essential practical implementation features of the method include interleaving supervised learning with trajectory optimization, injecting noise during training, training for unexpected changes in the task specification, and using the trajectory optimizer to obtain optimal feedback gains in addition to optimal actions. One exciting application of this to biological movement control is to investigate how animal motion is influenced by considerations of mental effort and capacity (as opposed to considering just physical metabolic energy expenditure), as well as sensory noise and delays. Our current ongoing work also includes incorporation of complex high-dimensional sensory modalities, such as vision and touch to the movement policies and inspecting how they influence behaviour.

III-23. The inevitability of probability: probabilistic inference in generic neural networks

Emin Orhan EORHAN@CNS.NYU.EDU
Wei Ji Ma WEIJIMA@NYU.EDU

New York University

Animals have been shown to perform near-optimal probabilistic inference in a wide range of psychophysical tasks, from causal inference to cue combination to visual search. On the face of it, this is surprising because optimal probabilistic inference in each case is associated with highly non-trivial behavioral strategies. Yet, typically animals receive little to no feedback during most of these tasks and the received feedback is not explicitly probabilistic in nature. How can animals learn such non-trivial behavioral strategies from scarce non-probabilistic feedback? We show that generic feed-forward and recurrent neural networks trained with a relatively small number of non-probabilistic examples using simple error-based learning rules can perform near-optimal probabilistic inference in standard psychophysical tasks. For inference problems involving continuous variables, the hidden layer of trained networks with a linear readout develop a novel sparsity-based probabilistic population code. We show that the networks do not have to be very finely tuned as random networks with only output weights trained can perform probabilistic inference too, albeit usually much less efficiently than fully trained networks. For fully trained networks, performance asymptotes at very small network sizes ($\sim O(10)$ hidden units for most tasks) due to the low computational complexity of the typical probabilistic psychophysical tasks. For the same reason, in many cases, the trained networks also display remarkable generalization to stimulus conditions not seen during training. Our results suggest that far from being difficult to learn, optimal probabilistic inference in standard psychophysical

tasks emerges naturally and robustly in generic neural networks trained with error-based learning rules, even when neither the training objective nor the training examples are explicitly probabilistic.

III-24. Neural classifiers with limited connectivity and local recurrent readouts

Lyudmila Kushnir KUSHNIRLV@GMAIL.COM
Stefano Fusi SF2237@CUMC.COLUMBIA.EDU
Columbia University

For many neural network models that are based on perceptrons, the number of activity patterns that can be classified is limited by the number of plastic connections that each neuron receives, even when the total number of neurons is much larger. This poses the problem of how the biological brain can take advantage of its huge number of neurons given that the connectivity is extremely sparse, especially when long range connections are considered. One possible way to overcome this limitation in the case of feed-forward networks is to combine multiple perceptrons together, as in committee machines. The number of classifiable random patterns would then grow linearly with the number of perceptrons, even when each perceptron has limited connectivity. However, the problem is moved to the downstream readout neurons, which would need a number of connections that is as large as the number of perceptrons. Here we propose a different approach in which the readout is implemented by connecting multiple perceptrons in a recurrent attractor neural network. We show with analytical calculations that the number of random classifiable patterns can grow unboundedly with the number of perceptrons, even when the connectivity of each perceptron remains finite. Most importantly both the recurrent connectivity and the connectivity of a downstream readout are also finite. Our solution requires that the input neural representations are sparse, which is surprising given the limited connectivity constraint. Our study shows that feed-forward neural classifiers with numerous long range connections connecting different layers can be replaced by networks with sparse long range connectivity and local recurrent connectivity without sacrificing the classification performance. Our strategy could be used in the future to design more general scalable network architectures with limited connectivity, which resemble more closely brain neural circuits dominated by recurrent connectivity.

III-25. Time-complexity and accuracy in neural winner-take-all computation

Rishidev Chaudhuri RCHAUDHURI@AUSTIN.UTEXAS.EDU
Birgit Kriener BIRGIT@MAIL.CLM.UTEXAS.EDU
Ila R Fiete ILAFIETE@MAIL.CLM.UTEXAS.EDU
The University of Texas at Austin

Selecting the largest number in a set is a ubiquitous computational problem, for machines and brains. For instance, to select an option associated with the highest payoff amongst many competing possibilities can involve solving a version of the problem. In neural models, winner-take-all (WTA) networks perform this computation on a set of inputs to different neurons, amplifying the response of the most strongly driven neuron and suppressing the rest. It is usually noted that WTA convergence in rate-based networks is slow. Spike time-based WTA networks can converge much more rapidly, but have mostly been studied with noiseless inputs. Here we study convergence time in WTA networks, focusing on the scaling with task complexity. We also examine the accuracy of the computation. We show analytically and numerically that for N inputs drawn at random from a real interval, convergence time scales only logarithmically in N. Closer analysis reveals that this Log[N] scaling arises from the difference in size of the two largest inputs. If this gap is held fixed, convergence time becomes independent of N, growing only as the gap shrinks, i.e., with the desired resolution of the task. By contrast, finding the maximum element (max) in an array has a time-complexity of O(N) in computer science, independent of resolution. In the presence of ongoing input noise, there is no strict WTA convergence. If convergence is defined by a threshold crossing, noise contributes to a speed-accuracy tradeoff. Noise-driven fluctuations can speed up selection of a winner but also

lead to erroneous selection of a weaker average input. We obtain speed- accuracy tradeoff curves for Gaussian and Poisson noise. These analyses reveal that temporal WTA schemes are subject to the same tradeoff, and that parallel neural WTA networks have better scaling on noisy max than serial strategies from computer science.

III-26. Deep convolutional neural network models of the retinal response to natural scenes

Lane McIntosh LMCINTOSH@STANFORD.EDU
Niru Maheswaranathan NIRUM@STANFORD.EDU
Aran Nayebi ANAYEBI@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Stephen Baccus BACCUS@STANFORD.EDU

Stanford University

A central challenge in sensory neuroscience involves understanding neural computations and circuit mechanisms underlying responses to ethologically relevant, natural stimuli. However, the ubiquity of cascaded nonlinear processes like synaptic transmission and spiking dynamics in multilayered circuits has presented significant obstacles to the goal of learning accurate computational models of circuit responses to natural stimuli from neural recordings. To address this, we employ deep convolutional neural networks (CNNs), which demonstrate success at many pattern recognition tasks (LeCun et al. 2015). These models cascade multiple layers of filtering and rectification—exactly the elementary computational building blocks thought to underlie complex functional responses of sensory circuits. Previous work utilized these models to understand properties in IT cortex (Yamins et. al. 2013), but not in early sensory areas where knowledge of neural circuitry can provide important validation for such models. We demonstrate that CNNs are considerably more accurate at capturing retinal responses to held-out natural scenes than linear-nonlinear (LN) models and related models, such as generalized linear models (GLMs). Furthermore, we find CNNs generalize significantly better across classes of stimuli (white noise vs. naturals scenes) they were not trained on. Remarkably, analysis of these CNNs reveals internal units selective for visual features on the same small spatial scale as the main excitatory interneurons of the retina, bipolar cells. Moreover, probing the model with reversing gratings, paired flashes, and contrast steps reveals that the CNN learns nonlinear retinal response properties such as frequency doubling and adaptation, even though the CNNs were not trained on such stimuli. Overall, this work demonstrates the power of CNNs to not only accurately capture sensory circuit responses to natural scenes, but also uncover the circuit's internal structure and function. Moreover, our methods can be readily generalized to other sensory modalities and stimulus ensembles.

III-27. Interactions between circuit architecture and plasticity in a closed-loop system

Hannah Payne[1] HPAYNE@STANFORD.EDU
Jennifer Raymond[1] JENR@STANFORD.EDU
Mark Goldman[2] MSGOLDMAN@UCDAVIS.EDU

[1]Stanford University
[2]University of California, Davis

Plasticity at specific synapses is thought to underlie memory, and can be probed in the intact animal by recording changes in neural activity during behavior. However, identifying sites of plasticity has been a challenge, because observed changes in neural activity do not necessarily reflect plasticity of inputs to that neuron, particularly in circuits with feedback loops, or in closed-loop paradigms in which an animal's actions affect its sensory input. To analyze the interaction between feedback loops and plasticity, we fit a set of firing-rate models to neural and behavioral data from cerebellum-dependent vestibulo-ocular reflex (VOR) learning, which calibrates the eye move-

ment responses to vestibular stimuli. Despite the relative simplicity of the circuits mediating this task, previous studies disagree on the loci of plasticity underlying the learned changes in behavior, and make opposite predictions about the sign of synaptic changes at a given locus. We compared models of the VOR circuit that differed in the strength of positive feedback conveyed to the cerebellum through an efference copy of motor commands, the key parameter distinguishing previous models. All models were fit to both neural and behavioral data. The goodness-of-fit was similar with and without strong positive feedback. However, models with weak to moderate positive feedback predicted LTD of vestibular inputs to Purkinje cells, whereas models with strong positive feedback predicted LTP. All models predicted LTP in non-cerebellar pathways. Surprisingly, all models with LTD of vestibular inputs exhibited a paradoxical increase in modulation of neural activity during a closed-loop paradigm intended to isolate the vestibular inputs. Finally, we used transient electrical and optogenetic perturbations of Purkinje cells to test different models' predictions. The results demonstrate how systematic comparison of circuit models with varying architectures can constrain the sites and direction of neural plasticity underlying behavior.

III-28. (Dis)inhibition as a binary switch of excitatory synaptic plasticity

Katharina Wilmes[1] KATHA@BCCN-BERLIN.DE
Jan-Hendrik Schleimer[1] JH.SCHLEIMER@HU-BERLIN.DE
Henning Sprekeler[2] H.SPREKELER@TU-BERLIN.DE
Susanne Schreiber[1] S.SCHREIBER@HU-BERLIN.DE

[1]Humboldt-Universitaet zu Berlin
[2]Technische Universitaet Berlin

Synaptic plasticity shapes functional circuits in the brain. To ensure the stability of established circuits, plasticity of synaptic connections should be controled. Currently, inhibition is discussed as a mechanism to selectively switch off plasticity of excitatory synapses (Bar-Ilan et al., 2012; Mullner et al., 2015). In particular, inhibition has been shown to block coincidence signals required for Hebbian forms of plasticity: backpropagating action potentials (bAPs) and calcium spikes. Such a regulatory mechanism, however, also requires that the forward-directed information flow via excitatory postsynaptic potentials (EPSPs) is preserved. Here, we address the question whether it is possible to cancel bAPs while still maintaining the forward-directed information flow from dendritic excitatory synapses to the soma. Our analysis of physiologically-constrained neuron models shows that with appropriate timing, EPSPs can indeed trigger somatic action potentials although the corresponding backpropagating signals are blocked. A control of bAPs and calcium spikes is possible in an all-or-none manner, enabling a binary switch of coincidence signals and hence plasticity. The required temporal precision for cancelation of distal calcium spikes is lower than that required for the annihilation of bAPs (1ms). We demonstrate that this seemingly high precision can be provided by a common feedforward inhibitory circuit. We further propose a mechanism that allows to learn the appropriate strength and timing of inhibition needed to exert control over excitatory plasticity: an anti-Hebbian learning rule for inhibitory synapses. This rule ensures a tight onset of inhibition with respect to the postsynaptic action potential and still allows the forward-directed EPSPs to pass. Interestingly, the proposed inhibitory learning mechanism is self-limiting, terminating once cancelation of bAP occurs. Our study provides experimentally testable predictions and demonstrates that the inhibitory switch of excitatory plasticity can serve as a robust mechanism that selectively regulates plasticity in functional circuits.

III-29. Intrinsic plasticity for optimal learning of variable stimulus intensities

Cristina Savin CSAVIN@IST.AC.AT
Travis Monk TRAVIS.MONK@UNI-OLDENBURG.DE
Joerg Luecke JOERG.LUECKE@UNI-OLDENBURG.DE

IST Austria

In many situations the meaning of a stimulus is the same despite fluctuations in its overall strength. A visual scene's content does not depend on light intensity, or a word utterance should be recognised irrespective of its loudness. Nonetheless, gain fluctuations are an integral part of the input statistics and they can help differentiate between stimuli. In the visual domain, for instance objects of the same class are likely to have similar surface properties, resulting in a distinct distribution of light intensities. Light intensities can therefore help identify objects. The neural underpinnings of such computation are unclear. Existing models discard gain information by ad hoc preprocessing (Nessler, 2009; Keck, 2012) or by divisive normalisation (Schwarz, 2001) before learning the input statistics from normalized data. Overall, it is unknown how neural circuits can robustly extract statistical regularities in their inputs when the overall intensity of stimuli is variable. Here we develop a principled account of unsupervised learning in the face of gain variations. We introduce a novel generative mixture model (Product-Poisson-Gamma) that explicitly models the statistics of stimulus intensity, and we derive a biologically-plausible neural circuit implementation for inference and learning in this model. We find that explicitly taking into account gain variations improves the robustness of unsupervised learning, as differences in input strength help distinguish between classes with similar features but different gain statistics. From a biological perspective, the derived neural circuit, in which feature-sensitive neurons are equipped with a specific form of intrinsic plasticity (IP), provides novel insights into the interaction between Hebbian and IP during learning. Furthermore, our results imply that neural excitability reflects nontrivial input statistics, in particular the intensity of the features to which a neuron is sensitive.

III-30. A three-state model helps to find anatomical correlates of perceptual learning

Rohan Gala[1] R.GALA@NEU.EDU
Daniel Lebrecht[2] DANIEL.LEBRECHT@UNIGE.CH
Anthony Holtmaat[2] ANTHONY.HOLTMAAT@UNIGE.CH
Armen Stepanyants[1] A.STEPANYANTS@NEU.EDU

[1]Northeastern University
[2]University of Geneva

In this study, we attempt to uncover the link between structural plasticity and perceptual learning. To this end, we image a sparse population of mouse somatosensory neurons which co-express channelrhodipsin-2 (ChR2) and two fluorescent labels, GFP and Synaptophysin-TdTomato. We train the mice to report the optical micro-stimulation evoked activity in ChR2-expressing neurons. During the two-month span of the experiment, labeled axons are imaged in vivo with two-photon microscopy at 4-day intervals, and the size changes of en passant boutons are used as proxies for circuit modification. Quantifying minute changes in time-lapse image stacks in a reliable and high-throughput manner presents a major obstacle. To overcome this challenge, we developed a semi-automated methodology and generated a dataset consisting of more than 50,000 putative boutons located on 115 axons from 5 animals. In order to find structural correlates of learning, we consider a three-state model in which individual boutons can transition among the following states: not present, learning state, and memory state. The learning and memory states in the model are hidden states, not necessarily related to bouton intensities. By fitting the three-state model to the data, we show that animals that learned the perceptual task well had a significantly higher fraction of boutons in the learning state, as compared to the animals that fail to achieve above-chance level of performance.

III-31. Serotonin and dopamine neurons signal distinct prediction errors during reversal learning

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

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

Serotonin is an important neuromodulator implicated in the regulation of many physiological and cognitive processes, but its primary role in behavior is still not understood. Serotonin is thought to oppose and / or complement dopamine's role in learning and behavioral control: while phasic dopamine release invigorates behavior and drives learning by signaling reward prediction errors, serotonin has been implicated in behavioral inhibition and aversive processing. Specifically, serotonin is thought to be necessary for preventing perseverative responses in changing environments. However, whether or how serotonin neurons signal such changes is not clear. To study the activity of serotonergic neurons in behaving mice and to compare it with dopaminergic activity, we used genetically encoded calcium indicators and fiber photometry to monitor neural activity from these populations. We used a reversal learning task in which mice first learned to associate different odor cues with specific outcomes and then we unexpectedly reversed these associations. We show that dorsal raphe serotonin neurons, like midbrain dopamine neurons, are specifically recruited following prediction errors that occur after reversal. Yet, unlike dopamine neurons, serotonin neurons are similarly activated by surprising events that are both better and worse than expected. Dopamine and serotonin responses both track learned cue-reward associations, but serotonin neurons adapt slower to the changes that occur at reversal. The different dynamics of these neurons following reversal creates an imbalance that favors dopamine activity when invigoration is needed to obtain rewards and serotonin activity when behavior should be inhibited to prevent perseveration. Our data support a model in which serotonin acts as a surprise signal, reporting mismatches between expectations and reality. This signal is related to uncertainty and is in a privileged position to drive learning and contribute to behavioral flexibility. Thus, our data support a concept of serotonin based on primary functions in prediction, control and learning rather than affect and mood.

III-32. Evidence for temporally sensitive and non-linear processing of spiking signals for behavioral update

Laureline Logiaco[1,2] LAURELINE.LOGIACO@EPFL.CH
Rene Quilodran[3] RENE.QUILODRAN@UV.CL
Emmanuel Procyk[4] EMMANUEL.PROCYK@INSERM.FR
Angelo Arleo[5] ANGELO.ARLEO@UPMC.FR

[1]University Pierre and Marie Curie
[2]Ecole Polytechnique Federale de Lausanne
[3]Universidad de Valparaiso
[4]Universite Lyon 1, Inserm U846
[5]Univ Paris 06/UMR_S 968/Institut de la Vision

Frontal cortical areas control behavioral adaptation to environmental rules. In particular, the dorsal Anterior Cingulate Cortex (dACC) is thought to signal the worth of updating the behavioral strategy to new evidence. Downstream areas would then process and store this signal to ultimately trigger the behavioral change when the opportunity arises (potentially several seconds later). Many studies have described the activation of dACC neurons in different contexts, but the neuronal mechanisms involved in the downstream processing of these signals are poorly understood. We investigated two candidate downstream processing mechanisms: (i) linear integrator networks

with a long timescale (to handle delays); (ii) non-linear, spike-timing sensitive networks capable of short-term memory. We considered the processing of spike trains recorded from dACC while monkeys received feedbacks (reward or no-reward) which instructed the appropriate behavioral strategy update. To gather evidence for either candidate processing mechanism, we emulated decoders with varying timescales. First, decoding different dACC spike trains with a slow timescale (similar to processing by an integrator network) was less efficient than decoding with a faster timescale. Second, the relation between dACC spike trains and monkey's response time during the behavioral update argued against an integrator decoder. If the decoder were an integrator, increased firing of dACC neurons would be expected to trigger a faster response, but we did not find robust evidence for such a relation. In contrast, when considering a non-linear and spike-timing sensitive processing of dACC spike trains, we did find a robust association to the behavioral response. This study is, to our knowledge, the first to use a decoding timescale analysis to identify different processing mechanisms which are relevant for high-level behavioral adjustments. Our results strongly suggest a non-linear and temporally sensitive processing of a multi-dimensional dACC signal, strengthening the hypothesis that dACC specifies the behavioral strategy's identity.

III-33. A novel measure of surprise with applications for learning within changing environments

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

Surprise is informative because it drives attention and modifies learning. Correlates of surprise have been observed at different stages of neural processing, and found to be relevant for learning and memory formation. Although surprise is ubiquitous, there is neither a widely accepted theory that quantitatively links surprise to observed behavior, such as the startle response, nor an agreement on how surprise should influence learning speed or other parameters in iterative statistical learning algorithms. Building on and going beyond earlier surprise measures, we propose a novel information theoretic measure for calculating surprise in a Bayesian framework so as to capture uncertainty of the world as well as imperfections of the subjective model of the world, two important aspects of surprise. The principle of future surprise minimization leads to a learning rule that can be interpreted as a surprise modulated belief update suitable for learning within changing environments. Importantly, we do not need an assumption about how quickly the world changes. We apply our surprise-modulated learning rule to an exploration task in a maze-like environment and to a dynamic decision making task. Our results are consistent with the behavioral finding that surprising events induce humans and animals to learn faster and to adapt more quickly to changing environments.

III-34. Observing seeking observing: How prediction errors boost anticipation

Kiyohito Iigaya[1] KIIGAYA@GATSBY.UCL.AC.UK
Zeb Kurth-Nelson[2] ZEBKURTHNELSON@GMAIL.COM
Giles Story[2] G.STORY@UCL.AC.UK
Ray Dolan[3] R.DOLAN@UCL.AC.UK
Peter Dayan[1] DAYAN@GATSBY.UCL.AC.UK

[1]Gatsby Computational Neuroscience Unit, UCL
[2]University College London
[3]Max Planck Computational Psychiatry Centre

When people anticipate possible future outcomes, they often prefer to know their fate in advance. More formally, gambling-based experiments into what is called ’observing’, or ’information-seeking’, indicate that delayed stochastic rewards become more attractive when the uncertainty about the outcome is resolved in advance by a predictive cue. Recent monkey experiments showed that neurons in the lateral habenula responded differently to the same predictive cues depending on how expected these cues were (Bromberg-Martin and Hikosaka, 2011). No existing model offers a satisfactory account for these data. This includes ideas related to Shannon information, as targets conveying less information can be more attractive (Roper and Zentall, 1999). Further, the preference appears to depend on the delay length between predictive cues and rewards (Spetch et al., 1990), a factor of little informational consequence. The notion that subjects are differentially engaged whilst waiting for the outcome according to the cues (Beierholm and Dayan, 2010), foundered on the recordings in lateral habenula (Bromberg-Martin and Hikosaka, 2011). Here, following Loewenstein (1987), we considered the delay to a future reward as itself being appetitive (utility of anticipation or savouring). In addition, inspired by animals’ excitement following the predictive cues (Spetch et al., 1990), we hypothesize that the attractive force of the anticipation is boosted by the reward prediction error (RPE) occasioned by the predictive cue. Our model accounts for a wide range of existing neuronal and behavioral data, including those in (Bromberg-Martin and Hikosaka, 2011), without appealing to an explicit value for information. We also confirmed the model’s central predictions in our new human empirical studies, in which subjects strikingly preferred delayed consumption of revealed stochastic rewards. Further studies also show that subjects seem to be willing to sacrifice rewards to purchase revelation. We suggest that RPE-boosted anticipation plays a crucial role in driving risk-seeking behaviors.

III-35. Blink different! Blink rate reflects individual differences in directed exploration

Siyu Wang[1] SYWANGR@EMAIL.ARIZONA.EDU
Jonathan Cohen[2] JDC@PRINCETON.EDU
Robert Wilson[1] BOB@EMAIL.ARIZONA.EDU

[1]University of Arizona
[2]Princeton University

Any organism that can learn needs to balance the competing demands of exploring to learn more and exploiting what it already knows. Previous work has suggested that humans solve such ‘explore-exploit dilemmas’ using a mixture of two strategies: an information seeking strategy, known as ‘directed exploration’, and a stochastic strategy, known as ‘random exploration’. In this work we explored the neural correlates of the two strategies by asking whether individual differences in directed and random exploration correlate with individual differences in spontaneous blink rate, an indirect measure of central dopamine. We used a modified version of the ‘Horizon Task’ (Wilson et al. 2014) to measure individual differences in exploration. In this task, directed and random exploration are easily quantified as changes in behavior with ‘time horizon’, the number of choices a participant will make in the future. Directed exploration is measured as the increase in information seeking from short to long horizon, while random exploration is measured as the change in behavioral variability. Across 33 subjects we found a strong correlation ($\rho(32) = -0.54, p = 0.001$) between directed exploration and blink rate, but no correlation ($\rho(32) = -0.21, p = 0.23$) between random exploration and blink rate. These results provide initial evidence for a dissociation between the neural systems responsible for directed and random exploration, with individual differences in dopamine reflecting individual differences in directed but not random exploration. In addition, the negative sign of the correlation with directed exploration is surprising, suggesting that more dopamine is associated with a reduction in directed exploration. More detailed analysis suggests that this negative correlation may reflect two competing effects of increased dopamine: 1) increased risk seeking at short horizons and 2) decreased information seeking (perhaps due to a shorter planning horizon for impulsive individuals) in the long horizon condition.

III-36. Serotonin stimulation modulates waiting through direct and associative learning effects

Madalena Fonseca MADALENA.FONSECA@NEURO.FCHAMPALIMAUD.ORG
Masayoshi Murakami MASAYOSHI.MURAKAMI@NEURO.FCHAMPALIMAUD.ORG
Eran Lottem ERAN.LOTTEM@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Centre for the Unknown

Neuromodulators can affect behavior by directly modulating neural circuits and by indirectly shaping them through learning. This is exemplified by the dual roles played by dopamine (DA): directly energizing behavior and indirectly reinforcing appetitive actions. The case for other neuromodulators is less clear. Photostimulation of dorsal raphe nucleus (DRN) 5-HT neurons enhances patient waiting, biasing the competition between competing patient and impulsive actions. These effects are rapid and transient and have been interpreted as complementing DA's direct function. 5-HT is also known to modulate cortical plasticity and to contribute to reversal learning tasks, but 5-HT stimulation does not appear to drive appetitive or aversive learning. Hence the involvement of 5-HT in associative learning processes remains unclear. Here, we sought to test whether 5-HT stimulation effects include an associative component. We reasoned that if 5-HT acted solely through an immediate and direct effect, then its effects should be independent of context and stimulation history. To test this, we modified the previously studied waiting paradigm by training mice to wait for randomly delayed tones at two spatially distinct ports. In one port, waiting was paired with optogenetic activation of DRN 5-HT neurons in 80% of trials and in the other only 20% of trials. Surprisingly, preliminary results show that waiting developed to be longer in the high probability port and that this difference reversed upon reversal of stimulation probabilities (n=6 mice). Control experiments showed that the effects were not accounted for by change in success, and thus reward, rate. These results suggest that 5-HT modulation of waiting may not be fully explained by immediate and direct effects on action circuits and that, like DA, 5-HT may also contribute to driving contingency-sensitive associative learning processes.

III-37. A cell-type-specific brainstem pathway for basal ganglia control of locomotion.

Thomas Roseberry[1,2] THOMAS.ROSEBERRY@UCSF.EDU
Andrew Moses Lee[1] ANDREWMOSES.LEE@UCSF.EDU
Anatol Kreitzer[2] AKREITZER@GLADSTONE.UCSF.EDU

[1]University of California, San Francisco
[2]Gladstone Institute for Neurological Disease

Locomotion is among the most fundamental actions initiated by the vertebrate brain. The basal ganglia (BG) and its phylogenetically-conserved brainstem targets, including the mesencephalic locomotor region (MLR), are known to play critical roles in the execution of goal-directed locomotion. Locomotion can be initiated or suppressed through stimulation of the BG direct or indirect pathway, respectively, while within the MLR electrical stimulation causes running and lesions result in immobilization. The MLR contains three types of neurochemically distinct neurons: cholinergic, GABAergic and glutamatergic. While it is clear that the MLR receives strong input from the BG, the function of MLR cell types and their control by the BG during locomotion remains a mystery. Here we demonstrate that activity explicitly in glutamatergic neurons of the MLR is required for both spontaneous and BG-driven locomotion. We find that optogenetic stimulation of the glutamatergic population alone is sufficient to drive locomotion, whereas stimulation of the cholinergic neurons only modulates running speed. In contrast, stimulation of GABAergic cells strongly suppresses locomotion. Optogenetic inhibition of MLR glutamatergic neurons demonstrates that they are necessary for spontaneous locomotion, and identified single-unit recordings indicates that subtypes of glutamatergic neurons encode locomotor state and speed. Cell-type-specific monosynaptic rabies tracing reveals that MLR glutamatergic neurons, but not GABAergic neurons, receive a dense input from neurons in BG output nucleus substantia nigra pars reticulata (SNr). Recordings from optogenetically-identified MLR glu-

tamatergic neurons reveals increased or decreased firing with direct or indirect pathway stimulation, respectively. Finally, we demonstrate that MLR glutamatergic neurons are necessary for locomotion initiated by direct pathway stimulation and sufficient to overcome immobilization induced by indirect pathway stimulation. These findings provide a fundamental understanding of how the BG can initiate or suppress a motor program through cell-type specific regulation of its targets.

III-38. Sensory suppression as a natural consequence of Bayesian processing.

Frederic Crevecoeur[1] FREDERIC.CREVECOEUR@UCLOUVAIN.BE
Konrad P Kording[2] KOERDING@GMAIL.COM

[1]Institute of Neuroscience
[2]Northwestern University

When we move, sensory sensitivity is strongly reduced around the time of movements. In the case of saccades, previous work has suggested that saccadic suppression contributes to maintaining spatial consistency during movement. However, that idea cannot explain the facts that saccadic suppression considerably precedes saccade onset, and lasts about twice as long as the saccade duration ($>$150ms). Here we show how these effects emerge naturally by simply considering that signal dependent noise makes delayed sensory signals less reliable around the time of movement. Thus, when we move, sensory signals should have less weight relative to other sources of information, such as internal priors. One easy way for the brain to implement this is to lower the gain of neural activities related to such sensory input. Thus, single modality sensation should be biased downwards during movements, as is experimentally observed. According to our view, it is the sharing of neural resources between single modality estimates and the combined estimates needed for movement that gives rise to the observed suppression. We capture these effects by constructing a model that makes estimates based on the optimal weighting of delayed feedback and its dependence on the noise associated with motor commands. More precisely, we derive an analytical expression of the conditional distribution of the present state given delayed sensory feedback, and show that the variance associated with this estimate scales with the intensity of motor commands. As a result, the weighting of sensory feedback is reduced to maintain statistically efficient control. Simulations from this model generate an important reduction of visual weight that quantitatively matches the behavioral and neural dynamics of saccadic suppression. These results suggest that reducing sensory sensitivity may be a simple mechanism whereby online motor commands are statistically well tuned to the actual state of the body during movement.

III-39. Bidirectional kinetic control and reinforcement in the basal ganglia

Eric Yttri[1,2] NERDOSCIENCE@GMAIL.COM
Joshua Dudman[1,2] DUDMANJ@JANELIA.HHMI.ORG

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

For voluntary actions such as singing or pitching a baseball it is critical that we can learn to modify the underlying movements (e.g. the pitch of a note or velocity of the arm) to improve the outcomes of the action. The basal ganglia are an evolutionarily conserved subcortical circuit, composed of two opponent pathways, direct and indirect, that promote actions that elicit positive outcomes or oppose actions that do not, respectively. This circuitry plays an indispensible role in selecting actions based upon their expected outcomes; however, it is unknown whether basal ganglia circuits are also sufficient to select movement parameters such as velocity through reinforcement. Furthermore, the mechanism by which action selection is shaped (change in gain, recruitment of neurons from specific populations, etc) is unknown. We used cell-type specific photostimulation delivered in closed-loop during

movement to demonstrate that activity in either pathway is sufficient to produce sustained, opposing changes in specific movement kinetics without generalized changes in action selection or motivation. During stimulation sessions, mice increased or decreased the probability of rapid velocity reaching movements in order to trigger more direct pathway stimulation or less indirect pathway stimulation, independent of the direction of the kinetic modulation. These behavioral changes accumulated over tens of trials, persisted after the cessation of stimulation, and were abolished in the presence of dopamine antagonists. Using a combination of computational modeling and multiunit electrophysiological recordings, we discovered a putative mechanism for these seemingly paradoxical effects based upon a BCM-type plasticity rule. Our results indicate that either pathway is sufficient to selectively reinforce specific movement parameters to control the efficiency of reward-seeking actions, demonstrating an unprecedented combination of specificity and flexibility in the control of volition by the basal ganglia.

III-40. An optimal control model of the compensatory eye movement system

Peter Holland[1] P.HOLLAND@ERASMUSMC.NL
Tafadzwa Sibindi[2] T.SIBINDI@ERASMUSMC.NL
Mark Ginzburg[1] M.GINZBURG@BGU.AC.IL
Opher Donchin[1] DONCHIN@BGU.AC.IL
Maarten Frens[2] M.FRENS@ERASMUSMC.NL

[1]Ben-Gurion University of the Negev
[2]Erasmus MC

Compensatory eye movement (CEM) is a general term for several reflexes whose goal is to maintain stable image on the retina during movements of the head. In afoveate animals like mice, the vestibulo-ocular reflex (VOR) uses vestibular input to compensate retinal slip and the optokinetic reflex (OKR) compensates based on the retinal slip itself. Due to the well described anatomical connections, minimal degrees of freedom in the output and a clear time-invariant goal the CEM system is an ideal candidate for quantitative modelling. Optimal control models are the dominant paradigm in current studies of motor control. They combine strong theoretical foundations with elegant explanatory power. However, it has proved difficult to build optimal control models that make specific predictions for real, physiological motor circuits. Here we present results from a working version of the Frens and Donchin state predicting feedback control (SPFC) scheme of the CEM system (Frens & Donchin, 2009). The model architecture is based on the known anatomical connections and physiological properties of individual components. The majority of parameters included in the model are either experimentally derived by ourselves or others. We challenge our model in a broad range of experimental conditions and compare the output to the eye movements of mice (n=13). In addition to reproducing the behaviour of the two primary reflexes (OKR and VOR) the model is also able to accurately mimic the physiological response in situations in which both reflexes function simultaneously. Furthermore, our model replicates a fundamental aspect of the system, adaptation of the VOR, via error driven adaptation of a single parameter. Moreover, the anatomical basis for the model is further evidenced by the ability to mimic the behavioural effect of lesions of specific nuclei in mice by removing the output of the corresponding computation within the model.

III-41. Deconstructing spike timing codes in single premotor neurons using Bayesian feature selection

Damian Hernandez Lahme HERNANDEZ@EMORY.EDU
Sam Sober SAMUEL.J.SOBER@EMORY.EDU
Ilya Nemenman ILYA.NEMENMAN@EMORY.EDU

Emory University

Important questions in computational neuroscience are whether, how much, and how information is encoded in the precise timing of neural action potentials. While these questions have been investigated in detail in sensory neuroscience, studies of cortical motor systems typically rely exclusively on neural firing rates to investigate how the brain controls behavior. We recently demonstrated that, during vocal control in songbirds, spike timing is far more informative about upcoming behavior than is spike rate (Tang et al., 2014). However, current computational methods do not allow us to identify the vocabulary of specific spike timing patterns that control behavior. Here, we present a Bayesian approach to decipher such codes for individual neurons in the songbird premotor area RA, an analog of mammalian primary motor cortex. Specifically, we analyze which multispike patterns of neural activity predict the pitch of the upcoming vocalization, and hence are features, or important codewords, of the neural code. As the size of the possible neural vocabulary is much larger than the number of samples, we consider a strong Bayesian regularization that allows us to approximate the posterior probability that multispike patterns are the features of the code. We use a recently introduced Bayesian Ising Approximation, which properly accounts for the fact that many such possible words possibly overlap and hence are not independent. Our results show that complex, temporally precise multispike combinations are likely used by individual neurons to control acoustic features of the produced song, and that these code words are different in individual neurons and individual animals.

III-42. Optimization costs underlying movement sequence chunking in basal ganglia

Pavan Ramkumar[1] PAVAN.RAMKUMAR@GMAIL.COM
Daniel Acuna[1] DANIEL.ACUNA@NORTHWESTERN.EDU
Max Berniker[2] MAX.BERNIKER@GMAIL.COM
Scott Grafton[3] GRAFTON@PSYCH.UCSB.EDU
Robert Turner[4] RTURNER@PITT.EDU
Konrad Kording[1] KK@NORTHWESTERN.EDU

[1]Rehabilitation Institute of Chicago
[2]University of Illinois, Chicago
[3]University of California, Santa Barbara
[4]University of Pittsburgh

We routinely execute complex movement sequences with such effortless ease that the optimization cost of making them efficient is often under-appreciated. A common facet of many such movements is that they tend to be discrete in nature, i.e. they are often executed as 'chunks'^1–3. Here, we ask (1) what computational principles and (2) what neural mechanisms does the brain use to organize movement sequences into chunks? (1) We framed movement chunking as the result of a trade-off between the desire to make efficient movements and the computational costs of optimizing them. We found that monkeys learning a reaching sequence adopt a cost-effective strategy to deal with this tradeoff. By modeling chunks as minimum-jerk trajectories4, we found evidence that cumulative optimization costs are kept in check over the course of learning. The decision to chunk must thus be based on the expected cost of optimizing the entire sequence, relative to the cost of executing each element of the sequence as a separate chunk. (2) To ask if the brain encodes this cost, we recorded from the Globus Pallidus (GP), a region in the basal ganglia (BG) implicated in habit learning and movement sequencing^5–7. We hypothesized that neurons would encode cost, and that this cost encoding would depend on whether the sequence was novel or habitual. Using generalized linear models (GLMs) of spiking activity, we found that $\sim 25\%$ of 375 GP neurons from two monkeys, encode the expected optimization cost of movements. Further, we found that the percentage of neurons encoding cost was significantly lower for habitual relative to novel sequences, suggesting that cost may only have a weak influence on chunking decisions while executing habitual movements. Our results suggest an important role for BG in the tradeoff between optimization cost and efficiency during movement sequencing.

III-43. Temporal estimation and proactive control in the stop-signal task: a Bayesian model-based fMRI study

Omri Raccah[1] ORACCAH@UCSD.EDU
Jaime Ide, Ph.D.[2] JAIME.IDE@STONYBROOK.EDU
Ning Ma[1] NIMA@ENG.UCSD.EDU
Sien Hu, Ph.D.[2] SIEN.HU@YALE.EDU
Chiang-shan Li, Ph.D.[2] CHIANG-SHAN.LI@YALE.EDU
Angela Yu, Ph.D.[1] AJYU@UCSD.EDU

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

Inhibitory control is an important cognitive function commonly studied using the stop-signal task, in which subjects inhibit a prepotent go response upon detecting a stop signal. Although traditionally it was thought that an individual's stopping ability is simply determined by his stop-signal reaction time (SSRT), more recent data show systematic contextual modulations. Previously, we showed that the observed increase of go response time (RT) and decrease in stop error rate, in proportion to the frequency of stop trials in recent trial history, are rational consequences of larger Bayes-estimated prior probabilities of a stop signal, P(stop) (Shenoy and Yu, 2011; Ide et al, 2013), and that BOLD response in the dorsal anterior cingulate cortex reflects an unsigned prediction error associated with P(stop) (Ide et al, 2013). Recently, we augmented across-trial learning with a Kalman Filter to also predict expected stop-signal delay (SSD), the delay between go and stop stimulus onsets, based on recently experienced SSD. We showed Go RT to be independently modulated by both P(stop) and E[SSD] (Ma and Yu, 2015), indicating stopping behavior to be under the contextual influence of both trial-type and temporal statistics. Here, we examine the neural basis of temporal statistics tracking and prediction, by using E[SSD] and its prediction error, SSD-E[SSD], as parametric regressors in an fMRI GLM analysis (n=84). We find bilateral clusters in the somatomotor cortex, including the paracentral lobules, to be correlated negatively with E[SSD], and positively with its prediction. These areas are distinct from those previously found for P(stop), suggesting separate neural encoding of trial-type and temporal statistics. While further work is needed to identify and differentiate their exact functions, these areas appear critical for the preparatory and learning processes involved in timing-specific aspects of proactive inhibitory control.

III-44. Rapid prediction of movement states from neuronal activity in Parkinson's disease

Minkyu Ahn MINKYU_AHN@BROWN.EDU
Shane Lee LEE.SHANE@BROWN.EDU
Julie Guerin JULIE_GUERIN@BROWN.EDU
David Segar DAVIDSEGAR@GMAIL.COM
Tina Sankhla TINA_SANKHLA@BROWN.EDU
Wael Asaad WAEL_ASAAD@BROWN.EDU

Brown University

Deep Brain Stimulation (DBS) is frequently performed to treat Parkinson Disease (PD), which is the second most common neurodegenerative disease of the central nervous system after Alzheimer's disease [1]. In recent years, adaptive, closed-loop DBS (aDBS) has been proposed to overcome the drawbacks [2] of conventional, open-loop DBS, which include side effects, battery consumption and manual parameter settings. In conventional methods [3], motor deficit is typically assessed by the Unified Parkinson's disease rating scale (UPDRS) and discrete movement tests. These, however, are subjective tests, that prevent the identification of more subtle behavioral motor markers. In addition, they measure un-naturalistic and discontinuous movements that may not generalize to the patients' daily continuous limb movements. Lastly, using these tests makes it difficult to quantify varying motor impairments on a short time scale (i.e. several milliseconds to seconds). One requirement to implement aDBS is to identify specific neural features that can be reliably associated with time-varying continuous behavioral motor performance. In this study, we designed such a motor task, in which subjects track a moving target with a cursor controlled by a joystick. Two performance metrics were identified, including tremor value (TV), which quantifies physiological tremor, and vector difference (VD), which quantifies how well the patient tracks the moving object on an instant-by-instant basis. We observed that these metrics differentiate motor performance among healthy control subjects, patients in clinic, and intra-operative patients. The classification analysis between good and poor motor performance identified by VD and TV, demonstrated that the broad band frequency features (1-50 Hz) of local field potential (LFP) predicts behavioral motor performance with 70% and 76% accuracy, respectively.

III-45. Option coding in the movement system

Joshua Glaser[1] JOSHGLASER88@GMAIL.COM
Daniel Wood[1] DANIELKENTWOOD@GMAIL.COM
Matthew Perich[1] MPERICH@GMAIL.COM
Patrick Lawlor[1] P.N.LAWLOR@GMAIL.COM
Pavan Ramkumar[2] PAVAN.RAMKUMAR@GMAIL.COM
Lee Miller[1] LM@NORTHWESTERN.EDU
Mark Segraves[1] M-SEGRAVES@NORTHWESTERN.EDU
Konrad Kording[2] KK@NORTHWESTERN.EDU

[1]Northwestern University
[2]Rehabilitation Institute of Chicago

In the real world, there are countless options for the next movement. How does the brain represent its options before a decision is made? Previous studies have shown that, for arm movements, the motor cortex can represent a small number of pre-cued movement options. But does the brain represent movement options in a more natural setting when they are not prespecified? And is the brain able to represent a continuum of options? To explore these questions across the motor system, we analyzed 1) eye movements (saccades) while recording from the macaque frontal eye field (FEF), and 2) arm movements while recording from macaque dorsal premotor cortex (PMd). To understand eye movement decisions, monkeys freely searched for a target embedded in a natural scene. To understand arm movement decisions, monkeys continuously moved a manipulandum controlling a cursor to targets on a screen. In both cases, certain upcoming movements were feasible options, while others were not. For instance, if a monkey was looking (or had its cursor) at the bottom of the screen, any upwards movement would be possible, while many downwards movements would not be options. Thus, we could ask how the brain encodes movement options by examining how it encodes screen position. We found that both FEF and PMd represent upcoming movements (eye and arm, respectively) according to the likelihood that they will occur. Essentially, they encode a prior over possible movement options. For example, an FEF neuron with a preferred movement direction to the right would have a higher firing rate when the monkey is looking left than right, regardless of the true upcoming movement. This is because the upcoming saccade is more likely to be in the receptive field (right) when the monkey is looking left. Option representations may be a ubiquitous component of motor planning.

III-46. Local dynamics of trained neural networks

Alexander Rivkind ARIVKIND@TX.TECHNION.AC.IL
Omri Barak OMRI.BARAK@GMAIL.COM

Technion, Israel Institute of Technology

Recently, there has been large progress in utilizing trained Recurrent Neural Networks (RNN) as explanations of neural phenomena. The resulting networks tend to have distributed, context dependent representations—similar

to those observed in the brain using population recording methods. Despite these advances, there are many fundamental gaps in our understanding of these networks: Which tasks can be learned and which cannot? Are the solutions to a specific task obtained through training unique or invariant in any way? While theoretical results exist for naive RNNs, such results are mostly lacking for trained networks. Here, we develop a Mean Field Theory for Reservoir Computing networks trained to have multiple fixed point attractors. Our main result is a low order differential equation governing the network output dynamics. As an immediate application, stability of attractors can be assessed, predicting training success or failure. Furthermore, two properties of trained RNN, with possible implication on biological neuronal circuits are reported. First, RNN output robustness in the presence of variability of the internal neural dynamics is predicted. Second, state dependent frequency selectivity is shown to be an inherent outcome of training a network for multiple target outputs.

III-47. Encoding sensory and motor spatiotemporal “objects” as continuous trajectories in RNNs

Vishwa Goudar VISHWA.GOUDAR@GMAIL.COM
Dean Buonomano DBUONO@UCLA.EDU
University of California, Los Angeles

A highly influential theory in neuroscience is that memories are stored as point attractors within recurrent neural networks. Yet, many sensory and motor tasks that animals perform are inherently time-varying, e.g. speech processing or handwriting. Neurocomputational models commonly cope with this fact by spatializing time. Here we show how firing-rate based recurrent neural networks (RNNs) can encode both time-varying sensory and motor ‘objects’ as continuous neural trajectories. We propose that this framework provides a more natural representation of how the brain tells time: the processing of spatiotemporal objects emerges from the voyage through state-space rather than arrival at a single destination or sequence of point attractors. Some brain areas are known to process both sensory and motor information—e.g., the songbird HVc nucleus. Yet, how a single circuit performs such disparate time-varying computations is poorly understood. We show that a single RNN can discriminate time-varying stimuli and generate corresponding time-varying motor outputs. The RNN is presented with a spoken digit during a sensory epoch, and is required to report which digit was presented with a ‘handwritten’ motor (output) response during the subsequent motor epoch. To perform this, an initially chaotic random-RNN is trained with an ‘innate trajectory’ learning rule. When trained during both epochs, the network transcribes digits with a performance close to 100%. When the RNN is either untrained (reservoir network), or trained only during the motor epoch, performance is significantly worse. Analysis revealed that training "sculpts" the dynamics of the RNN, balancing the input and recurrent drives during the sensory epoch to improve between- and within-digit separation of trajectories. During the motor epoch, when the network operates autonomously, training results in the formation of dynamic attractors which encode written digits as stable and well-separated continuous trajectories.

III-48. On the role of assemblies of hub neurons in cortical dynamics

Hesam Setareh[1] HESAM.SETAREH@EPFL.CH
Moritz Deger[2] MDEGER@UNI-KOELN.DE
Carl Petersen[1] CARL.PETERSEN@EPFL.CH
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH
[1]Ecole Polytechnique Federale de Lausanne
[2]University of Cologne

Sensory information is processed in cortical microcircuits with a network structure that is different from classical Erdos-Renyi random networks. Recent experiments revealed deviations from randomness in wiring and local correlations of synaptic efficacy. To model these non-random network properties we propose to embed assemblies of densely connected hub neurons into a network of excitatory and inhibitory neurons while respecting experimental data on distribution of synaptic weights, connection probabilities, and single-neuron parameters. We show in simulations that in such a network, spiking neurons with spike-frequency adaptation can generate irregular up/down state oscillations, as observed in anesthetized cortex and slow-wave sleep. The presence or absence of hub neuron assemblies may further explain the different responses of infra-granular and supra-granular cortical layers to optogenetic stimulation.

III-49. Dynamics of balanced networks with excess bidirectional connectivity

Shrisha Rao SHRISHA.RAO@PARISDESCARTES.FR
Carl van Vreeswijk CORNELIS.VAN-VREESWIJK@PARISDESCARTES.FR
David Hansel DAVID.HANSEL@UNIV-PARIS5.FR

Universite Paris Descartes

Recent experiments reveal the existence of a fine connectivity structure in the cortical microcircuit. Bidirectional connectivity and motifs of 3 or more highly interconnected neurons are more prevalent than expected for directed Erdos-Renyi random connectivity. What effects does the fine organization of connectivity have on cortical dynamics? To address this question we study the dynamics of networks with excess bidirectionality operating in the balanced regime. We analytically investigate inhibitory networks of binary neurons in the mean field limit. We show that excess bidirectional connections slows down the fluctuations in the neuronal input. As a result, the autocorrelation of the activity decays more slowly than in the corresponding Erdos-Renyi network. These phenomena are due to the small loops that the bidirectionality induces in the network architecture. Together with the relatively strong synapses in balanced networks the small loops lead to a non-negligible effective delayed positive self-coupling. Using numerical simulations, we also investigate the effect of bidirectional connectivity in a balanced network of rodent V1 (Layer 2/3) with conductance-based spiking neurons. We show that this network behaves as expected from the binary network analysis. Bidirectional connections between E cells or between I cells increases the decorrelation time and increases the Fano factor of the spike count. Remarkably, bidirectional connections between I cells are much more efficacious in slowing down the dynamics than those between E cells. In contrast, bidirectional connections between E and I cells decrease both the decorrelation time and the Fano factor. We also show that bidirectional connections do not affect the functional properties of the network. Finally we investigate the dependence of this effect on the synaptic time constants and study how the spike irregularity is modified by 'sensory' stimulation of the network.

III-50. Cortical hierarchy underlies preferential connectivity disturbances in schizophrenia

Genevieve Yang[1] GENEVIEVEJYANG@GMAIL.COM
John D Murray[1] JOHN.MURRAY@YALE.EDU
Grega Repovs GREGA.REPOVS@PSY.FF.UNI-LJ.SI
Michael Cole MWCOLE@GMAIL.COM
Xiao-Jing Wang[2] XJWANG@NYU.EDU
David Glahn[1] DAVID.GLAHN@YALE.EDU
Godfrey Pearlson[1] GODFREY.PEARLSON@YALE.EDU
John Krystal[1] JOHN.KRYSTAL@YALE.EDU
Alan Anticevic[1] ALAN.ANTICEVIC@YALE.EDU

[1]Yale University
[2]New York University

The neocortex contains distributed large-scale networks of coupled excitatory and inhibitory cells. Excitation-inhibition (E-I) balance within these cortical circuits is considered critical for a variety of computations including gain control and sensory tuning. However, it is unclear whether some cortical circuits may utilize specialized computations that are particularly vulnerable to altered E/I ratio. For instance, the hierarchical organization of primate cortex, revealed from tracer studies, is accompanied by a corresponding gradient of intrinsic neural activity timescales across cortical areas. Recent work has linked such timescale differences to differing amounts of local recurrent excitation—a circuit parameter that may also affect network-level vulnerability to increased E/I ratio. NMDA receptor hypofunction, which increases E/I ratio, is thought to be a core pathophysiology in schizophrenia. We therefore hypothesized E/I disruptions in schizophrenia may preferentially impact regions whose local circuit properties are specialized for computations involving greater local recurrent excitation. To test this, we used resting-state fMRI in 161 patients and 164 controls to assess activity patterns across well-characterized functional networks, which we matched to a corresponding large-scale model of resting-state brain activity. To incorporate functional hierarchy, we set recurrent excitation in association areas to be greater than in sensory areas. Simulations of increasing E/I produced global elevations in model-generated fMRI signal covariance and variance, which preferentially affected association regions. The model also predicted that covariance and variance elevations would be positively correlated. We confirmed all of these predictions empirically. These results have two important implications: first, that a global E-I disruption may parsimoniously characterize cortical disruption in schizophrenia, linked to specific fMRI features, and second, that the spatial pattern of schizophrenia findings may arise as a consequence of pre-existing neural dynamics, a result that may inform biomarker refinement across multiple illnesses wherein the spatial pattern of pathology may interact with cortical hierarchy.

III-51. Maintaining balance in heterogeneous-degree networks

Ryan Pyle RYANPYLE1@GMAIL.COM
Robert Rosenbaum RROSENB1@ND.EDU

University of Notre Dame

The balanced network modeling paradigm captures many features of cortical circuits: strong excitatory and inhibitory synaptic currents, dense connectivity and asynchronous, irregular spiking activity. Computational studies of balanced networks often assume Erdss-Renyi style connectivity, where connection probability only depends on cell type (excitatory or inhibitory). This leads to a homogeneous network structure where neurons' in- and out-degrees are all approximately equal. This conflicts with data from biological cortical networks, which reveal a wide variability of in and out-degrees, including 'hub' neurons that have a much higher degree of connectivity than most surrounding cells. We combine heterogeneous mean-field theory with computer simulations to study balanced networks with various heterogeneous degree distributions. Consistent with previous theoretical work, we find heterogeneous in-degrees can break balance. We show that balance can be recovered if in- and out-degrees are correlated, but this requires the surprising result that neurons with a higher in or out-degree have a lower firing rate, consistent with recently reported data from mouse visual cortex. Our results provide a critical extension of the highly influential theory of balanced networks.

III-52. Finite size effects and rare events in balanced cortical networks with plastic synapses

Jeff Dunworth[1] JBD20@PITT.EDU
Bard Ermentrout[1] BARD@PITT.EDU
Michael Graupner[2] MICHAEL.GRAUPNER@PARISDESCARTES.FR
Alex Reyes[3] REYES@CNS.NYU.EDU
Brent Doiron[1] BDOIRON@PITT.EDU

[1]University of Pittsburgh
[2]Universite Paris Descartes
[3]New York University

Cortical neuron spiking activity is broadly classified as temporally irregular and asynchronous. Model networks with a balance between large recurrent excitation and inhibition capture these two key features, and are a popular framework relating circuit structure and network dynamics. Balanced networks stabilize the asynchronous state through reciprocal tracking by the inhibitory and excitatory population activity, leading to a cancellation of total current correlations driving cells within the network. While asynchronous network dynamics are often a good approximation of neural activity, in many cortical datasets there are nevertheless brief epochs wherein the network dynamics are transiently synchronized (Buzsaki and Mizuseki, 2014, Tan et.al., 2014). We analyze paired whole cell voltage-clamp recordings from spontaneously active neurons in mouse auditory cortex slices (Graupner and Reyes, 2013) showing a network where correlated excitation and inhibition effectively cancel, except for intermittent periods when the network shows a macroscopic synchronous event. These data suggest that while the core mechanics of balanced activity are important, we require new theories capturing these brief but powerful periods when balance fails. Traditional balanced networks with linear firing rate dynamics have a single attractor, and fail to exhibit macroscopic synchronous events. Mongillo et.al. (2012) showed that balanced networks with short-term synaptic plasticity can depart from strict linear dynamics through the emergence of multiple attractors. We extend this model by incorporating finite network size, introducing strong nonlinearities in the firing rate dynamics and allowing finite size induced noise to elicit large, yet infrequent, synchronous events. We carry out a principled finite size expansion of an associated Markovian birth-death process and identify core requirements for system size and network plasticity to capture the transient synchronous activity observed in our experimental data set. Our model properly mediates between the asynchrony of balanced activity and the tendency for strong recurrence to promote macroscopic population dynamics.

III-53. Spike triplet-dependent plasticity and spike train correlations in recurrent networks

Gabriel Ocker GABEO@ALLENINSTITUTE.ORG
Michael Buice MICHAELBU@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Plasticity models based on neurons' firing rates, like the BCM rule, have proved highly successful in explaining the development of stimulus selectivity in sensory pathways. How spike-time correlations affect learning in recurrent networks remains an open and active area of investigation. An STDP model based on spike triplets can account for a wide range of STDP experiments, including how STDP protocols depend on pre- and postsynaptic firing rates. In the absence of any spike-time correlations, it can be reduced to the BCM rule. Spike-time correlations inherited from external inputs can allow purely feedforward networks to learn input patterns based on higher-order correlations. These results neglect, however, the recurrence that characterizes cortical circuits, leaving open the question of how recurrent connectivity reflects and impacts the way cortical networks learn computations. We present a consistent and general framework for predicting the dynamics of triplet STDP in recurrent networks. Combining a separation of timescales between plasticity and spike train correlations with a path integral formalism for spike train correlations of arbitrary order provides a prediction for the evolution of individual synaptic weights. We use this theory to examine the relative contributions of second- and third-order internally generated spike-time covariability to the network's plasticity during spontaneous dynamics, and how internally-generated spike train correlations affect the formation of receptive field structure in neurons' recurrent inputs. Finally, we derive the plasticity of multi-synapse patterns of connectivity during learning and spontaneous activity. This formalism is a component of a broader theory that will provide predictions for the functional connectomics work ongoing at the Allen Institute.

III-54. A novel perspective on neural network structure: connections and dissections of homological features

Ann Sizemore[1] ANN.E.SIZEMORE@GMAIL.COM
Chad Giusti[1] CGIUSTI@SEAS.UPENN.EDU
Matthew Cieslak[2] MATTHEW.CIESLAK@PSYCH.UCSB.EDU
Scott Grafton[2] SCOTT.GRAFTON@PSYCH.UCSB.EDU
Danielle Bassett[1] DSB@SEAS.UPENN.EDU

[1]University of Pennsylvania
[2]University of California, Santa Barbara

A necessary initial step toward understanding neural systems dynamics is the development of effective tools for extracting relevant features from the underlying physical networks. Relying on vertex-centered or path-based statistics indicate many such structural networks, including white matter tracts in the human brain, are best described as small-world. Here, we use an alternative method, persistent homology, to measure fundamentally mesoscale features to present a very different view of the structural neural network and appropriate models. We are interested in both cliques (all-to-all connected subgraphs) and patterns of cliques called cycles which enclose a structural cavity. Sewing together these mesoscale features to comprehend global network structure, we compare this architecture with that of Watts-Strogatz, scale-free, modular, and random geometric model networks, and show that surprisingly the random geometric is most similar from this mesoscale perspective. Moreover we inspect individual cycles found and discuss the neural regions involved. Collectively our results suggest persistent homology is an effective network analysis tool, revealing previously unseen structural formations and offering an alternative view of network similarity.

III-55. Structural instability in linear working memory networks

Yashar Ahmadian YASHAR@UOREGON.EDU

University of Oregon

In recent years a number of studies [1,2,3] have underlined the utility of feedforward structures as substrates for working memory. Unlike networks achieving graded short-term memory via positive feedback, feedforward memory networks do not require fine-tuning and are seemingly robust to structural perturbations [2]. Another advertised feature of such networks is that the feedforward structure can be “hidden" within apparently recurrently connected networks. Memory in such structurally stable feedforward networks is however susceptible to corruption by neural dynamic noise. Studying linear networks, Ganguli et al. (2008) showed that to combat noise, the feedforward chain must effectively amplify input signals such that they grow supralinearly over time as they traverse the chain. Such transient amplification allows the network to achieve memory lifetimes scaling with the chain length, while the network remains dynamically stable as the amplified signal eventually dies at the chain's end. I will show that such networks are nevertheless extremely sensitive to perturbations in their connectivity: precisely due to their strong amplification of input signals, small changes in their connectivity can render them dynamically unstable. A trade-off emerges between the memory lifetime, controlled by the strength of amplification, and robustness to connectivity perturbations. The relevant perturbations effectively connect the final nodes in the long feedforward chain to early nodes, thus creating a loop through which the signal cycles and grows indefinitely. The structural sensitivity is more dramatic when the feedforward chain is hidden: for large hidden chains the mere deletion of a single randomly chosen connection can provide a relevant perturbation strong enough to render the network unstable. Beyond working memory applications, nonnormal linear systems and their feedforward structure have attracted attention lately in neuroscience; the findings here suggest that a subclass of these systems featuring long feedforward chains with strong amplification can be prone to structural instability.

III-56. Somatostatin interneurons drive cortical gamma rhythms that link global stimulus features

Julia Veit JVEIT@BERKELEY.EDU
Richard Hakim RHAKIM@BERKELEY.EDU
Hillel Adesnik HADESNIK@BERKELEY.EDU

University of California, Berkeley

Rhythmic activity is a ubiquitous feature of brain circuits, and has been linked to many aspects of sensory and cognitive function. Gamma band rhythms in particular may aid communication between distributed cell assemblies by providing an efficient conduit for information transfer between synchronously firing neurons. Most models pose that parvalbumin (PV) positive GABAergic interneurons are critical for driving gamma rhythms through somatic inhibition. Contrary to this prevailing notion, we show that dendrite-targeting somatostatin (SOM) interneurons mediate a narrow-band, visually-driven low gamma oscillation in the primary visual cortex (V1) of awake, freely locomoting animals. SOM driven oscillations are stimulus specific and depend on matched stimulus features across large regions of the visual field. Long-range horizontal excitatory circuits in L2/3 recruit SOM-mediated rhythmic inhibition to drive gamma rhythms and phase lock spatially separated ensembles. These data establish dendrite-targeting SOM neurons as key mediators of intra-areal neuronal synchronization. By operating through dendritic inhibition, SOM-driven neural oscillations provide a new mechanism for flexibly gating synaptic integration and plasticity rhythmically in time, potentially with sub-cellular resolution.

III-57. Slow waves propagation in the cortex: how wavefronts are shaped by the layer structure.

Cristiano Capone[1,2] CRISTIANO0CAPONE@GMAIL.COM
Beatriz Rebollo[3] BEATRIZ.REBOLLO.G@GMAIL.COM
Alberto Munoz-Cespedes[4] AMUNOZC@BIO.UCM.ES
Paolo Del Giudice[1] PAOLO.DELGIUDICE@ISS.INFN.IT
Maria V Sanchez-Vives[3] SANCHEZ.VIVES@GMAIL.COM
Maurizio Mattia[1] MAURIZIO.MATTIA@GMAIL.COM

[1]Italian Institute of Health
[2]University La sapienza
[3]IDIBAPS, Barcelona, Spain
[4]Universidad Complutense de Madrid

It is widely recognized that neuronal spontaneous activity can provide valuable information on the structure of the underlying neuronal network. With this aim, we focused on the spontaneous slow-wave activity generated and propagated in cortical slices, aiming at relating their spatio-temporal organization with the laminar structure. We studied the wavefront propagation analyzing, across the slice, the Multi-Unit Activity slow oscillation (SO) between Up and Down states. We found different propagation modes in terms of speed and direction. Despite such variability, we found a rather stereotyped early propagation strip (EPS), i.e. a strip where the more advanced part of the wavefronts is found. This strip is the same for each mode of propagation and it is almost parallel to the cortical surface. We evaluated the hypothesis that this is related to local excitability using a large scale simulation of a slice model, with cortical modules of spiking neurons arranged in a 2D lattice. Each module worked as a relaxation oscillator. We modeled the excitable strip with different shapes by gradually increasing self- and cross-modular connectivity of involved modules. In this way we have been able to reproduce the overlap between the most excitable strip and the EPS observed in experimental evidences. We found that an overlap between excitable strips and detected EPSs emerges only for an optimal set self- and cross-connectivity parameters, suggesting that suggesting that connectivity parameters of a cortical slice can be inferred relying on a maximum overlap criterion. Besides, SO features like maximum firing rate and longest Up state durations are expected to be found within the same excitable strip: a model prediction confirmed by experiments, further strengthening the relationship between

spontaneous activity and network structure. Matching anatomical layer distribution and EPS, we finally found a remarkable overlap between EPS and Layers 4 and 5.

III-58. Theta oscillations mediate the neural mechanism to generate future predictions

Karthik Shankar KSHANKAR79@GMAIL.COM
Inder Singh INDERDEEP15@GMAIL.COM
Marc Howard MARC777@BU.EDU

Boston University

Theta oscillations observed in rodent hippocampus have a systematic effect on the activity of place cells. The spiking activity of a place cell is known to synchronize with progressively earlier phases of the theta oscillation as the animal traverses across its place field, a phenomenon called phase-precession. It has been suggested that this phenomenon corresponds to the animal cognitively predicting or imagining its future trajectory. Consistent with this, we describe a mechanism that shifts the current representation of the spatio-temporal memory state into a future state systematically within each cycle of theta oscillation by periodically modulating synaptic conductances. These shifted memory states can then be availed to construct a timeline of future predictions. Critical to this mechanism is our hypothesis that the spatio-temporal memory representation in hippocampus is constructed by a network that encodes the Laplace transform of real time inputs. This mathematically facilitates the network to represent the memory in a scale-invariant way, and non-destructively time-translate it to a future state. This model ties together several crucial findings from hippocampal neurophysiology, namely phase-precession of place cells, the traveling wave of theta oscillations along the dorsoventral axis, and the increase in spatial scales of the place cells along the dorsoventral axis. It leads to novel testable predictions that phase-precession should also be observed in time cells in the hippocampus; it should also be observed in cells outside the hippocampus that code for future anticipation of a reward in a characteristic way that depends on the spatio-temporal scale of the task performed the animal. This could directly shed light on neural underpinnings of the cognitive act of imagining or predicting the future

III-59. Sparse components of sensorimotor ECoG signals are relevant for speech control

Kristofer Bouchard[1,2] KRISTOFER.BOUCHARD@GMAIL.COM
Alejandro Bujan[3] AFBUJAN@BERKELEY.EDU
Edward Chang[2] CHANGED@NEUROSURG.UCSF.EDU
Friedrich Sommer[3] FSOMMER@BERKELEY.EDU

[1]Lawrence Berkeley National Laboratory
[2]University of California, San Francisco
[3]University of California, Berkeley

The concept of sparsity has proven very useful to understanding elementary neural computations in sensory systems. Metabolic measurements indicate that neural activity must be sparse in the entire brain; however, the functional role of sparseness outside of sensory brain areas, such as motor regions, is not well understood. To address this basic question we investigated the functional properties of sparse structure in neural activity from human speech sensorimotor cortex. We focused on speech because, although linguistics has provided excellent descriptions of the motor output, little is known about the underlying neural processes that generate speech. High-density electrocorticography (ECoG) from speech sensorimotor cortex (vSMC) in neurosurgical patients is a powerful tool to record broad-coverage, high-temporal resolution neural activity at meso-scale spatial resolution. However, ECoG recordings, like all field potentials, contain a mixture of different brain signals, which makes both the interpretation and extraction of specific features challenging. Utilizing efficient coding methods (independent components analysis: ICA, and convolutional sparse coding: CSC) we decomposed the complex spatial-temporal patterns of high-gamma activity during speech production into separate, temporally sparse signal components. Notably, we show that these components are reliably activated across trials of the same utterance (Fig.1a). Additionally, we found different components corresponding to the major oral articulators (e.g. Coronal Tongue, Dorsal Tongue, Lips), which were selectively activated during all utterances that engaged that articulator (Fig. 1b). Some of the components corresponded to spatially sparse activations (e.g. Fig. 1c), while others were more spatially distributed. Features with similar properties were also extracted using CSC, and required less data pre-processing. Finally, decoding of individual utterances from vSMC ECoG recordings improves consistently when linear classifiers are trained using the sparse codes generated by CSC. Together, these results suggest that sparse coding may be an important framework and tool for understanding sensory-motor activity generating complex behaviors.

III-60. Decomposition of the neural co-variations in a population of rat hippocampal place cells

Tao Tu TT2531@COLUMBIA.EDU
Lars Buesing LARS@STAT.COLUMBIA.EDU
Paul Sajda PS629@COLUMBIA.EDU

Columbia University

Multi-unit recording techniques have made it possible to study a population of neurons firing in concert, where the collective firing activity often exhibits shared variability across multiple neurons. Here we analyzed the co-variations in the firing rates of a population of place cells in the CA1 region of rat's hippocampus when the rat ran back and forth on a linear track. Previous studies have shown that the spatial location of the rat on the linear track, the running velocity, and the phase of the theta oscillation in the local field potential are known external factors that modulate the firing of place cells. Here we employed a Poisson Linear Dynamical System (PLDS) model to capture the unexplained variability in the total shared variability that presents in the simultaneous firing of multiple place cells, after accounting for the modulation of these known external factors. Using this model, first we show that common external factor that modulates the firing activity of multiple place cells can be accurately captured by the latent factors identified from the data, when we explicitly not account for it in the model. Then we show that a spatial location dependent between-trial variation is captured by one latent factor that represents a slow fluctuation, after taking into account all external factors. Furthermore, the factor loadings corresponding to this slow latent factor precisely reflects the magnitude and polarity of the influence of the slow trend on each individual neuron. Lastly, to confirm our model accurately captures most of the shared variability across neurons, we combined the identified slow trend latent factor and all external factors to generate synthetic spike trains for these neurons via the PLDS model. The structured spatial location dependent between-trial variations estimated from this low-rank approximation of data is in great agreement with those obtained from experimental data.

III-61. Cortical communication via randomized dimensionality reduction with local synaptic connections

Christopher Rozell[1] CROZELL@GATECH.EDU
Ninghao Liu[1,2] NHLIU@GATECH.EDU

[1]Georgia Institute of Technology
[2]Texas A&M University

Cortical regions transmit information through axonal projections that often form a bottleneck, having many fewer fibers than the number of cells encoding information in the region. An example is the optic nerve where infor-

mation from $\sim 100M$ photoreceptors is transmitted to the thalamus using $\sim 1M$ axons. While physiological constraints may require a compressive communication scheme, the data compaction must retain critical sensory information and the specific wiring patterns are likely unknown to the receiving population. Previous work proposed the randomized dimensionality reduction method known as compressed sensing (CS) as a model for these communication bottlenecks. In this model, every projection encodes a random combination of activity from the input population and CS theory guarantees information is preserved for signals that are sparse in a basis. This model was previously tested in the context of the sparse coding hypothesis by using completely unsupervised learning to train a dictionary for compressed images. While the learned dictionary appears random, the mapped receptive fields (RFs) of these putative cells resemble the RFs observed in V1. While encouraging, this basic CS model unrealistically requires every projection fiber to aggregate activity from every input cell. In recent work we have shown that randomized block-diagonal matrices can still satisfy the CS requirements and preserve signal information. Importantly, these operators require local wiring so that each projection fiber only aggregates activity from a local subset of input neurons. In this work we show that these biophysically plausible CS operators can be used to model the compressive information bottleneck from retina to V1. Specifically, we demonstrate that there is little degradation of the learned representation in a sparse coding framework relative to a CS model with global wiring. This work establishes biophysically plausible randomized dimensionality reduction as a potential model for cortical communication pathways.

III-62. A neural model for self-localization in an ambiguous world

Ingmar Kanitscheider IKANITSCHEIDER@MAIL.CLM.UTEXAS.EDU
Ila R Fiete ILAFIETE@MAIL.CLM.UTEXAS.EDU

The University of Texas at Austin

To self-localize while moving about, you must deal with erroneous motion cues and partial spatial information from landmarks. Often, landmarks at distinct locations look similar (two glades in a forest or doors along a hallway). Encountering one such landmark after starting from an uncertain location is a good indicator that you are at one of the corresponding locations, but not which one. In such scenarios, accurate localization is known to require sequential updating of multiple hypotheses, implemented by engineers using asymptotically Bayesian methods like the particle filter. Animals excel at navigating with high uncertainty, but we know little about how brains perform the hard computations that are involved. We set a problem of inferring location in environments with several identical landmarks at known locations and a noisy velocity input, starting from unknown start locations. We take a model-free approach to generate neurally plausible solutions by training recurrent neural networks, then scrutinize their performance, errors, and dynamics. The networks learn to update their estimates through velocity integration, construct relational information about landmark positions, and use memory of the last landmark encounter to choose between competing location hypotheses. They generalize to new environments with different specified landmark configurations, while matching particle filter performance. Output units encode multiple location hypotheses after the first landmark and before the second, when location ambiguity is resolved, a prediction for neural representations under similar experimental paradigms. In any trial, the recurrently coupled hidden units encode the distance run from the first landmark; absolute location tuning only emerges after two or more landmark encounters. Most units exhibit conjunctive speed tuning. These results demonstrate that sequential probabilistic inference can be solved by deterministic dynamics, and provide predictions for neural representations during real-world navigational challenges. We aim to build on this approach by adding spatial tasks and constraints on architecture.

III-63. Understanding principles of encoding navigationally-relevant variables in entorhinal cortex

Kiah Hardcastle KIAHHARDCASTLE@GMAIL.COM
Niru Maheswaranathan NIRUM@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Lisa Giocomo GIOCOMO@STANFORD.EDU

Stanford University

To survive, animals must maintain an internal representation of position and movement at all times. Medial entorhinal cortex (MEC) likely supports this representation, as MEC neurons modulate their activity with the animal's position, head direction, and running speed. In particular, the superficial layers of MEC have attracted attention for containing strikingly regular cell types encoding a single spatial variable, namely position-encoding grid cells. Despite this attention, many ($>$50%) of cells in this region remain uncharacterized by conventional tuning curve-based methods, indicating that current heuristics for identifying cell types may only reveal the tip of the iceberg of MEC coding properties. Specifically, uncharacterized cells may irregularly encode single spatial variables, as well as encode conjunctions of multiple variables, contrasting with the classical view that superficial MEC neurons encode only single variables. To investigate this possibility, we employ a statistical approach to identify single-cell encoding properties. We fit a nested series of generalized linear models (GLMs) containing various combinations of position, head direction, and speed information, and used principled hierarchical probabilistic model selection methods to detect which variables each cell encodes. We confirm that classical heuristics miss many important features of entorhinal coding. First, we detect more navigationally-relevant neurons: of 799 neurons recorded from mice during open field navigation, we find that 71% encode at least one variable, while classical metrics based on tuning curves detect only 45%. Second, we observe increased multiple-variable encoding: we find that 37% of cells encode multiple variables, while classical methods report 6%. Lastly, we find that the fraction of multiple-variable cells increases with running speed, consistent with previous information theoretic analyses suggesting that conjunctive cells are advantageous when sensory inputs vary rapidly (Finkelstein et al., 2015). Overall, our principled methods successfully confront MEC heterogeneity and uncover remarkable, adaptive behavioral-state dependent changes in its spatial coding properties.

III-64. Efficient coding with time-varying stimuli & noise

Kamesh Krishnamurthy[1] KAMESHKK@GMAIL.COM
Barry Wark[2] WARK@U.WASHINGTON.EDU
Adrienne Fairhall[2] FAIRHALL@U.WASHINGTON.EDU
Jonathan W Pillow[3] PILLOW@PRINCETON.EDU

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

The efficient coding hypothesis (ECH) prescribes that sensory systems should use their dynamic range efficiently by matching their input-output relation to the statistics of the stimuli. Classical ECH successfully predicted that with small output noise, the optimal steady-state input-output function has the same form as the cumulative distribution of the stimulus. However, the ECH in this form does not account for two salient features of neural responses: i) stimuli vary in time, and neural systems respond to these changes by adjusting their input-output functions; ii) neural responses are noisy and this noise can significantly influence the optimal input-output function. Here we extend the ECH to include temporally modulated stimuli and multiple output noise models. In particular, we calculate how neurons should adjust their input-output function to tractable, but realistic, modulated stimuli commonly used in experiments. For the case of stimuli that switch between two distributions, the optimal function is a time-dependent linear combination of two optimal steady-state solutions, and the dynamics of this linear combination are consistent with experiments. We also compare the predicted dynamics of the slow firing rate adaptation to that of responses observed in mouse retinal ganglion cells. Our results suggest that the adaptation of neural responses is consistent with a system which maintains efficient encoding in dynamic environments. For the case of noisy outputs, we can determine the optimal input-output function under Poisson noise for arbitrary values of the maximal firing rate via numerical optimisation. We show that, contrary to previous suggestions, these functions transition smoothly from being step-like to parabolic as the maximal firing rate increases. Our method can be easily extended to include other noise models and constraints on the average firing rate. Ongoing work systematically characterises these input-output functions and compares them to experimental data.

III-65. Dynamical constraints on improving coding fidelity through the 'sign rule'

Birgit Kriener KRIENER@UTEXAS.EDU
Ila R. Fiete ILAFIETE@MAIL.CLM.UTEXAS.EDU

The University of Texas at Austin

It has been suggested that noise correlations, appropriately tuned, can substantially increase coding fidelity in neural representations compared to if noise is independent across neurons. For stimulus discrimination, correlations that decrease noise fluctuations along the stimulus-response direction improve information transfer compared to the uncorrelated case. This observation is called the 'sign rule' in two dimensions, because noise correlations are of opposite sign from the stimulus-evoked (signal) correlations. Finally, large gains in coding capacity can result from realistically small microscopic correlations, if of the right variety. We ask whether it is plausible to independently tune signal and noise correlations according to the sign rule, considering that in reality both are shaped by the recurrent network dynamics. We investigate simple models of neural networks in quasi-2D configurations, in which we derive the signal and noise correlations as a function of network parameters, including weights, for both stationary stimuli (discrimination task scenario) and time-dependent stimuli (generalized regression maximizing correlation between stimulus and response). In linear networks, optimizing weights to flatten noise along the stimulus direction results in vanishing signal-to-noise ratio (SNR), while maximizing the time-varying stimulus-content in the response yields finite fluctuation size along the stimulus direction that is independent of network parameters and can thus not be further minimized. We extend our analysis to balanced random networks of binary and leaky integrate-and-fire neurons. These networks can assume asynchronous-irregular activity states akin to what is observed in cortex, and pairwise correlations are inherently decreased by the ongoing recurrent inhibitory feedback. Though this helps to reduce noise correlations, it also limits the achievable amount of stimulus content in the network response. We conclude that the form of correlation shaping required for the sign rule also modifies the stimulus response to the detriment of the SNR, at least in our network architectures.

III-66. Decoding position from single field and multi field cells in the dentate gyrus

Fabio Stefanini[1] FS2545@COLUMBIA.EDU
Mazen Kheirbek[1] MK3156@CUMC.COLUMBIA.EDU
Lyudmila Kushnir[1] LK2511@CUMC.COLUMBIA.EDU
Joshua Jennings[2] JOSHJENNINGS@STANFORD.EDU
Charu Ramakrishnan[2] CHARUR@STANFORD.EDU
Karl Deisseroth[2] DEISSERO@STANFORD.EDU
Garret Stuber[3] GARRET_STUBER@MED.UNC.EDU
Rene Hen[1] RH95@CUMC.COLUMBIA.EDU
Stefano Fusi[1] SF2237@CUMC.COLUMBIA.EDU

[1]Columbia University
[2]Stanford University
[3]University of North Carolina

Dentate gyrus granule cells (DG GCs) of the hippocampus process information arising from the Entorhinal Cortex (EC) and transmit this to CA3. Previous electrophysiological studies (Leutgeb et al., 2007) have indicated DG GCs mai encode spatial information through the tuning of their firing fields, where $\sim 30\%$ of DG GCs have been reported to have a single firing field (single field cells, SFCs) while $\sim 60\%$ multiple firing fields (multi-field cells, MFCs). However, it remains unclear whether these tuning properties are stable enough to be used to decode the animal's position. To understand how position is encoded by DG GCs, we use miniaturized head-mounted microscopes to perform functional Calcium imaging of DG GCs as mice foraged in an open field. We show for the first time that position can be accurately decoded using population activity in the DG with a precision that is comparable to the animal's body size. This could be achieved both with a linear and a nonlinear decoder using 100 to 600 simultaneously recorded cells. We then identified SFCs, which, similar to that seen in electrophysiological studies, constituted $\sim 1/3rd$ of the recorded neurons. We decoded position from the SFCs and separately from an equal number of MFCs. The cells used for the comparison are the best in each category and were selected using the Lasso algorithm. The decoding accuracy for the MFCs was comparable or better than that for SFCs. This indicates that multi-field cells can be important for encoding position despite their promiscuous spatial tuning. This multi-field encoding strategy may be similar to that of grid cells, which in theoretical studies has been shown to be more efficient than single field cells in encoding space (Mathis et al. 2012).

III-67. An improved approach for assessing the role of correlated noise using copula models

Shaobo Guan SHAOBO_GUAN@BROWN.EDU
Ruobing Xia RUOBING_XIA@BROWN.EDU
David Sheinberg DAVID_SHEINBERG@BROWN.EDU

Brown University

The structure of noise can have a significant impact on the amount of information coded by a neural population. Recent advances in simultaneous recording techniques have inspired a large number of studies looking at correlated noise. However, to assess its impact on a large neural population is very challenging. The popular approaches either involve parametric models like the multivariate gaussian (MVG) distribution, or discriminative classifiers like SVM. In this study, we present an improved approach, which uses a copula distribution to explicitly model the population response, and evaluate the impact of correlated noise by information metrics and decoding performance through a Bayesian decoder. Popularized in finance and insurance modeling, copulas are multivariate distributions that flexibly link arbitrary univariate marginal distributions via certain dependency structures. To model the population spike counts, we chose the gaussian copula on negative binomial marginals. We developed methods that effectively fit the model parameters, estimate information, and perform decoding for a population of tens to hundreds of neurons. Tested on neural data collected in macaque area V4 in response to natural images, our approach works better than the MVG and SVM approaches. It effectively incorporates the advantages and addresses shortcomings of them. Like the MVG, the parameters of our model explicitly correspond to important summary statistics including the mean, variance and pairwise noise correlation; but better than like MVG, our model captures the discrete, non-negative, and positively skewed nature of spike counts, yielding a more accurate descriptor. When evaluating decoding performance, our model is comparable to the state-of-the art SVM decoder, while being less prone to over-fitting, and it allows inspection of the contribution of every summary statistic whereas SVM does not. Furthermore, various regularized correlation estimation methods can be easily implemented in the model. The Matlab code for the model will be released.

III-68. Hierarchical differences in population coding in auditory cortex

Josh Downer JDOWNER@UCDAVIS.EDU
Mitchell Sutter MLSUTTER@UCDAVIS.EDU

University of California, Davis

Various models for population coding in auditory cortex (AC) have been advanced and most have focused on A1. Thus, our understanding of how the neural code for sounds progresses along the cortical hierarchy remains obscure. We recorded single neurons from A1 and a secondary field, middle-lateral belt (ML) from rhesus macaque AC to address this issue. We presented amplitude-modulated (AM) noise during both passive and engaged conditions. In both fields, neurons exhibit monotonic tuning to AM depth, with A1 neurons mostly exhibiting positive rate-depth functions and ML neurons evenly distributed between positive and negative functions. We measured noise correlation between simultaneously recorded neurons and found that, whereas engagement decreased average noise correlation in A1, engagement increased average noise correlation in ML. This finding surprised us considering that attentive states most commonly decrease average noise correlation in sensory cortex. We therefore constructed models of A1 and ML populations to simulate the effect of noise correlation on coding accuracy in each field. We modeled each population as an array of filters (fit to the recorded neurons' rate-depth functions) and simulated different magnitudes of average noise correlation across all pairs (from ~ 0 to ~ 0.2). Using a decoder that ignores noise correlation, we found that A1 coding accuracy is decreased as noise correlation increases whereas noise correlation has little effect on ML accuracy. We argue that this difference can be explained by the difference in the distribution of single neuron AM tuning between fields: since most A1 neurons exhibit positive rate-depth slopes, noise correlation can't average out. On the other hand, in ML the correlated activity across the population can act as common noise rejection. These results demonstrate striking changes in population coding along the sensory hierarchy, as well as provide data to support diverse effects of noise correlation depending on single neuron tuning.

III-69. Understanding the MT representation of speed and speed changes in a structurally simple neural model

Oana Constantin[1] OANNA.CTIN@GMAIL.COM
Lisa Bohnenkamp[1,2] LISA.BOHNENKAMP@UNI-BREMEN.DE
Detlef Wegener[2] WEGENER@BRAIN.UNI-BREMEN.DE
Udo Ernst[1] UDO@NEURO.UNI-BREMEN.DE

[1]Institute for Theoretical Physics
[2]Institute for Theoretical Neurobiology

In the primate visual system, area MT is closely linked to the processing and perceptual representation of motion in visual scenes. Neurons in MT are strongly tuned to the direction and speed of moving stimuli, and they exhibit pronounced transients in their firing rates after changes in visual stimulation. These transients increase the sensitivity of neurons and they are closely correlated to behavioral per-formance. Understanding the neural dynamics shaping these responses, and their effects on infor-mation transmission of arbitrary time-varying signals, is key to understanding how the visual system copes with dynamic scenes. In our contribution, we extend a well-established framework based on divi-sive, global inhibition acting on an excitatory neuron population of direction-tuned cells, and develop a dynamical model for MT neurons that reproduces detailed characteristics of experimentally observed transients. In its simplest form, our model consists of only two coupled differential equations with two time constants for inhibition and excitation. Using an optimization procedure, we fit the model to dif-ferent single cell MT recordings and find a perfect fit of the model's activity to both, transient and sus-tained activity patterns. The model reliably reproduces MT cell responses to arbitrary accelerations and decelerations of a moving stimulus, starting from both low and high base speeds (recent unex-plained experimental data). If the inhibitory time constant is a multiple of the excitatory time constant, the model turns out to be analytically tractable for a piecewise constant input current: The analytical solution allows to quantify the transients' magnitude as a function of general neuron parameters such as response gain and time constants, and to predict the cell's response to arbitrary stimulus changes.

III-70. Redundancy reduction, efficiency and prediction: towards a unified theory

Matthew Chalk MATTHEWJCHALK@GMAIL.COM
Olivier Marre OLIVIER.MARRE@GMAIL.COM
Gasper Tkacik GTKACIK@IST.AC.AT
IST Austria

A long-standing idea—the 'efficient coding' principle—posits that neurons are adapted to maximise encoded information. In the low-noise limit, this predicts that neurons should remove statistical redundancies in their inputs, generating responses that are decorrelated and sparse. A limitation of efficient coding, however, is that it only considers 'how much' but not 'what' information is encoded. In contrast, it has been argued that sensory systems should prioritise stimuli that are informative about the future, and thus can help guide behaviour. How does the objective of encoding predictive information relate to established notions of redundancy reduction and efficiency? We propose that, to be optimally predictive, neurons must trade-off competing demands. On the one hand, they should preferentially encode information about slow/temporally correlated features. On the other hand, to encode this information efficiently requires temporally decorrelating their responses. A natural way to capture predictive coding is the information bottleneck (IB) framework, maximising the information encoded about future stimuli, constrained on information encoded about the past. While in general, IB problems are notoriously difficult to solve, we develop an approximative framework that can be solved tractably. In contrast to previous work, this allows us to consider the information encoded by the entire history of neural responses, and generalises easily to populations of non-linear neurons and arbitrary stimulus/response statistics. We find that the optimal tradeoff between prediction and redundancy reduction depends crucially on the stimulus statistics. For certain stimuli, neurons should decorrelate their responses, while for others, all resources should be put into encoding predictive features, at the cost of highly redundant responses. Finally, we show that the framework provides a direct bridge between predictive coding, and established models of efficient coding (e.g. sparse coding and ICA), which emerge as special cases of the theory.

III-71. Emergence of neuronal signals supporting naturalistic texture discrimination

Corey M Ziemba[1] CZIEMBA@NYU.EDU
Robbe LT Goris[1,2] ROBBE.GORIS@NYU.EDU
Gabriel Stine[1] GABRIEL.STINE@NYU.EDU
Eero Simoncelli[1,2] EERO.SIMONCELLI@NYU.EDU
J Anthony Movshon[1] MOVSHON@NYU.EDU
[1]New York University
[2]Howard Hughes Medical Institute

In anesthetized macaque monkeys, neuronal sensitivity to complex features of visual textures is substantially stronger in V2 than in V1, suggesting that V2 neuronal activity might directly support texture perception. To test this idea, we investigated the relationship between single neuron activity in V1 and V2 and simultaneously measured psychophysical judgments of texture. We generated stimuli along a continuum between naturalistic texture and spectrally-matched, phase-randomized noise, and trained macaque monkeys to classify them according to which end of this 'naturalness' continuum they resembled. Monkeys learned to discriminate briefly presented (500 ms) near-peripheral stimuli well, with performance rivaling that of human observers. The firing rates of single V1

and V2 neurons carried much less information about texture naturalness than the behavioral reports, although V2 neurons were significantly more sensitive than V1 neurons on average. Selectivity for naturalness evolved with a time course that differed across areas. In V2, selectivity emerged early and peaked 50 ms after the initial transient. In V1, selectivity was initially absent, but increased gradually throughout the stimulus period. Both V1 and V2 showed significant decision-related activity: choice probability was modest but significant, and of similar magnitude in V1 and V2. Mirroring the difference in the dynamics of selectivity, decision-related activity in the two areas evolved with different time courses: choice probability emerged early in V2 and remained stable, but emerged later and more gradually in V1. Together, these results suggest that V2 establishes sensitivity to naturalistic visual structure, and feeds this information back to V1. The comparatively weak sensitivity of single neurons suggests that texture perception likely arises from the combined activity of many V2 neurons, and may be consolidated further downstream.

III-72. Direct measurement of correlation responses in Drosophila direction-selective neurons

Emilio Salazar Cardozo[1] EMILIO.SALAZARCARDOZO@YALE.EDU
Juyue Chen[2] JUYUE.CHEN@YALE.EDU
Matthew S Creamer[1] MATTHEW.CREAMER@YALE.EDU
Omer Mano[1] OMER.MANO@YALE.EDU
Catherine Matulis[1] CATHERINE.MATULIS@YALE.EDU
Joseph Pottackal[1] JOSEPH.POTTACKAL@YALE.EDU
James E Fitzgerald[3] JAMES.ELIOT.FITZGERALD@GMAIL.COM
Damon Clark[1] DAMON.CLARK@YALE.EDU

[1]Yale University
[2]Yale
[3]Harvard University

Animals estimate motion in their visual field to accomplish tasks ranging from course correction during navigation to avoiding predators and hunting prey. In order to estimate visual motion, animals must integrate visual information over time and space, and use nonlinear processing, which makes visual motion estimation a paradigm of neural computation. In insects, including the fruit fly Drosophila, the neural and behavioral responses to visual motion have historically been well predicted by a model known as the Hassenstein-Reichardt correlator (HRC). To detect motion, this model correlates two inputs by delaying one before multiplying them together. However, the operations that implement the neuronal motion detector have yet to be quantified. In this study, we used calcium imaging to measure the response properties of single local motion detectors in Drosophila's visual circuits. We found that the spatial receptive fields of the cells were broader than expected, and their linear response properties hinted at both their direction-selectivity and at their selectivity for light or dark edges. By presenting flies with stimuli containing specific, imposed pairwise correlations, we found that these cells perform correlations on timescales faster than expected from their linear receptive fields, on the timescale of 20 ms. Strikingly, these cell types respond to negative correlations in a manner not predicted by the HRC. Lastly, we presented cells with pairs of adjacent bars, which turned white or black in sequence. Responses to these bar pairs showed that the elementary motion detectors in the fly are sensitive to inputs at three distinct points in space, rather than the two envisioned by the HRC, and that their responses to pairwise contrast changes are highly structured and complementary. Our results show that fly direction-selective cells are responsive to a rich and complementary set of fast timescale correlations.

III-73. Synaptic rectification controls nonlinear spatial integration of natural visual inputs

Max Turner[1] MHTURNER@UW.EDU
Fred Rieke[1,2] RIEKE@UW.EDU

[1]University of Washington
[2]Howard Hughes Medical Institute

A central goal in the study of any sensory system is to predict neural responses to complex inputs, especially those encountered during natural stimulation. Nowhere is the transformation from stimulus to response better understood than the vertebrate retina. Nevertheless, descriptions of retinal ganglion cell (RGC) computation are largely based on stimulation using artificial visual stimuli, and it is not clear how these descriptions map onto the encoding of natural stimuli. One such description concerns nonlinear integration over visual space, a phenomenon that has been demonstrated in many classes of RGC using artificial stimuli like gratings. Despite this, many RGC receptive field (RF) models, including classic difference of Gaussians models and more recent linear-nonlinear or generalized linear models, assume linear integration across visual space. It is unknown whether nonlinear responses are primarily elicited by artificial stimuli designed to test spatial linearity or are instead elicited by natural inputs. If the latter, nonlinear spatial integration may be an essential feature of models for RGC responses that successfully generalize across a broad range of stimuli, including natural images. Using an in vitro primate retinal preparation, we show that nonlinear responses are driven by natural stimuli in Off but not On parasol (magnocellular-projecting) RGCs. We then construct two classes of RF model: a spatially linear model and a model that contains spatially offset nonlinear subunits. Using these models alongside electrophysiological recordings, we show that 1) nonlinear parasol responses to natural images are due to the presence of excitatory subunits within the RF; 2) the degree of rectification of subunit output determines whether a nonlinear RF can be approximated as spatially linear in the context of natural stimuli; and 3) accounting for excitatory nonlinear subunits can substantially improve models that predict neural responses to natural images.

III-74. Joint coding of shape and blur in area V4: toward a sufficient representation of natural scenes

Timothy Oleskiw T.D.OLESKIW@GMAIL.COM
Amy Nowack ALN@U.WASHINGTON.EDU
Anitha Pasupathy PASUPAT@U.WASHINGTON.EDU

University of Washington

Computational and psychophysical studies have argued for the importance of boundary contours in visual processing, showing how the statistics of edges can be exploited to form efficient representations of natural scenes. Specifically, the blur of an edge, i.e., the contrast gradient across a boundary, reveals information useful in object segmentation, such as 3D structure and depth. Moreover, it has been shown that edge information, including the magnitude of blur at each edge, is sufficient for encoding naturalistic images. While joint coding of orientation and spatial frequency is well-established in V1, it is unclear whether or how blur information is represented by complex form-sensitive neurons along the ventral pathway. This is largely due to the study of cortical areas beyond V1 utilizing synthetic stimuli with sharp boundaries, or naturalistic stimuli without readily quantified object shape. To further understanding of how blur information is encoded by intermediate stages of the ventral stream, we perform a targeted study of shape-selective neurons in macaque V4. After manual RF characterization, shape stimuli were presented under Gaussian blur, i.e., low-pass filtered at multiple spatial frequencies. Surprisingly, we find neurons tuned to blur, with individual cells exhibiting peak responses for intermediate blur levels. Control experiments demonstrate confounds of stimulus size and intensity cannot account for this modulation. Further, a detailed analysis of boundary curvature suggests blur and shape selectivity as distinct computations, since tuning for shape in V4 cannot be explained by spatial frequency preferences alone. Thus, we report a population of V4 neurons that jointly encode both object shape and boundary blur. A novel model is proposed to explain our data,

giving insight into how V4 responses can form a complete representation of natural scenes. Finally, we report rich dynamics of shape-dependent blur selectivity, consistent with the coarse-to-fine processing of object features within the ventral pathway.

III-75. A canonical circuit for object constancy across visual modalities

David Mely[1] DAVID_MELY@BROWN.EDU
Thomas Serre[2] THOMAS_SERRE@BROWN.EDU

[1]Brown University
[2]Brown Institute for Brain Science

The ability to achieve object constancy over the many image transformations that can affect the appearance of objects (e.g., changes in illumination, viewing distance, etc.) is a hallmark of primate vision. Understanding the neural circuits underlying perceptual constancy is a major goal of visual neuroscience. Starting with idealized populations of unit responses tuned to different modalities (orientation, color, binocular disparity, etc.), we have found that real-world object transformations map onto a reduced set of 'canonical deformations' of those populations (i.e., stretch, shift, gain control). We show that an extension of the normalization model to include contextual interactions via recurrent tuned connections outside the classical receptive field (cRF) (excitatory in the near surround and inhibitory in the far surround) achieves a higher degree of constancy across visual modalities. In addition, the proposed circuit is shown to account for a variety of contextual phenomena across modalities with minimal parameter tuning: color and disparity induction, orientation tilt effects, feature-specific gain control. Our approach is amenable to predictions as it binds meaningful parameters to distinct regimes of behavior found in cellular electrophysiology: e.g., the relative extents of the near vs. far surrounds control the balance between contrast (repulsion by inducer stimuli outside the cRF) and assimilation (attraction by those stimuli).

III-76. Modeling high acuity vision in the presence of fixational eye movements

Alexander Anderson AGA@BERKELEY.EDU
Bruno Olshausen BAOLSHAUSEN@BERKELEY.EDU

University of California, Berkeley

There is strong evidence that our visual system is able to extract out stable features from the constantly jittering images that fall onto our retina without direct knowledge of the eye movements themselves (eg. without an efferent copy of the ocular motor signal). Further, recent research suggests that involuntary eye movements during visual fixation enhance our ability to detect high spatial frequencies. But a crucial question remains: what is the biological mechanism that the brain uses to recover the spatial pattern landing on our retina in the presence of fixational eye movements and imprecise neurons? We propose a computational model based on a Bayesian ideal observer that attempts to estimate the spatial pattern on the retina. From this emerges a neural model containing two populations of cells which we hypothesize to exist in primary visual cortex: one which encodes the spatial pattern using a sparse code and another which tracks eye position and is used to dynamically route information coming from LGN afferents feeding into the pattern cells. Our work extends a previous model (Burak et al. 2010) by incorporating smooth eye movements, continuous valued pixels, a better motion prior, and a more biologically plausible representation of the incoming image. Finally, we propose an experiment to test the core assumptions of our model and comment on how we can evaluate the signal-processing benefits of fixational eye movements.

III-77. A cortical neural network model of visual motion perception for decision-making and navigation

Michael Beyeler[1] MBEYELER@UCI.EDU
Micah Richert[2] RICHERT@BRAINCORPORATION.COM
Nicolas Oros[3] NICOLASYOROS@GMAIL.COM
Nikil Dutt[1] DUTT@UCI.EDU
Jeffrey Krichmar[1] JKRICHMA@UCI.EDU

[1]University of California, Irvine
[2]Brain Corporation
[3]BrainChip Corporation

Behavioral data suggests that humans rely on optic flow to traverse cluttered environments, avoid obstacles, and track objects. Although it is generally assumed that vector-based representations of retinal flow are highly accurate in cortical areas V1 and MT, in order to generate accurate motion estimates the aperture problem must be solved. Here we present a two-stage spiking model of MT that solves the aperture problem, where component-direction-selective (CDS) cells in MT linearly combine inputs from V1 cells that have spatiotemporal receptive fields according to the motion energy model (Simoncelli and Heeger, Vision Research, 1998), and pattern-direction-selective (PDS) cells in MT are constructed by pooling over MT CDS cells with a wide range of preferred directions (Rust et al., Nature Neuroscience, 2006). The full network, which consisted of 153,216 spiking neurons and 40 million synapses, processed 20 frames per second of 40x40 pixel video in real-time using a single off-the-shelf GPU (Beyeler et al., Neuroinformatics, 2014). Simulated neural activity matched direction and speed tuning curves commonly found for neurons in macaque V1 and MT (Rodman and Albright, Vision Research, 1987, Movshon and Newsome, Journal of Neuroscience, 1996). The behavioral response of the network in a motion discrimination task emulated human choice accuracy and the effect of motion strength on reaction time (Roitman and Shadlen, Journal of Neuroscience, 2002). In addition, when used with a physical robot performing a visually-guided navigation task in the real world (Beyeler et al., Neural Netw, 2015), the model produced behavioral trajectories that closely matched human psychophysics data (Fajen and Warren, J Exp Psychol Hum Percept Perform, 2003). The present study demonstrates how a model of MT might build a cortical representation of optic flow in the spiking domain, and shows how these motion signals might relate to perceptual decision-making as well as active steering control.

III-78. Relative weighing of visual features and running speed varies across mouse visual cortex

E. Mika Diamanti EFTHYMIA.DIAMANTI.11@UCL.AC.UK
Aman Saleem AMAN.SALEEM@UCL.AC.UK
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

Neurons in primary visual cortex (V1) of the mouse code not only for features of visual stimuli but also for the animal's running speed. While running, peripheral visual fields experience higher speed changes compared to central visual fields. Do V1 regions analyzing peripheral or central visual fields differ in the way they code for self-motion and visual features? To answer this question we measured calcium signals from V1 using wide-field imaging, while head-fixed mice (Emx1-GCaMP6f) traversed a corridor in virtual reality. We first imaged each animal in closed-loop, where the speed of the virtual corridor (visual speed) matched the animal's running speed. Animals were presented with gray screen in 10% of the trials. We then imaged each mouse in open-loop, by replaying previous sessions to the animal regardless of its behavior. To assess the dependence of calcium signals on virtual-position and speed, we obtained maps of calcium intensity as a function of the two variables. The dependence of visual responses on virtual-position and speed varied across regions. Dependence on speed

better explained responses in the posteromedial region of V1, which analyzes the upper and lateral visual field. Instead, signals in the rest of V1 were more dependent on virtual- position. To assess whether the dependence was on visual or running speed, we turned to open-loop and gray-screen trials. In open-loop, V1 responses were weakly modulated by visual speed. In gray-screen trials, running speed strongly modulated responses all across the visual field. We conclude that in 'closed-loop' conditions, such as during exploration of an environment, responses of V1 are more strongly modulated by the position of visual cues appearing in front and below the animal, whereas regions of V1 analyzing stimuli from the periphery depend more on speed. Open-loop data suggest that visual speed provides a small contribution to this dependence.

III-79. The stabilized supralinear network (SSN) model explains feature-specific surround suppression in V1

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

Columbia University

We have described a mechanism, the Stabilized Supralinear Network (SSN) (D. Rubin, S.D. Van Hooser, K.D. Miller, Neuron 2015 ; Y. Ahmadian, D. Rubin, K.D. Miller, Neural Computation 2013) that can describe a large set of phenomena in V1 and more generally in sensory cortex, including surround suppression, normalization, and their dependence on contrast. While the SSN describes a basic mechanism, results depend on the details of model connectivity. Understanding the underlying circuit is crucial to understanding fundamental brain computations. For a large-scale network we showed (D. Obeid, K.D. Miller, Cosyne Abstracts 2015) that the decrease in inhibition with surround suppression (Ozeki et al., Neuron 2009) required the network to have stronger local connectivity than our previous studies, but all the previous results are preserved. Here we address two new phenomena: (1) the strongest suppression arises when the surround orientation matches that of the center stimulus (Shushruth et al., J. Neurosci. 2012; A.R. Trott, R.T. Born, J. Neurosci. 2015); and (2) a surround with orientation matching one orientation of a plaid center stimulus specifically suppresses that component of response (A.R. Trott, R.T. Born, J. Neurosci. 2015). To match (1), we require local connectivity that is dense as in (D. Obeid, K.D. Miller, Cosyne Abstracts 2015) and more broadly tuned for orientation than that we studied previously; both this connectivity and the connectivities we studied previously match (2). We also study transient responses that emerge from the network in response to center and center surround stimuli.

III-80. Serotonin linearly transforms visual responses in macaque V1

Corinna Lorenz[1] CORINNA.LORENZ@STUDENT.UNI-TUEBINGEN.DE
Lenka Hruba[1] LENKA.HRUBA@CIN.UNI-TUEBINGEN.DE
Torben Ott[1] TORBEN.OTT@UNI-TUEBINGEN.DE
Andreas Nieder[1] ANDREAS.NIEDER@UNI-TUEBINGEN.DE
Paria Pourriahi[1] PARIA.POURRIAHI@CIN.UNI-TUEBINGEN.DE
Hendrikje Nienborg[2,1] HENDRIKJE.NIENBORG@CIN.UNI-TUEBINGEN.DE

[1]University of Tuebingen
[2]Werner Reinhardt for Integrative Neuroscience

Serotonin (5HT) has been implicated in the modulation of a wide variety of behavioral and brain states such as sleep, mood, patience, aversion and reward expectation, and a number of hallucinogenic drugs have high affinity for serotonin receptors (Gonzalez-Maeso and Sealfon, 2009). Nonetheless the functional role of 5HT is still largely unknown and controversial. Serotonin receptors are significantly expressed in macaque V1 (Watakabe et al, 2009), predominantly in the input layers. This suggests a modulatory role of serotonin on visual information processing as early as the cortical input stage. Here, we show that visual responses in macaque V1 are system-

atically altered by the application of 5HT. We measured orientation tuning of single units in the fixating macaque using either drifting luminance gratings (n=39) or by performing orientation subspace reverse correlation (Ringach et al, (1997); n=23). The effect of serotonin on orientation tuning was examined by block-wise iontophoretic application of serotonin. In control experiments we iontophoretically applied pH-matched NaCl (n=13, n=4, respectively). We found that for both, orientation subspace maps and drifting gratings 5HT predominantly decreased the visual responses without systematically affecting the tuning properties such as the orientation tuning width and preferred orientation. The changes were well captured by multiplicative and modest additive transformations of the tuning curves. A simple threshold-linear model in which 5HT induced a subtractive shift of the membrane potential, e.g. reflecting increased un-tuned inhibition, could replicate the experimental findings. Our results may also provide an explanation for visual hallucinogenic effects involving serotonin receptors. Theory proposes that hallucinations arise when the balance between top-down internally generated beliefs and feed-forward sensory signals is shifted towards top-down beliefs, e.g. (Jardri and Deneve, 2013). Our results suggest that the activation of serotonin receptors in V1 may produce such a shift by decreasing the gain of the feed-forward visual input.

III-81. Mixed E-I interactions between V1 cells may reflect a Bayesian edge probability calculation

Gabriel Mel[1] MELDEFON@USC.EDU
Chaithanya A Ramachandra[2] CHAITHANYA@GMAIL.COM
Bartlett W Mel[1] MEL@USC.EDU

[1]University of Southern California
[2]Eyenuk, LLC.

A key computation in visual cortex is the extraction of object contours, where the first stage of processing - local edge detection - is commonly attributed to orientation-tuned V1 simple cells. However, using a labeled natural image database, we found that a conventional model of a simple cell, consisting of a Gabor-like linear filter stage followed by a divisive normalization, shows poor edge detection performance in natural images. To improve upon this performance, the cortical circuit can in principle, integrate information from a local population of cells with overlapping receptive fields. It is not clear, however, how a population of oriented cells should be decoded to calculate local edge probability, nor what type of circuit is capable of doing so. To gain insight into this population decoding problem, we collected filter statistics from neighborhoods in natural images when an edge was present or not present at a ‘reference’ location. We then examined the form of the log-likelihood (LL) ratios for surrounding filters, which, within a Bayesian formulation, capture whether a given neighbor cell contributes positively or negatively to the reference location edge probability. We found the LL ratios, considered as functions of neighboring cell responses, were nearly all bump shaped, indicating that a cell at the reference location should receive an intensity-tuned input from each simple cell. Intensity-tuned interactions between cells can be produced by subtraction of sigmoidal excitatory and inhibitory curves with shifted thresholds. Based on this mechanism, we propose a simple cortical circuit model in which neighboring simple cells excite and inhibit a decoding neuron to implement the intensity-tuned interactions specified by the LL ratios. According to this model, the mixed excitatory-inhibitory interactions frequently found between nearby cells in V1 could reflect an underlying Bayesian edge probability calculation.

III-82. Emergence of an optimal command for orienting behavior

Fanny Cazettes[1] FANNY.CAZETTES@PHD.EINSTEIN.YU.EDU
Brian Fischer[2] FISCHER9@SEATTLEU.EDU
Jose Pena[1] JOSE.PENA@EINSTEIN.YU.EDU

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

How the brain, often a biased estimator, translates unreliable sensory information into adaptive behavioral responses is still a matter of debate. The systematic underestimation of peripheral sound locations by barn owls provides a means to address this question. We have previously shown that neurons in the owl's midbrain map of auditory space are tuned to the most reliable auditory cues and that the shape of tuning curves is informative about the degree to which a cue can be trusted. The next question is how the neural population is read out to capture cue reliability in the behavioral response. It has been proposed that integrating reliability and prior knowledge about the sensory input may make the owl's bias adaptive, in a statistically optimal fashion. However, a dedicated neural circuit decoding a Bayesian estimate has not been reported. Here, we tested the hypothesis that premotor neurons downstream from the midbrain map of auditory space guide the owl's head-orienting behavior by encoding a Bayesian estimate of sound direction. We found that neural activity in the premotor brainstem nucleus approximates the optimal orienting response given sound direction and cue reliability, and this response predicts the owl's behavioral bias. We demonstrate that this coding emerges from the weighted convergence of upstream projections. We further show that manipulating the sensory input yields changes in both premotor brainstem responses and the owl's behavioral bias consistent with Bayesian inference. This work demonstrates computationally and experimentally how a sensory representation can be read out to guide an adaptive behavior.

III-83. A new song mode in Drosophila melanogaster

Philip Coen[1,2] PIPCOEN@GMAIL.COM
Jan Clemens[1] CLEMENSJAN@GMAIL.COM
Mala Murthy[1] MMURTHY@PRINCETON.EDU

[1]Princeton University
[2]University College London

Acoustic communication signals are prevalent throughout the animal kingdom. Many species produce a limited number of acoustic patterns, and experimenters classify each signal as a particular syllable or mode. It is presumed that these distinct signal types arise from different, but reproducible, patterns of neural activity. Therefore, accurately establishing the acoustic repertoire of an organism, and correctly classifying each signal, is a critical step in dissecting the underlying neural circuits. For more than fifty years, researchers have separated Drosophila melanogaster courtship song into two modes—pulse and sine. Here, we use quantitative methods to demonstrate that 'pulse' is actually composed of two song modes which are distinct both in their acoustic structure and the sensory conditions that elicit them. During courtship, male flies use wing vibration to generate pulses trains with an inter-pulse interval of $\sim 35ms$. We analyzed >500,000 pulses produced by Drosophila melanogaster males from 8 geographically diverse regions. For all strains, pulse shapes separated into distinct clusters, a lower frequency pulse (P1, $\sim 210Hz$) and a higher frequency pulse (P2, $\sim 300Hz$). Pulse trains were primarily composed of a single pulse type, suggesting these pulses reflect two stable modes at the level of the central pattern generator. Furthermore, we demonstrate that males produce P1 and P2 under different sensory conditions, biasing toward P2 at larger distances from the female. As all strains produced P2 at a higher intensity, switching between pulse modes may allow the male to communicate over a larger range of distances. Overall, we establish that Drosophila melanogaster courtship song consists of three distinct song modes which are selected in response to changing sensory stimuli. This not only has important ramifications within the field, where all previous studies have decomposed melanogaster courtship song into two modes, but also highlights the importance of using computational methods to analyze auditory signals.

III-84. A model-based EEG approach for investigating the hierarchical nature of continuous speech processing

Giovanni Di Liberto DILIBERG@TCD.IE
Edmund Lalor EDLALOR@TCD.IE
University of Dublin

That cortical sensory systems are organized in a hierarchical structure is reasonably well established. In the context of human speech it has been suggested that such an organization could explain how acoustically variable inputs can be perceived as categorical speech units. A number of studies have been conducted to reveal the precise mechanisms that underlie this hierarchical system; however the analysis methodologies have limited the stimuli to unnaturalistic discrete units of speech, such as isolated syllables or words. An approach for indexing the neurophysiology of this hierarchical processing in the context of natural, continuous speech has been recently introduced (Di Liberto et al., Current Biology, 2015). Specifically, the relationship between continuous speech and low-frequency EEG responses was estimated using a multivariate regression model based on different speech representations. This mapping was shown to be best described when speech was represented using both its low-level spectrotemporal information and a categorical labeling of its phonetic features. While this approach results in a quantitative measure of scalp neural activity related to phonetic features, it remains unclear to what extent this measure reflects speech-specific processing. Here, we outline an experiment aimed at investigating the speech-specific nature of our model-based neural measure. The intelligibility of 10-s speech stimuli was degraded using noise vocoding. Each vocoded stimulus was presented twice with an intervening presentation of the original clean speech version of the same stimulus. As such, the second presentation of the vocoded stimulus was primed by the clean speech and was found to be significantly more intelligible on a match-to-sample task. Our model-based neural measure was found to be significantly correlated with a behavioral measure of intelligibility, suggesting that we have isolated a dependent measure of speech-specific processing at the phonetic level.

III-85. Learning and predicting the acoustic consequences of movement

David Schneider SCHNEIDER@NEURO.DUKE.EDU
Richard Mooney MOONEY@NEURO.DUKE.EDU
Duke University

We constantly make predictions, ranging from whether it will rain today to who will be the next president. But some of the most important predictions that we make are much less obvious, such as what my voice will sound like when I speak or what the next note will sound like when I strike a key on the piano. This ability to predict the acoustic consequences of our actions is vital for learning and maintaining complex behaviors such as speech. Yet we understand remarkably little about how the brain learns to predict that sounds that our movements make. To address this question, we have developed an acoustic virtual reality (VR) platform in which arbitrary sounds can be mapped onto arbitrary movements and with which we can study neural circuits in parallel with behavior as mice learn to predict the sound that a movement produces. In one such VR environment, a mouse produces a series of predictable tone pips while running on a treadmill, with a fixed pitch and with a rate proportional to the mouse's speed. Multi-electrode array recordings made after ~ 7 days in this VR environment show that responses of auditory cortical neurons to predictable tones are suppressed during running, revealing the operation of an experience-dependent filter that is dynamically engaged during movement. Predictive suppression in the auditory cortex appears to be largely subtractive, suggesting that unique interneuron populations might be important for suppressing self-generated sounds. Preliminary experiments indicate that this predictive suppression is not present in the auditory thalamus, suggesting that it arises de novo in the cortex. Finally, longitudinal 2-photon calcium imaging shows that auditory cortical neurons adapt in parallel with sensorimotor experience over several days to selectively suppress self-generated sounds. These experiments provide a platform for exploring how neural circuits learn, store and engage movement-related predictions.

III-86. Sensitivity to sound texture statistics in auditory cortex

Richard McWalter[1] RMCW@ELEKTRO.DTU.DK
Jens Hjortkjaer[1] JHJORT@ELEKTRO.DTU.DK
Hartwig Siebner[2] HARTWIG.SIEBNER@DRCMR.DK
Torsten Dau[1] TDAU@ELEKTRO.DTU.DK
Kristoffer Madsen[2] STOFFER@DRCMR.DK

[1]Technical University of Denmark
[2]Danish Research Centre for Magnetic Resonance

Natural sounds, like birds chirping or a stream flowing, are characterized by higher-order statistical regularities that are stationary over long time windows. Behavioral evidence presented by McDermott and Simoncelli (2011) suggests that such 'sound textures' can be recognized and stored in memory via their statistics, but potential cortical mechanisms for encoding such statistics have not been explored. Using fMRI and sound texture synthesis, we investigated the neural processing of sound textures in human auditory cortex. We used a biophysically inspired model of the auditory periphery to analyze the statistics of recorded natural sounds and then inverted the model to synthesize sounds by matching the time-averaged statistics. In a first experiment, we contrasted BOLD responses between synthesized sound textures and spectrally matched noise. Across a range of different sound textures, we found that regions of auditory cortex, including Heschl's Gyrus and planum temporale responded selectively to stimuli with higher-order statistics. In both primary and secondary auditory cortex, we also observed a pattern of BOLD-signal adaption to repeating presentations of texture exemplars with identical statistics that was absent for repetitions of spectrally matched noise. To quantify these responses further, a second experiment was conducted with sounds that were synthesized by cumulatively introducing the statistics. We found that voxels in planum temporale and posterior superior temporal gyrus responded parametrically to the introduction of higher-order statistics across different sound textures. Moreover, the neural response to different subsets of statistics correlated with their relevance in a behavioral identification task conducted separately, where performance increased as more statistics were included. The results point towards a cortical mechanism that identifies sound textures via their time-averaged statistics.

III-87. Facilitation of auditory cortical responses in mice playing the mouseophone

Uri Livneh URILIV@GMAIL.COM
Anthony M Zador ZADOR@CSHL.EDU

Cold Spring Harbor Laboratory

During many tasks, including those that require attention, the perception of behaviorally relevant sensory inputs is prioritized over irrelevant inputs, but the neural mechanisms that implement this prioritization remain unknown. A central challenge in investigating the neural mechanisms that underlie sensory prioritization is the difficulty in training mice to flexibly prioritize one sensory component over another. Here we present a novel computer-controlled 'musical' instrument—the 'mouseophone'—that enables the study of sensory prioritization in head-fixed mice. Mice were trained to move a rod to a target position using auditory feedback. The auditory feedback was determined by a transfer function (TF) that mapped each rod position into the repetition rate (12-96Hz) of brief (10ms) sound. Perturbation of the auditory feedback impaired performance, indicating that trained mice indeed used the auditory feedback to guide movement. To test whether auditory feedback was behaviorally prioritized, we introduced a distractor. The distractor, presented at the same time as the feedback signal, was not coupled to the movement of the mouse. Performance was not affected by the presence of the distractor, suggesting that mice used the feedback to guide movement. We next used extracellular recording to investigate neural correlates of prioritization in the auditory cortex of mice performing the task. We found that in a surprisingly high fraction (46%) of neurons, responses evoked by the acoustic feedback were selectively facilitated when compared to the responses evoked to the very same auditory stimulus presented as a distractor. These modulations were

often strong enough to allow classification of the frequency of the feedback stimuli that was used in single trials, indicating that sensory prioritization for motor-locked feedback is fundamentally intertwined with the representation of sounds in the auditory cortex. These results suggest that motor-related inputs play a central role in controlling sensory prioritization in the auditory cortex.

III-88. The effect of resonance properties on network oscillations through electrical gap junction coupling

Xinping Li[1,2] XL365@NJIT.EDU
Yinbo Chen[1] AIGNCY@GMAIL.COM
Horacio G Rotstein[1] HORACIO@NJIT.EDU
Farzan Nadim[1,2] FARZAN@NJIT.EDU

[1]New Jersey Institute of Technology
[2]Rutgers University, Newark

Neurons often produce a maximal subthreshold voltage response to oscillatory current inputs at a non-zero input frequency (fres), a property known as membrane potential resonance. Typically, resonance is measured by using the impedance Z(f) of a neuron over a range of frequencies (f). Recent studies have suggested that fres of constituent neurons is a strong indicator of the network frequency (fnet). Nevertheless, how the resonance property influence fnet remains unclear. We focus on networks of neurons that are electrically coupled and examined the hypothesis that, in such a network, biophysical parameters that shift fres also shift fnet in the same direction. Additionally, we proposed that an increase in the resonance power (Q=Zmax-Z(0)) of participating resonator neurons increases this dependence. We tested our hypothesis in an electrically coupled network consisting of a pacemaker neuron (having intrinsic oscillations) and a resonator that does not necessarily have intrinsic oscillations. Using a two-cell model of such a network, we show that only the impedance profile of the resonator, and not the specific biophysical parameters, influenced fnet. Additionally, this influence was independent of oscillator type. To examine our hypothesis experimentally, we connected the pacemaker neurons in the oscillatory pyloric network of the crab C. borealis to a model resonator via electrical coupling using the dynamic clamp technique. The attributes fres and Zmax of the resonator can be varied independently by changing the model parameters. We found that by shifting the fres of the resonator, fnet shifted in the same direction and this effect was enhanced by increasing Zmax when Z(0) was fixed. Our results provide experimental support that resonance frequency and power can strongly influence the network oscillation frequency and therefore modulators may directly target these attributes in order to influence network activity.

III-89. Intensity invariant readout of olfactory bulb output is facilitated by an interglomerular circuit

Arkarup Banerjee[1,2] ABANDYOP@CSHL.EDU
Honggoo Chae[1] HGCHAE@CSHL.EDU
Dinu F Albeanu[1] ALBEANU@CSHL.EDU

[1]Cold Spring Harbor Laboratory
[2]Watson School of Biological Sciences

Perception is often invariant for stimulus identity (e.g: face recognition) even when other stimulus features (e.g: light level, viewing angle) vary widely. In the mammalian olfactory system, little is known about the neural mechanisms that extract odor identity, while tolerating fluctuations in concentration. As a first step, we hypothesized that a specific gain-control circuit in the olfactory bulb (OB) mediated by DAT+ cells, reformats the OB output to facilitate efficient read-out of concentration-invariant odor identity by target cortical areas. To test this, we monitored the activity of numerous OB output (Mitral/Tufted) cells to odors across varying concentrations (3 orders of

magnitude) using 2-photon microscopy. Individual mitral cells exhibited a diverse range of concentration response profiles (CRFs)—some monotonically increased, or decreased with increasing concentration, while others peaked at intermediate concentrations. Importantly, the mean population response increased only modestly across this large concentration range. A simple model assuming divisive normalization implemented by the DAT+ cells was able to generate the diversity of the experimentally observed CRFs. We visualized the neural population trajectories using PCA and found that for a given odor, all sampled concentrations spanned low-dimensional manifolds. Additionally, across a large odor panel, mitral cells exhibited significantly larger dimensionality compared to that of tufted cells, highlighting the differences between these two OB output channels. To assess the ability of cortical targets to correctly identify odors, we trained a linear decoder with sparse and non-negative weights to classify odor identity irrespective of concentration. The cross-validated performance was $>$80%, which significantly dropped ($<$60%) when DAT+ cells were specifically ablated. We conclude that a specific interneuron circuit in the glomerular layer formats the OB population output so as to facilitate concentration invariant odor identification by the cortex. This may be the first of many such transformations that ultimately lead to perceptually stable behavior.

III-90. Active sensation disrupts correlations in S1 and M1 networks in the mouse neocortex–a sensorimotor account

Gregory I Telian[1] GITELIAN@BERKELEY.EDU
Mayur Mudigonda[2] MUDIGONDA@BERKELEY.EDU
Jesse Livezey[1] JESSE.LIVEZEY@BERKELEY.EDU
Ryan Zarcone[1] ZARCONE@BERKELEY.EDU
Michael DeWeese[1] DEWEESE@BERKELEY.EDU
Hillel Adesnik[1] HADESNIK@BERKELEY.EDU

[1]University of California, Berkeley
[2]Redwood Center for Theoretical Neuroscience

Animals function in a 3D world where repeatable, robust action drives their survival. Consequently, it is of great importance to understand sensorimotor representations and how sensory stimuli are represented and transformed into motor actions. Recent work [Matyas et al, 2010] has shown that there exists a very tight coupling between primary somatosensory (S1) and motor (M1) neurons in the mouse cortex but very little has been done to explain what sort of computations these populations of neurons might be performing. We present a preliminary analysis of new experimental data that was collected simultaneously from mouse S1 and M1. Network connectivity of active neurons was inferred from the coupling matrix of an Ising model fit to binary spiking data. We find that when mice actively palpate an object using multiple whiskers, S1 units become weakly coupled while M1 units appear to reorganize their couplings. These stimulus induced network decorrelations may be carrying out the computations involved in sensorimotor transformations and warrant further study.

III-91. The motor side of the sensory loop: Fast changes in active touch after sensory cortex stimulation

Jason Ritt[1] JRITT@BU.EDU
Joseph Schroeder[1] JBS7@BU.EDU
Gregory I Telian[2] GITELIAN@BERKELEY.EDU
Vincent Mariano[1] VINN1017@BU.EDU

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

In active sensing, afferent input must be integrated with self-motion to be properly interpreted by the organism, incorporating closed-loop behavioral selection of information during sensory acquisition. The rodent whisker tactile system is a key model for active touch. During exploration, mice primarily employ 5-20 Hz whisking, but adjust

sensing motions based on goals and recent afferent information. We examined head and whisker motions of unconstrained mice performing tactile search for randomly located rewards, and found that mice select from a range of active sensing strategies, based at least in part on the behavioral context of whisker contacts. In particular, mice selectively employed a strategy we term contact maintenance, where whisking is modulated to counteract head motion and sustain repeated contacts, but only when doing so is likely to be useful. The selection of sensing strategies and timing of whisker repositioning prior to head motion suggests the possibility of higher level control, beyond solely reflexive mechanisms. To investigate possible primary somatosensory cortex (SI) influences on whisk-by-whisk motion, we delivered optogenetic feedback to SI, time locked to whisking as estimated from facial electromyography. We found stimulation increased the frequency and regularity (or periodicity) of whisking, emulating behavior when actively contacting objects. Additionally, stimulation induced small, short latency motions similar to those in a previous study in quiescent head fixed animals, except that we induced protractions rather than retractions. We found a substantial reduction in SI responsiveness during retractions, possibly encoding that contact is more likely during forward motion of the whiskers, and downstream areas may show greater sensitivity to SI activity during protractions. These findings challenge the traditional notion of sensory coding in primary cortical areas as derived solely from feed forward inputs from the periphery, and address a causal role of sensory cortex activity in guiding sensing motions during active touch.

III-92. Optimal learning with redundant synaptic connections

Naoki Hiratani[1] HIRATANI@BRAIN.RIKEN.JP
Tomoki Fukai[2] TFUKAI@RIKEN.JP

[1]The University of Tokyo
[2]RIKEN Brain Science Institute

Recent experimental studies suggest that, in cortical microcircuits of the mammalian brain, the majority of neuron-to-neuron connections are realized by multiple synapses. For instance, in the barrel cortex of juvenile mice, mean number of synapses per connection is estimated to be around 10 [1]. However, little is known on the functional benefit of having such redundant synaptic connections. Here, we show that redundant synaptic connections enable near-optimal learning in cooperation with synaptic rewiring. By constructing a simple dendritic neuron model, we demonstrate that, in multi-synaptic connections, synaptic plasticity approximates a particle-filtering algorithm, and wiring plasticity corresponds to its resampling process. The derived synaptic plasticity rule reconciles with dendritic position dependence of spike-timing-dependent plasticity observed in previous experiments [2]. In particular, our study reveals a functional merit of anti-Hebbian plasticity at distal synapses. The model also explains why two synapses projected to different dendritic branches from the same axon seldom show spine-size correlation, while those on the same branch exhibit tight size correlation [3]. The proposed framework is also applicable to unsupervised learning in recurrent circuits. In conclusion, our study provides a novel conceptual framework for synaptic plasticity and rewiring by focusing on redundancy in synaptic connections. [1] Markram et al., Cell 163, 456 (2015). [2] Letzkus et al., J Neurosci 26, 41 (2006). [3] Bartol et al., BioRxiv (2015).

III-93. Two types of cortical interneurons differentially modulate behavioral frequency discrimination acuity

Jennifer Blackwell BLAJE@MAIL.MED.UPENN.EDU
Mark Aizenberg MARKAIZ@MAIL.MED.UPENN.EDU
Laetita Mwilambwe-Tshilobo LAETITIA@MAIL.MED.UPENN.EDU
Sara Jones JONESSAR@SAS.UPENN.EDU
Ryan G Natan RNATAN@MAIL.MED.UPENN.EDU
Maria Geffen MGEFFEN@MAIL.MED.UPENN.EDU

University of Pennsylvania

The ability to discriminate between tones of different frequencies is fundamentally important for everyday hearing. Primary auditory cortex (A1) regulates behaviors that rely on frequency discrimination (Aizenberg and Geffen, 2013), but the underlying neural mechanisms are poorly understood. Frequency tuning of cortical excitatory neurons are thought to be shaped by the interplay of excitatory and inhibitory inputs. In the cortex, the two most common classes of inhibitory interneurons are parvalbumin-positive (PV) interneurons and somatostatin-positive (SOM) interneurons. PVs target the soma and initial axon segment, while SOMs target distal dendrites. Therefore, these two interneuron classes may differentially affect responses of excitatory neurons. We recently found that photo-activation of PVs enhanced tone-evoked responses of excitatory neurons, which was correlated with an improvement in behaviorally measured frequency discrimination acuity (Aizenberg et al., PLoS Biology, in press). We now find that photo-activation of SOMs diminished tone-evoked responses in the excitatory neurons, by suppressing tone-evoked responses more strongly than spontaneous activity. Interestingly, photo-activation of SOMs also increased the frequency selectivity of excitatory neurons and improved behaviorally measured frequency discrimination acuity. Combined, we find that activation of both PVs and SOMs improves frequency discrimination acuity, but exerts a differential effect on frequency responses of excitatory neurons. These findings are consistent with the interpretation that PVs and SOMs carry out complementary roles in shaping frequency selectivity in the auditory cortex.

III-94. A feedback disinhibition circuit for behavioral choice revealed by connectomics

Casey Schneider-Mizell[1,2] SCHNEIDERMIZELLC@JANELIA.HHMI.ORG
Tihana Jovanic[1] JOVANICT@JANELIA.HHMI.ORG
Mei Shao[1,2] SHAOM@JANELIA.HHMI.ORG
Marta Zlatic[1,2] ZLATICM@JANELIA.HHMI.ORG
Albert Cardona[1,2] CARDONAA@JANELIA.HHMI.ORG

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

Experimental evidence has suggested two approaches for how nervous systems select and implement behavioral sequences. In one, input activates the first in a chain of neuronal populations, each of which drives one behavior while also activating the population that drives the next. In the other, neuronal populations for many behaviors are co-excited, but heterogeneities in input gain and lateral inhibition produce sequential activation. However, identifying the specific neuronal circuitry implementing either model has proved elusive due to difficulties in mapping the circuit context of appropriate neurons. Here, we identify key circuitry underlying sequence generation and behavioral choice in the Drosophila larval reaction to an air puff by combining a synaptic-level reconstruction of neuronal circuits from electron microscopy (EM) with cell-type specific genetic manipulation. We identified two 2nd order relay neurons that drive distinct, temporally ordered behavioral components and a complex network of inhibitory interneurons well-posed to differentially modulate their activity. The neuron driving the earlier component received more excitatory synaptic inputs from sensory neurons and fewer inhibitory inputs than the neuron driving the later component, consistent with higher input gain eliciting earlier activation. However, EM reconstruction also identified feedback interneurons downstream of the relay neurons that could disinhibit this feedforward inhibition. Using a computational model based on the EM reconstruction, we propose that feedback disinhibition can dynamically increase the gain of the second component in the sequence, resulting in sharper selection of and transition between behaviors. Additionally, inactivation of interneurons in the model predicts changes observed after similar manipulations in behaving animals. This work links specific neuronal circuits to models of behavioral choice and sequence generation. We further speculate that disinhibitory chains, conceptually similar to excitatory chains, are a powerful circuit motif and could be found more widely in circuits underlying sequential behavior.

III-95. Model-based evaluation of the role of nonlinear dendrites in cortical computations

Balazs Ujfalussy[1,2] BALAZS.UJFALUSSY@GMAIL.COM
Mate Lengyel[3] M.LENGYEL@ENG.CAM.AC.UK
Tiago Branco[2] TBRANCO@MRC-LMB.CAM.AC.UK

[1]Lendulet Laboratory of Neuronal Signaling
[2]MRC Laboratory of Molecular Biology
[3]University of Cambridge

Establishing the input-output function of single neurons is a key step for understanding neural circuit computations. Neurons have dendrites with a large repertoire of biophysical mechanisms supporting diverse forms of nonlinear integration, but it is unclear how these properties contribute to the overall input transformation of the cell under realistic input conditions with complex spatio-temporal structure. Current experimental techniques allow the simultaneous monitoring of a substantial fraction of the synaptic inputs and the somatic output of a single neuron in awake, behaving animals, but standard approaches characterising input-output transformation along a single dimension (e.g. stimulus strength or inter-stimulus interval) do not generalise to complex and dynamically changing stimuli. Here we present a novel model-based approach based on a hierarchical extension of the generalised linear model (hGLM) that predicts both the sub-threshold and spiking response of neurons to arbitrary and complex spatio-temporal patterns of synaptic inputs. We validated our approach by studying dendritic integration in a biophysically detailed compartmental model that reproduces the main features of dendritic and somatic voltage activity recorded in vivo. Fitting the hGLM is a non-convex optimization problem and we took an iterative approach to facilitate the convergence to the global optimum and the interpretation of the optimal parameters. We found that around 5-15% of the total variance of the somatic membrane potential arises from dendritic nonlinearities and is only captured by models with multiplexed nonlinearities (multiple nonlinear channels per dendritic subunit) as models with simple nonlinearities only marginally outperformed point neuron models. Our approach provides an intuitive description of dendritic information processing performed by neurons receiving large barrages of synaptic inputs and thus paves the way for understanding the role of dendrites in the computation performed by neuronal circuits.

III-96. Strong functional connectivity of parvalbumin-expressing cortical interneurons

Dimitri Yatsenko[1] DVYATSEN@BCM.EDU
Emmanouil Froudarakis FROUDARA@BCM.EDU
Alexander Ecker ALEXANDER.ECKER@UNI-TUEBINGEN.DE
Robert Rosenbaum[2] RROSENB1@ND.EDU
Kresimir Josic[3] JOSIC@MATH.UH.EDU
Andreas Tolias[1] ASTOLIAS@BCM.EDU

[1]Baylor College of Medicine
[2]University of Notre Dame
[3]University of Houston

The morphological and electrophysiological properties of parvalbumin-expressing inhibitory interneurons (PV+ neurons) suggest their role as synchronizers and normalizers of the local cortical microcircuit. PV+ cells are thought to average the local activity and dynamically regulate its overall level. In apparent agreement with this model, previous studies have shown stable patterns of correlations of the spiking activity of the PV+ neurons among themselves and with the local excitatory cells. However, we have previously shown that, in sufficiently dense recordings, estimates of the partial pairwise correlations of the spiking activity can yield a more insightful picture of interactions in the circuit, or its functional connectivity. Using high-speed 3D two-photon imaging of calcium signals and genetically encoded fluorescent markers of PV+ neurons, we recorded the activity of the majority of neurons in 200 um x 200 um x 100 um volumes in layers 2/3 and 4 of mouse visual cortex during visual stimulation. If PV+ neurons simply pooled the activity of the local circuit, their activity would be predicted from the local circuit and the partial correlations among the PV+ neurons would all but vanish. Surprisingly, we found that the partial pairwise correlations among the PV+ cells were exceptionally high. In fact, the partial pairwise correlations enhanced the differentiation of PV+ neurons from other cell types. The average partial pairwise correlation between PV+/PV+ pairs was 4.9 times higher than between PV-/PV- pairs whereas the average noise correlations differed by the factor of 1.5. This effect was insensitive to the choice of the temporal scales of correlation analysis. Although other explanations cannot yet be excluded, the present finding may suggest that the correlations among the PV+ neurons are shaped predominantly by structured input from outside the local circuit such as, for example, by input from layer 5.

III-97. Population density techniques for modeling neural populations

Yi Ming Lai SCSYML@LEEDS.AC.UK
Marc de Kamps M.DEKAMPS@LEEDS.AC.UK
University of Leeds

Many mesoscopic models of the brain are based on mean-field models, such as rate-based approaches or population density techniques (PDTs). PDTs have a long-standing history in computational neuroscience, and form a useful bridge between simulating networks of individual spiking neurons and neural mass models. They are rigorously derived from spiking neural models and therefore provide a level of biological realism, while being lightweight enough that one can develop circuits of them and create complex cognitive models that can be used to make high-level predictions for experiments. However, existing techniques for modelling population densities are not without shortcomings. Here we present a novel quasi-analytic method for point neuron models that extends their domain of applicability in several important ways. We demonstrate it with 1-dimensional neuron models, although it is in principle also applicable to higher dimensional models. We are able to deal with arbitrary distributions of synaptic efficacies (including large jumps which go beyond Fokker-Planck or Langevin equations), excitatory and inhibitory connections, non-Markovian statistics, as well as transient behaviour. This is particularly important as the experimental literature often reports a wide range of spike train statistics, and recent work suggests that this can cause significant qualitative differences in the behaviour of neuronal networks. As such, we are developing a versatile open-source framework which will allow users to simulate time-dependent behaviour of large-scale networks of populations, with minimal changes in the code required to switch between neuronal models, different spike train statistics or various other parameters. As an example, we simulate the transient response of a population to spike trains from both a Poisson process as well as gamma-distributed inter-spike intervals, which is analytically very difficult, and discover that both transient behaviour and steady-state firing rates can be affected by the inter-spike statistics even when the mean input intensity is identical.

III-98. Whole-brain dynamics and statistics in freely moving C. elegans

Ashley Linder[1] LINDER@PRINCETON.EDU
Jeffrey Nguyen[1] JNGUYEN@EXCHANGE.PRINCETON.EDU
Joshua Shaevitz[1] SHAEVITZ@EXCHANGE.PRINCETON.EDU
Andrew Leifer[1,2] LEIFER@PRINCETON.EDU

[1]Princeton University
[2]Lewis Sigler Institute

Understanding the relationships between structural connectivity, brain-wide neural dynamics, and resulting behavior requires the ability to simultaneously record activity from all neurons with cellular resolution and observe animal behavior that this activity generates, while also having knowledge of underlying neural connectivity. We have developed a new instrument to obtain such datasets for the nematode C. elegans, and we use this instrument to look at the relationship between activity of the majority of neurons in the head, their functional connectivity, and the behavior that is generated by these global brain dynamics1. Using spinning disk-confocal microscopy, we record fluorescence from GCaMP6s with single cell resolution at 6 head volumes per second, while also recording the behavior of the animal as it moves freely. We observe calcium transients from more than 100 neurons with cellular resolution for more than four minutes in a given animal, and correlate this to the animal's naturalistic, spontaneous behavior. Across worms, multiple neurons show significant correlations with modes corresponding to forward, backward, and turning locomotion. Using a regularized least squares approach, we estimate the functional connectome from the activity of all the neurons, finding that the functional connectivity is not as sparse as the known structural connectome would suggest. This finding indicates that neural activity is distributed widely throughout the brain, and so consequentially we performed principal component analysis on the neural activity to describe underlying dynamics in a low-dimensional state space. We find that certain behaviors are correlated with specific components of the trajectory, while other behaviors occur throughout the entire space. This is a first step towards a quantitative understanding that relates neural structure, neural function, and behavior in awake and unrestrained animals.

III-99. Accuracy of circuit reconstruction from spikes in the memory versus sensory regime

Abhranil Das ABHRANIL.DAS@UTEXAS.EDU
Ila R. Fiete ILAFIETE@MAIL.CLM.UTEXAS.EDU

The University of Texas at Austin

Recording spikes is substantially easier than measuring connectivity in neural circuits, yet understanding mechanism ultimately requires knowledge of connectivity. Therefore, methods for network reconstruction from spiking data are powerful tools in the neuroscientist's arsenal. Such methods include maximum entropy-based inverse Ising inference and generalized linear models, both quasi-Bayesian techniques that can 'explain away', i.e. correctly infer the non-existence of synapses between correlated but unconnected nodes. They have been applied to low-level sensory systems, e.g. retinal data, with success measured by accurate prediction of unseen stimulus-response data. What can we hope to learn about mechanism or connectivity in non-sensory systems using such approaches? Memory systems can exhibit recurrently stabilized activity patterns with large, long-range correlations that are harder to explain away. To address this question, we consider inverse Ising inference on a simple 1D pattern-forming recurrent network with local connections, where it is easy to differentiate between circuit reconstruction noise and systematic failures in explaining away. We show that at low recurrent weights, reconstruction accuracy is limited only by noise. At strong weights, neurons without direct connections strongly correlate via pattern formation, resulting in systematic inference of artifactual long-range couplings that are relatively impervious to more data. Consistent with literature on the temperature-dependence of inverse Ising inference, accuracy is maximized at an intermediate critical point with weights that lie in a strong-amplification sensory regime. Requirement for data increases sharply into the self-sustained memory regime, a serious practical problem; with

realistic amounts of data, estimates will be strongly biased toward high connectivity. We further investigate the known problem of inference from subpopulations failing to explain away correlations through the mutual missing nodes. We demonstrate that the problem is exacerbated in memory systems compared to sensory systems. These results generalize to different circuit inference techniques and network architectures.

III-100. Stitching neural activity in space and time: theory and practice

Marcel Nonnenmacher[1,2] MARCEL.NONNENMACHER@CAESAR.DE
Lars Buesing[3] LARS@STAT.COLUMBIA.EDU
Artur Speiser[1] ASPEISER@UNI-OSNABRUECK.DE
Srinivas Turaga[4,5] TURAGAS@JANELIA.HHMI.ORG
Jakob H Macke[6] JAKOB.MACKE@CAESAR.DE

[1]research center caesar, Bonn
[2]Bernstein Center Tubingen
[3]Columbia University
[4]Janelia Farm Research Campus
[5]Howard Hughes Medical Institute
[6]Neural Systems Analysis, caesar, Bonn

Simultaneous recordings of the activity of large neural populations are extremely valuable as they can be used to infer the dynamics and interactions of neurons in a local circuit, shedding light on the computations performed. It is now possible to measure the activity of hundreds of neurons using in-vivo 2-photon calcium imaging. However, this experimental technique imposes a trade-off between the number of neurons which can be simultaneously recorded, and the temporal resolution at which the activity of those neurons can be sampled. Previous work (Turaga et al 2012, Bishop & Yu 2014) has shown that statistical models can be used to ameliorate this trade-off, by 'stitching' neural activity from subpopulations of neurons which have been imaged sequentially with overlap, rather than simultaneously. This makes it possible to estimate correlations even between non-simultaneously recorded neurons. In this work, we make two contributions: First, we show how taking into account correlations in the dynamics of neural activity gives rise to more general conditions under which stitching can be achieved, extending the work of (Bishop & Yu 2014). Second, we extend this framework to stitch activity both in space and time, i.e. from multiple subpopulations which might be imaged at different temporal rates. We use low-dimensional linear latent dynamical systems (LDS) to model neural population activity, and present scalable algorithms to estimate the parameters of a globally accurate LDS model from incomplete measurements. Using simulated data, we show that this approach can provide more accurate estimates of neural correlations than conventional approaches, and gives insights into the underlying neural dynamics.

III-101. Automated unsupervised decoding of long-term, naturalistic human neural recordings with video

Xin Ru Wang NANCYWANG1991@YAHOO.CA
Jared Olson JAREDOL@UW.EDU
Jeffrey Ojemann JOJEMANN@UW.EDU
Rajesh Rao RAO@CS.WASHINGTON.EDU
Bing Brunton BBRUNTON@UW.EDU

University of Washington

Fully automated decoding of human activities and intentions from direct neural recordings is a tantalizing challenge in brain-computer interfacing. Most ongoing efforts have focused on training decoders on specific, stereotyped tasks in laboratory settings. Implementing brain-computer interfaces in natural settings requires adaptive

strategies and scalable algorithms that work with minimal supervision. Here we propose an unsupervised approach to decoding neural states from human brain recordings acquired in a naturalistic context. We demonstrate our approach on continuous long-term electrocorticographic (ECoG) data recorded over many days from the brain surface of subjects in a hospital room, with simultaneous audio and video recordings. We first discovered clusters in the high-dimensional ECoG recordings and then annotated coherent clusters using correlations with speech and movement labels extracted automatically. To our knowledge, we are the first to leverage techniques from computer vision and speech processing for purposes of natural ECoG decoding. Our results show that this entirely unsupervised approach can discover distinct behaviors in ECoG data, including moving, speaking and resting. We verify the accuracy of our approach with comparison to manual annotation. By mapping the discovered cluster centers back onto the brain, this technique opens the door to automated functional brain mapping in natural settings.

III-102. A Bayesian approach to structured sparsity for fMRI decoding

Anqi Wu[1] ANQIW@PRINCETON.EDU
Oluwasanmi Koyejo[2] SANMI@STANFORD.EDU
Jonathan W Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Stanford University

In many problem settings, parameter vectors are not merely sparse, but dependent in such a way that non-zero coefficients tend to cluster together. We refer to this form of dependency as ‘region sparsity’. Classical sparse regression methods, such as the lasso and automatic relevance determination (ARD), model parameters as independent it a priori, and therefore do not exploit such dependencies. Here we introduce a novel, flexible method for capturing dependencies in sparse regression problems, which we call dependent relevance determination (DRD). Our approach uses a Gaussian process to model dependencies between latent variables governing the prior variance of regression weights. We combine this with a structured model of the prior variances of weights’ Fourier coefficients, which eliminates unnecessary high frequencies. The resulting model captures sparse, local structures in two different bases simultaneously, yielding estimates that are sparse as well as smooth. Our method extends previous work on automatic locality determination (ALD) and Bayesian structure learning (BSL), both of which described hierarchical models for capturing sparsity, locality, and smoothness. Unlike these methods, DRD can tractably recover region-sparse estimates with multiple regions of non-zero coefficients, without pre-definining number of regions. We also develop two efficient approximate inference methods, Laplace approximation and MCMC sampling, to deal with the intractable posterior inference for weights. In the experiment, we show substantial improvements over comparable methods (e.g., group lasso and smooth RVM) for both simulated and real datasets from brain imaging. Specifically for fMRI decoding, we can use DRD to discover multiple regional sparsities meanwhile maintaining a best regression performance against the task response.

III-103. Development and application of computational methods for high-throughput optical electrophysiology

Eli Weinstein ELI.N.WEINSTEIN@GMAIL.COM
Evangelos Kiskinis EVANGELOS.KISKINIS@NORTHWESTERN.EDU
Joel Kralj JOELKRALJ@GMAIL.COM
Peng Zou ZOUPENG85@GMAIL.COM
Kevin Eggan EGGAN@MCB.HARVARD.EDU
Adam Cohen COHEN@CHEMISTRY.HARVARD.EDU

Harvard University

A new all-optical electrophysiology technique, 'Optopatch', enables high-throughput functional characterization of individual neurons. This offers the opportunity for large scale studies of heterogenous culture populations like those derived from human induced pluripotent stem cell (iPSC) disease models, but application of Optopatch is hindered by significant computational challenges in signal extraction and interpretation. Here we develop a complete pipeline for Optopatch data analysis and apply it to an iPSC model of amyotrophic lateral sclerosis (ALS). We built image segmentation techniques for unmixing complex movies of overlapping cells and developed, via model-reduction techniques, a parameterized description of individual cell's firing patterns. We then constructed statistical metrics of both short and long time-scale spiking behavior that are robust to noise sources inherent in the Optopatch method. These techniques were applied to compare motor neurons with an ALS-causing mutation to their genome-corrected controls. The mutants showed elevated spike rates under weak or no stimulus, and greater likelihood of entering depolarization block under strong Optogenetic stimulus. We compared these results to numerical simulations of simple conductance-based neuronal models. Our data and simulations unify seemingly inconsistent literature results under a single mechanistic explanation involving deficits in potassium channels in ALS. This analysis not only provides novel insight into a specific disease but also demonstrates a framework for extracting and interpreting high-throughput optical electrophysiology data.

III-104. Active learning of psychometric functions with multinomial logistic models

Ji Hyun Bak JHBAK@PRINCETON.EDU
Jonathan W Pillow PILLOW@PRINCETON.EDU

Princeton University

As new technologies expand the capacity for making large-scale measurements of neural activity, there is a growing need for methods to rapidly characterize behavior and its dependence on stimuli. In typical experiments, an animal is presented with a stimulus on each trial and has to select a response among several options. Since experiments are costly, a problem of practical importance is to learn the animal's psychometric choice functions from a minimal amount of data. Here we show that one can achieve substantial speedups over traditional randomized designs via active learning, in which stimuli are selected adaptively on each trial according to an information-theoretic criterion. Specifically, we model behavior with a multinomial logistic regression model, in which the probability of each choice given a stimulus depends on a set of linear weights. Our work extends previous work on this problem in several important ways: (1) we incorporate an explicit lapse rate to account for the fact that observers may occasionally make errors on "easy" trials due to lapses in concentration or memory; (2) we develop an efficient method based on Markov Chain Monte Carlo (MCMC) sampling that is accurate in settings in which the log-likelihood is not concave (e.g., as in the presence of lapse rates); and (3) we extend consideration for multi-alternative responses, extending previous work for binary responses. We compare the performance of our method to one based on a local (Laplace) approximation to the posterior, and show that failure to incorporate lapse rates can have deleterious effects on the accuracy of inferred parameters under both methods. We discuss the comparative advantages and disadvantages of the different methods, and how one might adapt these algorithms to achieve best results.

III-105. Neural mass spatio-temporal modeling from high-density electrode array recordings

Robert C Sumner[1,2] RCSUMNER@UCSC.EDU
Mojtaba Sahraee-Ardakan[3,2] MOJTABASAH@GMAIL.COM
Michael Trumpis[4] MTRUMPIS@GMAIL.COM
Michele Insanally[5] MN1@NYU.EDU
Robert Froemke[5] ROBERT.FROEMKE@MED.NYU.EDU
Jonathan Viventi[4] J.VIVENTI@DUKE.EDU
Alyson K Fletcher[3,2] AKFLETCH@UCSC.EDU

[1]University of California, Santa Cruz
[2]University of California, Los Angeles
[3]University of California, Berkeley
[4]Duke University
[5]New York University

Neural mass models provide an attractive framework for modeling complex behavior in cortical circuits. Fitting these models to electrode array recordings can provide insight into connectivity and structure of neural circuits as well as the response of these circuits to stimuli. However, neural mass models are fundamentally nonlinear dynamical systems with large numbers of hidden states, and validating the models on actual recordings and estimating the key parameters remains challenging. This work presents a novel and computationally efficient method for systematically identifying a very general class of neural mass models particularly well-suited for high-density micro-electrocorticographic data. The methodology requires minimal assumptions on the model, and can automatically uncover the underlying components in the neural populations and their interactions. Importantly, the methodology can identify nonlinear dynamics as well as pre-filtering and dimensionality reduction of the stimuli that may occur prior to processing in the cortical region of interest. The procedure is validated on in vivo recordings of the rat primary auditory cortex from a recently-developed high-density 60-channel flexible electrode array with 406 um inter-electrode spacing

Author Index