COSYNE

CONFERENCE

Lisbon, Portugal

28 feb–3 mar

2019

cosyne.org

GATSBY

Google

SIMONS FOUNDATION

BURROUGHS WELLCOME FUND

Program Summary

Thursday, 28 February

4:00p **Registration opens**

4:45p **Welcome reception**

6:15p **Opening remarks**

6:30p **Session 1: Feedforward networks, deep and shallow; Keynote lecture**
Invited speaker: Yann LeCun; 3 accepted talks

8.30p **Poster Session I** *(Ball Room of the Four Seasons Hotel Ritz Lisbon)*

Friday, 01 March

9:00a **Session 2: Rewards and neuromodulators**
Invited speaker: Ilana Witten; 3 accepted talks

11:00a **Session 3: Dimensionality, high and low**
Invited speakers: Eric Shea-Brown, Kenneth Harris

12:30p **Lunch break**

3:00p **Session 4: Cortico-ception and -action I**
Invited speaker: Karel Svoboda; 3 accepted talks

5:00p **Session 5: Deliberation, exploitation, and exploration**
Invited speaker: Kenji Doya; 3 accepted talks

6:30p **Dinner break**

8:30p **Poster Session II** *(Ball Room of the Four Seasons Hotel Ritz Lisbon)*

Saturday, 02 March

9:00a **Session 6: Navigating space and time; Gatsby lecture**
Invited speaker: Evard Moser; 3 accepted talks

11:30a **Session 7: Estimating movement**
Invited speaker: Kathleen Cullen; 3 accepted talks

1:00p **Lunch break**

3:00p **Session 8: Sensory, motor, and in-between**
Invited speakers: Gwyneth Card, Sara Solla

5:00p **Session 9: Strategies of decision-making**
Invited speaker: Bruno Averbeck; 3 accepted talks

6:30p **Dinner break**

8:30p **Poster Session III** *(Ball Room of the Four Seasons Hotel Ritz Lisbon)*

Sunday, 03 March

9:00a	**Session 10: Biophysics of learning and computing**	
	Invited speaker: Yiota Poirazi; 3 accepted talks	
11:00a	**Session 11: Cortico-ception and -action II**	
	Invited speaker: Sonja Hofer; 3 accepted talks	
12:30p	**Lunch break**	
3:00p	**Session 12: Into the latent space**	
	Invited speaker: Maneesh Sahani; 3 accepted talks	

ALLEN INSTITUTE *for*
BRAIN SCIENCE

ALLEN BRAIN ATLAS RESOURCES

The Allen Institute for Brain Science provides a variety of unique online public resources that give users a powerful way to search and view extensive gene expression data, neuron projection pathways, detailed molecular characterization of single cells and neuroanatomy – all openly accessible via the Allen Brain Atlas data portal. Accelerate your research today by visiting **brain-map.org**.

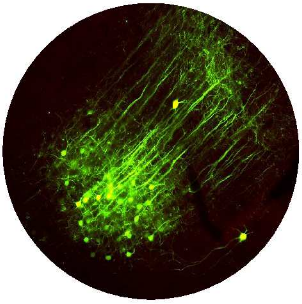

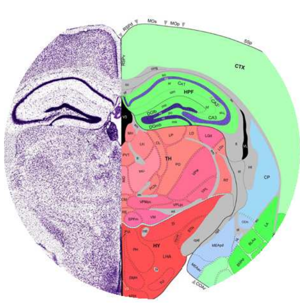

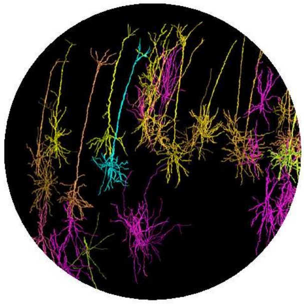

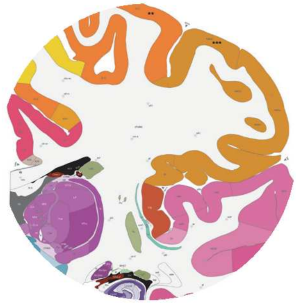

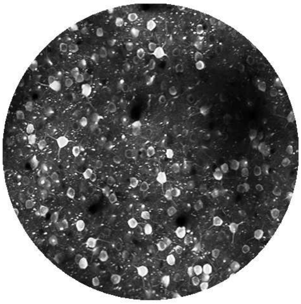

To learn more and access resources, visit

brain-map.org

If you could change **one thing** about publishing, what would it be?

flashPub


Bernstein Conference
Berlin, Sept 17-20, 2019
Confirmed speakers
Dora Angelaki (USA)
Matthias Bethge (Germany)
Matthew Botvinick (USA)
Nicolas Brunel (USA)
Claudia Clopath (UK)
Hopi Hoekstra (USA)
Gilles Laurent (Germany)
Eve Marder (USA)
Haim Sompolinsky (Israel/ USA)
Gašper Tkačik (Austria)
Nachum Ulanovsky (Israel)
Back to back with
CCN
www.ccneuro.org
Sept 13-16
Call for Workshops
Deadline: March 27, 2019
www.bernstein-conference.de

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 2019 Leadership

Organizing Committee

General Chairs

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

Program Chairs

Eugenia Chiappe (Champalimaud) and Christian Machens (Champalimaud)

Workshop Chairs

Catherine Hartley (New York University) and Ralf Haefner (University of Rochester)

Undergraduate Travel Chairs

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

Communications Chair

Xaq Pitkow (Rice University)

Publicity Chair

Il Memming Park (Stony Brook University)

Development Chair

Michael Long (New York University)

Executive Committee

Stephanie Palmer (University of Chicago)

Zachary Mainen (Champalimaud)

Alexandre Pouget (University of Geneva)

Anthony Zador (Cold Spring Harbor Laboratory)

Program Committee

Eugenia Chiappe (Champalimaud), co-chair
Christian Machens (Champalimaud), co-chair
Jennifer Bizley (Universtiy College London)
Laura Busse (Ludwig-Maximilians University Munich)
Daniel Butts (University of Maryland)
Steve Chase (Carnegie Mellon University)
Damon Clark (Yale University)
Saskia de Vries (Allen Institute for Brain Science)
Long Ding (University of Pennsylvania)
Sean Escola (Columbia University)
Rainer Friedrich (Friedrich Miescher Institute for Biomedical Research)
Julijana Gjorgjieva (Max Planck Institute for Brain Research)
Benjamin Grewe (ETH Zurich)
Richard Hahnloser (ETH Zurich)
Timothy Hanks (University of California, Davis)
Catherine Hartley (New York University)
Christopher Harvey (Harvard University)
Na Ji (University of California, Berkeley)
Andrew M Leifer (Princeton University)
Matthieu Louis (Centre for Genomic Regulation)
Jakob Macke (Caesar)
Arianna Maffei (Stony Brook University)
Cristopher Niell (University of Oregon)
Daniel O'Connor (Johns Hopkins University)
Srdjan Ostojic (Ecole Normale Superieure)
Robert Rosenbaum (University of Notre Dame)
Marshall Schuler (Johns Hopkins University)
Samuel Sober (Emory University)
Tatjana Tchumatchenko (Max Planck Institute for Brain Research)
Srini Turaga (Janelia Farm Research Campus)
Joel Zylberberg (University of Colorado)

Cosyne 2019 reviewers

Hillel Adesnik, Laurence Aitchison, Alireza Alemi, Johnatan Aljadeff, Elizabeth Amadei, Ana Amador, Todd Anthony, Ezequiel Arneodo, Hiroki Asari, Bassam Atallah, Asli Ayaz, Behtash Babadi, Andrea Barreiro, Felix Bartsch, Brice Bathellier, Aaron Batista, Daniel Bendor, Gordon Berman, Arjun Bharioke, Yazan Billeh, William Bishop, Yael Bitterman, Bart Borghuis, Hannah Bos, Yves Boubenec, Farran Briggs, Braden Brinkman, Andre E Brown, Christian Burgess, Laura Busse, Adam Calhoun, Natasha Cayco Gajic, Matthew Chalk, Rishidev Chaudhuri, Hannah Choi, SueYeon Chung, Damon Clark, Albert Compte, Christine Constantinople, Frédéric Crevecoeur, Michael Crickmore, Rajan Dasgupta, Saskia de Vries, Daniel J Denman, Eric DeWitt, Takahiro Doi, Laura Driscoll, Chunyu Duan, Eva Dyer, Alexander Ecker, Tatiana Engel, Bernhard Englitz, Jeffrey Erlich, Sean Escola, Claudia Feierstein, Gidon Felsen, Katie Ferguson, James Fitzgerald, Katrin Franke, Kevin Franks, Rainer Friedrich, Jonathan Fritz, Terufumi Fujiwara, Aleena Garner, Marcelo Matheus Gauy, Julijana Gjorgjieva, Michael Goard, Matthew D Golub, Pedro J Goncalves, Andrew Gordus, Nathan Gouwens, Michael Graupner, David Greenberg, Benjamin Grewe, Zengcai Guo, Gabrielle Gutierrez, Richard Hahnloser, Bilal Haider, Timothy Hanks, Christopher Harvey, Guillaume Hennequin, Matthias Hennig, Christian Henning, James Herman, Chengcheng Huang, Daniel Huber, Laurence Hunt, Alexandre Hyafil, Ramakrishnan Iyer, James Jeanne, Xiaoxuan Jia, Maximilian Jösch, Jonathan Kadmon, Naama Kadmon Harpaz, Matthias Kaschube, Leor Katz, Matthew Kaufman, Hokto Kazama, Peter Keating, Sander Keemink, Georg Keller, Simon Kelly, Ann Kennedy, Adil Khan, Zachary Kilpatrick, Shinichiro Kira, Dmitry Kobak, Sepp Kollmorgen, Charles Kopec, Jens Kremkow, Alex Kwan, Sung Eun Kwon, Josef Ladenbauer, Guillaume Lajoie, Stephan Lammel, Crystal Lantz, Arthur Leblois, Peter Ledochowitsch, Wei-Chung Lee, Nuo Li, Scott Linderman, Ashok Litwin-Kumar, Haixin Liu, Laureline Logiaco, Matthieu Louis, Christian Machens, Jakob Macke, Emily Mackevicius, Arianna Maffei, Valerio Mante, Jonathan Mapelli, David Markowitz, Jesse Marshall, Francesca Mastrogiuseppe, Mackenzie Mathis, Luca Mazzucato, Matthew McGinley, Xiangying Meng, Natasha Mhatre, Daniel Millman, Wiktor Młynarski, Gabriela Mochol, Alice Mosberger, James Murray, John D Murray, Katherine Nagel, Kae Nakamura, Vijay Mohan K Namboodiri, Reviewer Name, Devika Narain, Matthew Nassar, Richard Naud, Eva Naumann, Cristopher Niell, Thierry Nieus, Daniel O'Connor, Gabriel Ocker, Srdjan Ostojic, Johannes Oswald, Marius Pachitariu, Chethan Pandarinath, Simon Peron, Catherine Perrodin, Alon Poleg-Polsky, Viola Priesemann, Andrew Pruszynski, Pavan Ramdya, Samuel Reiter, Alfonso Renart, Blake Richards, Mattia Rigotti, Nathalie Rochefort, Robert Rosenbaum, Douglas Ruff, Simon Rumpel, Eleonora Russo, Houman Safaai, Laura Sainz Vilalba, Aman Saleem, Yulia Sandamirskaya, Takashi Sato, Cristina Savin, Evan Schaffer, David Schneider, Geoffrey Schoenbaum, Monika Scholz, Annalisa Scimemi, Shan Shen, Marshall Hussain Shuler, Marion Silies, Julie Simpson, Annabelle Singer, Fabian Sinz, Matt Smear, Spencer Smith, Adam Snyder, Henning Sprekeler, Natalie Steinemann, Martin Stemmler, Carsen Stringer, Simone Carlo Surace, Lalitta Suriya-Arunroj, Tatjana Tchumatchenko, Corinne Teeter, Sina Tootoonian, Philip Tovote, Srinivas C Turaga, John Tuthill, Vivek Venkatachalam, Roberto Vincis, Kerry Walker, Lupeng Wang, Yingxue Wang, Adrian Wanner, Kate Wassum, Mike Wehr, Felix Weissenberger, Ben Willmore, Klaus Wimmer, Marina Elaine Wosniack, Qiaojie Xiong, Emre Yaksi, Guangyu Robert Yang, Tianming Yang, Friedemann Zenke, and Joel Zylberberg

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

Conference Support

Administrative Support, Registration, Hotels

Leslie Weekes, Cosyne

Social media policy

Cosyne encourages the use of social media before, during, and after the conference, so long as it falls within the following rules:

- Do not capture or share details of any unpublished data presented at the meeting.
- If you are unsure whether data is unpublished, check with the presenter before sharing the information.
- Respect presenters' wishes if they indicate the information presented is not to be shared.

Stay up to date with Cosyne 2019 #cosyne19

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 2019.

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

- The Gatsby Charitable Foundation
- Burroughs Wellcome Fund
- Simons Foundation
- Google

Cosyne Presenters Travel Grant Program

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

The 2019 recipients are

William Allen, John Berkowitz, Georgia Christodoulou, Alexis Dubreuil, Becket Ebitz, Jaclyn Essig, Katharina Glomb, HyungGoo Kim, Kathleen Martin, Michael Morais, Marin Ozaki, Dylan Paiton, Elizabeth Sachse, Nitzan Shahar, Jianghong Shi, Ethan Solomon, Sigrid Trägenap, Anqi Wu, and Ziniu Wu.

Cosyne New Attendees Travel Grant Program

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

The 2019 recipients are

Amanda Buch, Michael Caiola, Cheng Chen, Anna Ciaunica, Chase Clark, Aida Davila, Céline Drieu, Samuel Eckmann, Victor Ekuta, Jordan Elum, Soledad Gonzalo Cogno, Harsha Gurnani, Nora Alicia Herweg, Nikolaos Karalis, Vickie Li, Katie Long, James Marshel, Jorge Aurelio Menendez, Sima Mofakham, Josue Nassar, Luis Giordano Ramos Traslosheros Lopez, Kathryn Rothenhoefer, Karen Schroeder, Varsha Sreenivasan, Julia Steinberg, Marcus Triplett, and Chi Zhang

Cosyne Mentorship Travel Grant Program

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

The 2019 Cosyne Mentors and mentees are:

Nicola Grissom and Sijin Chen, Nils Kolling and Jan Grohn, Kishore Kuchibhotla and Sarah Elnozahy, John Murray and Jasmine Stone, Stephanie Palmer and Jennifer Ding, and Cengiz Pehlevan and Xinyuan Zhao.

Cosyne Undergraduate Travel Grant Program

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

The 2019 recipients are

Sharon Chen, Philip Cho, Jennifer DiSanto, Madison Hecht, Samo Hromadka, Fatima Irfan, Grace Irungu, Francisco Sacadura, Mengyu Tu, and Abby Wood.

Code of Conduct at Cosyne

Cosyne is adopting the practice of defining a meeting code of conduct and providing information for participants on our ethics policies as well as resources for filing complaints. This code of conduct is based on standards and language set at other meetings, whose organizing boards convened special working groups of scientific and legal experts to set their policies. We follow, in particular, those guidelines established for the Gordon Research Conferences and the Society for Neuroscience Annual Meeting.

The following code of conduct has been adapted from:

https://www.grc.org/about/grc-policies-and-legal-disclaimers

https://www.sfn.org/Membership/Professional-Conduct/Code-of-Conduct-at-SfN-Events

other online resources:

http://changingourcampus.org

https://www.sfn.org/Membership/Professional-Conduct/SfN-Ethics-Policy

At Cosyne, we strive for open and honest intellectual debate as part of a welcoming and inclusive atmosphere. Cosyne asks each session chair and organizing and reviewing committee member to promote rigorous analysis of all science presented for or at the meeting in a manner respectful to all attendees. To help maintain an open and respectful community of scientists, Cosyne does not tolerate illegal or inappropriate behavior at the main meeting or the workshops, including violations of applicable laws pertaining to sale or consumption of alcohol, destruction of property, or harassment of any kind, including sexual harassment.

Cosyne condemns inappropriate or suggestive acts or comments that demean another person by reason of their gender, gender identity or expression, race, religion, ethnicity, age or disability or that are unwelcome or offensive to other members of the community or their guests. Cosyne will review allegations of any such behavior on a case-by-case basis, and violations may result in the prohibition on future attendance by particular individuals.

Cosyne believes home institutions and employers of our event attendees, including research institutions, companies, and other organizations, are best equipped to investigate allegations or violations by their faculty members, trainees, employees or affiliated individuals, and to evaluate and determine appropriate actions consistent with their employment and academic obligations. In the event of an allegation of harassment at a Cosyne-organized event, Cosyne will document and will generally refer allegations to an attendee's organization for its review. To the extent possible, documentation will be held in confidence by Cosyne's executive staff.

If you believe you have been subjected to or have otherwise experienced behavior at a Cosyne event that violates Cosyne's Code of Conduct, please act promptly to report the issue so that steps may be taken to address the situation immediately. You may notify the chair of the session you are attending, or one of the meeting organizers, each of whom we have asked to address the matter promptly or refer it back to the Cosyne main office. If you are reluctant to speak with an organizer or session chair for any reason, you may notify Leslie Weekes, Cosyne's Executive Administrator, at the Cosyne main office by email to LESLIE.WEEKES@COSYNE.ORG.

Reports requesting anonymity will be respected, although Cosyne reserves the right to notify appropriate law enforcement if the allegations are deemed serious enough to warrant such notice. It should be noted that Cosyne's ability to investigate or address anonymous reports may be limited or otherwise affected by the need to balance concerns over privacy and fairness to all concerned. Following completion of Cosyne's inquiry, any action to be taken by Cosyne against the person accused of acting inappropriately will be determined by Cosyne in its sole discretion and may include discharge from the conference or restrictions on their future attendance.

Program

Note:

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

Thursday, 28 February

8:00a **Allen Institute Tutorial session**

12:00n **Cosyne 2019 Tutorial session sponsored by the Simons Foundation**

4:00p **Registration opens**

4:45p **Welcome reception**

6:15p **Opening remarks**

Session 1: Feedforward networks, deep and shallow; Keynote lecture

(Chair: Eugenia Chiappe, Christian Machens)

Friday, 01 March

Session 2: Rewards and neuromodulators

(Chair: Joseph Paton)

10:15a	Serotonin neurons track reward history and modulate flexible behavior. C. Grossman, B. Bari, J. Cohen, Johns Hopkins School of Medicine	35
10:30a	**Coffee break**	

Session 3: Dimensionality, high and low

(Chair: Srdjan Ostojic)

11:00a	What makes high-dimensional networks produce low-dimensional activity? Eric Shea-Brown, University of Washington **(invited)**	27
11:45a	$N_{neurons} \to \infty$ Kenneth Harris, University College London **(invited)**	28
12:30p	**Lunch break**	
2:00p–3:00p	***NIH Funding and Training Event (optional).*** **Please join NIH representatives to learn about current funding and training opportunities in computational neuroscience offered by the NIH and the BRAIN Initiative.**	

Session 4: Cortico-ception and -action I

(Chair: Leopoldo Petreanu)

3:00p	The multi-regional neural circuits underlying elementary cognitive behaviors Karel Svoboda, HHMI Janelia Research Campus **(invited)**	28
3:45p	Inhibition stabilization is a widespread feature of mouse cortical networks A. Sanzeni, B. Akitake, N. Brunel, M. Histed, NIMH/NIH	38
4:00p	Low threshold layer-specific neural ensembles contributing to the ignition of perception J. Marshel, T. Machado, S. Quirin, Y. Kim, B. Benson, J. Kadmon, C. Raja, A. Chibukhchyan, C. Ramakrishnan, M. Inoue, H. Kato, S. Ganguli, K. Deisseroth, Stanford University	39
4:15p	A premotor circuit responsible for flexible decision making Z. Wu, A. Litwin-Kumar, P. Shamash, A. Taylor, R. Axel, M. Shadlen, Columbia University	39
4:30p	**Coffee break**	

Session 5: Deliberation, exploitation, and exploration

(Chair: Zachary Mainen)

5:00p	Neural circuit for mental simulation Kenji Doya, OIST **(invited)**	28
5:45p	Anterior Cingulate Cortex Directs Exploration of Alternative Strategies D. Tervo, A. Karpova, E. Kuleshova, M. Manakov, M. Proskurin, M. Karlsson, HHMI Janelia Research Campus	36
6:00p	Amygdala ensembles encode behavioral states. J. Grundemann, Y. Bitterman, B. Grewe, T. Lu, S. Krabbe, M. J. Schnitzer, A. Lthi, University of Basel	37
6:15p	Exploration via disrupted sensorimotor control dynamics B. Ebitz, J. Cohen, T. Buschman, T. Moore, B. Hayden, University of Minnesota	37
6:30p	**Dinner break**	
8:30p	**Poster Session II** *(Ball Room of the Four Seasons Hotel Ritz Lisbon)*	

Saturday, 02 March

Session 6: Navigating space and time; Gatsby lecture

(Chair: Stephanie Palmer)

9:00a Neural coding of space and time in the entorhinal cortex **(Gatsby lecture)**
Edvard Moser, NTNU **(invited)** 29

10:15a Multiple maps of the same context stably coexist in the mouse hippocampus
A. Rubin, L. Sheintuch, N. Geva, H. Baumer, Y. Ziv, Weizmann Institute of Science 40

10:30a Representation of 3D space in the entorhinal cortex of flying bats
G. Ginosar, J. Aljadeff, Y. Burak, H. Sompolinsky, L. Las, N. Ulanovsky, Weizmann Institute of Science 40

10:45a Detailing the network structure and downstream function of a Drosophila ring attractor
D. Turner-Evans, S. Ali, A. Sheridan, T. Paterson, K. Jensen, R. Ray, S. Lauritzen, D. Bock, V. Jayaraman, HHMI Janelia Research Campus 41

11:00a **Coffee break**

Session 7: Estimating movement

(Chair: Saskia de Vries)

11:30a Neural representations of natural self-motion: Implications for perception & action
Kathleen Cullen, Johns Hopkins University **(invited)** 29

12:15p Population activity of cerebellar granule cells in awake mice
N. Cayco Gajic, R. A. Silver, F. Lanore, University College London 41

12:30p Head direction and orienting-suppression signals in primary visual cortex of freely-moving rats.
G. Guitchounts, J. Masis, S. Wolff, D. Cox, Harvard University 42

12:45p Predicting natural behavior from brain-wide neural activity
M. Scholz, A. Linder, F. Randi, A. Sharma, X. Yu, J. Shaevitz, A. Leifer, Princeton University 43

1:00p **Lunch break**

1:00p **Lunch—Workshop for equality and diversity in science**

Session 8: Sensory, motor, and in-between

(Chair: Andrew M Leifer)

3:00p Towards a brain architecture for visual behavior selection
Gwyneth Card, Janelia Research Campus **(invited)** 30

3:45p Stable manifold dynamics underlie the consistent execution of learned behavior
Sara Solla, Northwestern University **(invited)** 30

4:30p **Coffee break**

Session 9: Strategies of decision-making

(Chair: Richard Hahnloser)

5:00p	Neural systems underlying reinforcement learning Bruno Averbeck, NIMH **(invited)**	31
5:45p	Dissociating task acquisition from expression during learning reveals latent knowledge T. Hindmarsh Sten, E. Papadoyannis, K. Fogelson, S. Elnozahy, C. Drieu, R. K. Chillale, Y. Boubenec, R. Froemke, P. Holland, S. Ostojic, K. Kuchibhotla, Rockefeller University	43
6:00p	A value-based explanation for lapses in perceptual decisions S. Pisupati, L. Chartarifsky-Lynn, A. Khanal, A. Churchland, Cold Spring Harbor Laboratory	44
6:15p	Striatum encodes task adaptative expectations and mediates their impact on perceptual decisions A. Hermoso Mendizabal, A. Hyafil, P. E. Rueda Orozco, D. Robbe, J. de la Rocha, Idibaps	44
6:30p	**Dinner break**	
8:30p	**Poster Session III** *(Ball Room of the Four Seasons Hotel Ritz Lisbon)*	

Sunday, 03 March

Session 10: Biophysics of learning and computing

(Chair: Daniel Butts)

9:00a	Active dendrites and their role in neuronal and circuit computations Yiota Poirazi, IMBB-FORTH **(invited)**	31
9:45a	Measuring and modeling the weight dynamics of many synapses onto diverse cell-types in vivo A. Nayebi, J. Melander, B. Jongbloets, T. Mao, H. Zhong, S. Ganguli, Stanford University	45
10:00a	Is synaptic potentiation required for associative learning in adult fruit fly? R. Lonsdale, P. F. Jacob, B. Menzat, S. Waddell, T. P. Vogels, University of Oxford	46
10:15a	Flexible mapping of visual scenes to a heading representation via inhibitory Hebbian plasticity S. S. Kim, A. Hermundstad, L. Abbott, V. Jayaraman, HHMI Janelia Research Campus	46
10:30a	**Coffee break**	

Session 11: Cortico-ception and -action II

(Chair: Cristopher Niell)

11:00a	Making sense of what you see: Cortical and thalamic circuits for vision Sonja Hofer, Sainsbury Wellcome Centre, UCL **(invited)**	31
11:45a	Orientation preference maps in the mouse lemur confirm common design principles in primate visual cortex C. L. A. Ho, R. Zimmermann, J. Weidinger, B. Tia, F. Pifferi, F. Aujard, F. Wolf, D. Huber, University of Geneva	47
12:00n	Sustained and dynamic representations reduce interference in short-term memory A. Libby, T. Buschman, Princeton University	48
12:15p	Laminar-specific cortico-cortical loops in mouse visual cortex H. Young, B. Belbut, L. Petreanu, Champalimaud Foundation	48
12:30p	**Lunch break**	

Session 12: Into the latent space

(Chair: Sean Escola)

Poster Session I 8:30p Thursday 28 February

Poster Session II — 8:30p Friday 01 March

Poster Session III 8:30p Saturday 02 March

Cosyne Poster Sessions Walking Map

Directions from EPIC SANA Lisboa Hotel to the Four Seasons/ Ritz. Please use R. Castilho entrance.

A

EPIC SANA Lisboa Hotel, Avenida Engenheiro Duarte Pacheco, Lisbon, Portugal

B

Four Seasons/Ritz, R. Castilho 77C, 1070-051 Lisboa, Portugal

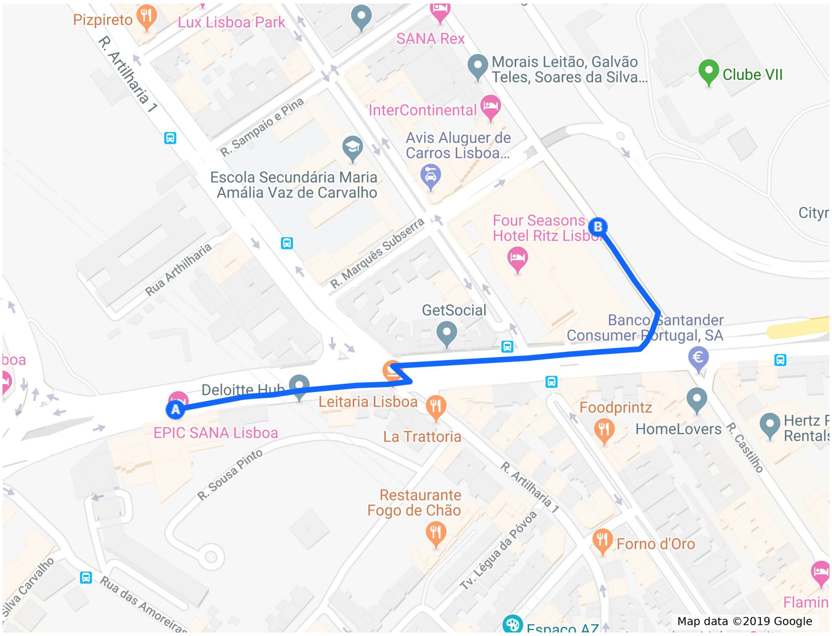

Abstracts

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. Cosyne 2019 keynote lecture

Yann LeCun
New York University

T-2. Functional organization of the dopamine system

Ilana Witten
Princeton University

The classic view of the striatal circuit in learning and decision making is that corticostriatal inputs encode specific actions or stimuli, and a homogeneous reward prediction error provided by dopamine neurons serves to modify the strength of those corticostriatal synapses, altering the behaviors which are most likely to subsequently occur. However, due to technical limitations, it has been difficult to test this idea rigorously. To address this gap, my lab has been using circuit dissection tools to record and manipulate activity in genetically and anatomically defined inputs to the striatum. For example, by comparing neural coding in anatomically-defined dopamine subpopulations, we discovered that dopamine neurons convey specialized and spatially organized information about movements, choices, and other behavioral variables to specific striatal subregions, in addition to encoding reward prediction error. These findings revise the classic view that dopamine neurons convey a spatially uniform reward prediction error signal to the striatum, and raise important questions regarding potential functions of non-reward signals in the dopamine system.

T-3. What makes high-dimensional networks produce low-dimensional activity?

Eric Shea-Brown
University of Washington

There is an avalanche of new data on the brain's activity, revealing the collective dynamics of vast numbers of neurons. In principle, these collective dynamics can be of almost arbitrarily high dimension, with many independent degrees of freedom — and this may reflect powerful capacities for general computing or information. In practice, neural datasets reveal a range of outcomes, including collective dynamics of much lower dimension — and this may reflect other desiderata for neural codes. For what networks does each case occur? Our contribution to the answer is a new framework that links tractable statistical properties of network connectivity with the dimension of the activity that they produce. In tandem, we study how features of connectivity and dynamics that impact dimension arise as networks learn to perform basic tasks. I'll describe where we have succeeded, where we have failed, and the many avenues that remain.

T-4. $N_{neurons} \to \infty$

Kenneth Harris

University College London

Simultaneous recordings from tens of thousands of neurons allow a new framework for characterizing the neural code at large scales. As the number of neurons analyzed increases, population activity approximates a vector in an infinite-dimensional Hilbert space. In this limit, the independent activity of any single neuron is of no consequence, and the neural code reflects only activity dimensions shared across the population. Analyzing the responses of large populations in mouse visual cortex to natural image stimuli revealed an unexpected result: signal variance in the n^{th} dimension decayed as a power law, with exponent just above 1. We proved mathematically that a differentiable representation of a d-dimensional stimulus requires variances decaying faster than $n^{-1-2/d}$. By recording neural responses to stimulus ensembles of varying dimension, we showed this bound is close to saturated. We conclude that the cortical representation of image stimuli is as high-dimensional as possible before becoming non-differentiable.

T-5. The multi-regional neural circuits underlying elementary cognitive behaviors

Karel Svoboda

HHMI Janelia Research Campus

Our goal is to uncover the principles by which neural circuits perform fundamental computations, including perception, decision-making, movement planning and motor control. All behaviors involve multi-regional neural circuits. Neural information is represented by action potentials in widely distributed ensembles of neurons. We are interested in the mechanisms shaping these neural representations at the level of the entire brain, and how these representations drive behavior. I will discuss our attempts to understand the diverse ways in which connected brain regions interact in behaving mice.

T-6. Neural circuit for mental simulation

Kenji Doya

OIST

The basic process of decision making can be captured by learning of action values according to the theory of reinforcement learning. In our daily life, however, we rarely rely on pure trial-and-error and utilize any prior knowledge about the world to imagine what situation will happen before taking an action. How such "mental simulation" is realized in the circuit of the brain is an exciting new topic of neuroscience. Here I report our works with functional MRI in humans and two-photon imaging in mice to clarify how action-dependent state transition models are learned and utilized in the brain.

T-7. Neural coding of space and time in the entorhinal cortex

Edvard Moser

Kavli Institute for Systems Neuroscience and Centre for Neural Computation (NTNU)

The medial entorhinal cortex (MEC) is part of a brain system for mapping of self-location during navigation in the proximal environment. In this system, neural firing correlates are so evident that cell types have been given simple descriptive names, such as grid cells, border cells, and head direction cells. Insights remain limited, however, by an almost exclusive reliance on recordings from rodents foraging in empty enclosures quite different from the richly populated, geometrically irregular environments of the native world. In the first part of my talk, I will show that in environments filled with discrete objects, a large number of MEC neurons, comparable to the number of grid cells, fire specifically at given distances and directions from the objects. These 'object-vector cells' are tuned equally to a spectrum of objects, irrespective of their location, as well as to a broad range of object dimensions and shapes, from point-like to extended. In the second part of the talk, I will show that this network of entorhinal cell types maintains its temporal correlation structure across environments, and that in grid cells, phase relationships are preserved from awake exploration to sleep, suggesting they are part of an interconnected, low-dimensional network with locational firing correlates. I will report work in which we address the dynamics of this network by calcium imaging from many dozens of MEC neurons while head-fixed mice run in darkness, with very limited external input to guide the formation of a coherent map of experience. Under these conditions, neural population activity is organized in stereotyped motifs, consisting of sequences that span time scales of tens of seconds or even minutes. Such motifs appear progressively during development, suggesting that the MEC network is organized into a rigid topology of subnetworks that emerges progressively during the first month of life of a mouse. Finally, in the last part of the talk, I will show how episodic temporal information is encoded across scales from seconds to hours within the overall population state of the lateral entorhinal cortex (LEC). This representation of time is dynamic and reflects the structure of the task, suggesting that populations of LEC neurons represent time inherently through the encoding of experience. In the hippocampus, this task-dependent representation of time may be integrated with spatial inputs from grid cells and object-vector cells in MEC, allowing it to store a unified representation of experience.

T-8. Neural representations of natural self-motion: Implications for perception & action

Kathleen Cullen

Johns Hopkins University

A fundamental question in neuroscience is how does the brain compute accurate estimates of our self-motion relative to the world and orientation relative to gravity in everyday life. In this talk, I will describe recent findings from my laboratory's research that have addressed this question and provided new insight into how vestibular pathways encode self-motion information to ensure accurate perception and motor control. First, we have recently examined the statistics of natural self-motion signals experienced by mice, monkeys, and humans, and then explored the neural coding strategies used by early vestibular pathways. Focusing on the relationships between neural variability, detection thresholds, and information transmission, our findings have revealed that two distinct sensory channels represent vestibular information at the level of the vestibular periphery. Notably, more regularly discharging afferents have better detection thresholds and use rate coding, while more irregular afferents take advantage of precise spike timing (i.e, temporal coding) and are better optimized for processing natural vestibular stimuli. Our research has also established that the neurons at the first central stage of vestibular processing are substantially less sensitive to active motion. Notably, this ability to distinguish between active and passive motion is not a general feature of early vestibular processing, but is instead a characteristic of a distinct group of neurons known to contribute to postural control and spatial orientation. Our most recent studies further indicate that multimodal integration within the vestibular cerebellum is required for this cancellation of self-generated vestibular information from the subsequent computation of orientation and posture control. Moreover, when unexpected vestibular inputs become persistent during active motion, this mechanism is rapidly updated to re-enable the vital distinction between active and passive motion to ensure the maintenance of posture and stable perception.

T-9. Towards a brain architecture for visual behavior selection

Gwyneth Card

HHMI Janelia Research Campus

Selecting the right behavior at the right time is critical for animal survival. Animals rely on their senses to deliver information about the environment to sensory processing areas in the brain that extract relevant features and form the perceptual representations that guide behavior. We aim to uncover the organization of this feature space and the neural mechanisms by which these cues are translated into dynamic motor activity. Our current focus is visually-driven behaviors of the fly. In particular, those driven by visual looming cues produced by an approaching predator or an imminent collision. The same looming stimulus can evoke a wide range of different behaviors, including a rapid escape jump, a slower, more stable takeoff sequence, or a landing response. We use whole-cell patch clamp physiology in behaving flies, calcium imaging, high-throughput/high-resolution behavioral assays, and genetic tools to examine the transformation of information from sensory to motor. I will discuss our recent work investigating the representation of ethologically-relevant visual features in the fly optic glomeruli and the mechanisms by which descending neurons read out this feature information to produce an appropriate behavioral choice.

T-10. Stable manifold dynamics underlie the consistent execution of learned behavior

Sara Solla

Northwestern University

For learned actions to be executed reliably, the cortex must integrate sensory information, establish a motor plan, and generate appropriate motor outputs to muscles. Animals, including humans, perform such behaviors with remarkable consistency for years after acquiring a skill. How does the brain achieve this stability? Is the process of integration and planning as stable as the behavior itself? We explore these fundamental questions from the perspective of neural populations. Recent work suggests that the building blocks of neural function may be the activation of population-wide activity patterns, the neural modes, rather than the independent modulation of individual neurons. These neural modes, the dominant co-variation patterns of population activity, define a low dimensional neural manifold that captures most of the variance in the recorded neural activity. We refer to the time-dependent activation of the neural modes as their latent dynamics. We hypothesize that the ability to perform a given behavior in a consistent manner requires that the latent dynamics underlying the behavior also be stable. A dynamic alignment method allows us to examine the long term stability of the latent dynamics despite unavoidable changes in the set of neurons recorded via chronically implanted microelectrode arrays. We use the sensorimotor system as a model of cortical processing, and find remarkably stable latent dynamics for up to two years across three distinct cortical regions, despite ongoing turnover of the recorded neurons. The stable latent dynamics, once identified, allows for the prediction of various behavioral features via mapping models whose parameters remain fixed throughout these long timespans. These results are upheld by an adversarial domain adaptation approach that aligns latent spaces based on data statistics rather than dynamics. We conclude that latent cortical dynamics within the task manifold are the fundamental and stable building blocks underlying consistent behavioral execution.

T-11. Neural systems underlying reinforcement learning

Bruno Averbeck
NIMH

I will discuss recent work on the neural circuitry underlying model-free and model-based reinforcement learning (RL). While there has been considerable focus on dopamine and its action in the striatum, particularly for model-free RL, our recent work has shown that the amygdala also plays an important role in these processes. We have further found that the amygdala and striatum learn in parallel. However, the amygdala learns more rapidly than the striatum. Therefore, each structure tends to be optimized for different reward environments. We also have studied prefrontal cortex's role in Bayesian, model based RL within the context of reversal learning paradigms. Simultaneous recordings from large populations of neurons in dorsal-lateral prefrontal cortex show that, when an animal detects a reversal in choice-outcome contingencies, prefrontal cortex responds with a phasic increase in activity, coincident with the animal's preference switch. Overall, the work in our lab outlines roles for multiple neural circuits spanning cortical-basal ganglia-thalamocortical loops, as well as the amygdala's interaction with these circuits, in RL.

T-12. Active dendrites and their role in neuronal and circuit computations

Yiota Poirazi
IMBB-FORTH

The goal of this presentation is to provide a set of predictions generated by biophysical and/or abstract mathematical models regarding the role of dendrites in information processing, learning and memory across different brain regions. I will present modelling studies from our lab –along with supporting experimental evidence- that investigate how dendrites may be used to facilitate the learning and coding of both spatial and temporal information at the single cell, the microcircuit and the neuronal network level. I will present the main findings of a number of projects in lab dealing with dendritic nonlinearities in PV interneurons and their consequences on memory encoding [1], the role of dendrites in pattern separation in the DG [2] and microcircuit contributions to place cell dynamics in the CA1 [3,4]. [1] Tzilivaki A., Poirazi P. Challenging the point neuron dogma: FS basket cells as 2-stage nonlinear integrators (in revision) [2] Chavlis, S., Petrantonakis, P., and Poirazi, P. "Dendrites of dentate gyrus granule cells contribute to pattern separation by controlling sparsity." Hippocampus. 2017 Jan;27(1):89-110. doi: 10.1002/hipo.22675. [3] Turi, G., Li, W-K., Chavlis, S., Pandi, I., O'Hare, J., Priestley, J.B., Grosmark, AD, Liao, Z., Ladow, M., Zhang, JF., Zemelman, BV., Poirazi, P. and Losonczy, A. "Vasoactive intestinal polypeptide-expressing interneurons in the hippocampus support goal-oriented spatial learning", in Press, Neuron, 2019. [4] Shuman, T., Aharoni, D. Cai, DJ, Lee, CR, Chavlis, S., Taxidis, J., Flores, S., Cheng, K., Javaherian, M, Kaba, CC, Shtrahman, M, Kakhurin, KI, Masmanidis, S., Khakh, BS, Poirazi, P., Silva, AJ, Golshani, P. "Breakdown of spatial coding and neural synchronization in epilepsy" (in revision)

T-13. Making sense of what you see: Cortical and thalamic circuits for vision

Sonja Hofer
Sainsbury Wellcome Centre, UCL

The classical model of sensory information processing is based on a hierarchical organization of feed-forward connections from one brain region to the next. However, perception is not only dependent on the sensory feed-forward input but also on the context in which a given stimulus occurs, such as an animal's behavioural state, its knowledge, expectations and actions. Such contextual, top-down information can strongly modulate sensory responses and influence how sensory information is interpreted and perceived. My lab studies the circuits supporting sensory information processing, how different signals are integrated by these circuits and the mechanisms by which context and behavioural relevance influence visual processing and visually-guided behaviour. My talk will cover our most recent findings on the principles of circuit organization that underlie visual processing in the neocortex as well as how visual cortex responses and interactions between excitatory and inhibitory cell types change during learning as sensory stimuli acquire behavioural relevance. Further, while sensory perception is thought to mainly rely on cortical circuits, higher-order sensory nuclei in the thalamus interconnect extensively with all sensory cortical and many subcortical areas. In the second part of my talk, I will present our efforts to understand the role of higher-order thalamocortical interactions during sensory processing, with a focus on what information these pathways convey to different cortical areas and whether they link sensory with internal, contextual signals.

T-14. Learning to compute with distributed distributional representations of uncertainty

Maneesh Sahani

Gatsby Computational Neuroscience Unit, UCL

To guide effective decisions and plans, the brain must combine (as best as possible) ambiguous and noisy sensory signals with imperfect knowledge of the structure of the environment and of the consequences of action. Optimal behaviour in many such noisy, ambiguous and uncertain settings can be formulated as (Bayesian) probabilistic computation, and indeed behavioural experiments have shown that humans and other animals often approach Bayes-defined optimality. How does the brain achieve this?

Numerous theories have been put forward for how neural activity might represent uncertainty, and how neural circuits might compute with these representations. But little of this work has asked how such representations and computations might be learnt using local learning rules. In this talk I will argue that one particular representation, which we call a distributed distributional code or DDC, is particularly well-suited to distributed computation and local learning. I will demonstrate various inferential computations that can be efficiently implemented and learnt with DDCs in a neurally plausible way, and show instances where such DDC-based computation and learning out-performs state-of-the-art variational methods.

T-15. Lagrangian dynamics of dendritic microcircuits enables real-time backpropagation of errors

Dominik Dold[1] DOMINIK.DOLD@KIP.UNI-HEIDELBERG.DE
Akos F Kungl[1] FKUNGL@KIP.UNI-HEIDELBERG.DE
Joao Sacramento[2] SACRAMENTO@INI.ETHZ.CH
Mihai A Petrovici[3] PETROVICI@PYL.UNIBE.CH
Kaspar Schindler[4] KASPAR.SCHINDLER@INSEL.CH
Jonathan Binas[5] JONATHAN.BINAS@UMONTREAL.CA
Yoshua Bengio[5] BENGIOY@IRO.UMONTREAL.CA
Walter Senn[3] SENN@PYL.UNIBE.CH

[1]Heidelberg University
[2]ETH Zurich
[3]University of Bern
[4]Inselspital Bern
[5]Universite de Montreal

A major driving force behind the recent achievements of deep learning is the backpropagation-of-errors algorithm (backprop), which solves the credit assignment problem for deep neural networks. Its effectiveness in abstract neural networks notwithstanding, it remains unclear whether backprop represents a viable implementation of cortical plasticity. Here, we present a new theoretical framework that uses a least-action principle to derive a biologically plausible implementation of backprop. In our model, neuronal dynamics are derived as Euler-Lagrange equations of a scalar function (the Lagrangian). The resulting dynamics can be interpreted as those of multi-compartment neurons with apical and basal dendrites, coupled with a Hodgkin-Huxley-like activation mechanism that undoes temporal delays introduced by finite membrane time constants. We suggest that a neuron's apical potential encodes a local prediction error arising from the difference between top-down feedback from higher cortical areas and the bottom-up prediction represented by activity in its home layer. This computation is enabled by a stereotypical cortical microcircuit, projecting from pyramidal neurons to interneurons back to the pyramidal neurons' apical compartments. When a subset of output neurons is slightly nudged towards a target behavior that cannot be explained away by bottom-up predictions, an error signal is induced that propagates back throughout the network through feedback connections. By defining synaptic dynamics as gradient descent on the Lagrangian, we obtain a biologically plausible plasticity rule that acts on the forward projections of pyramidal and interneurons in order to reduce this error. The presented model incorporates several features of biological neurons that cooperate towards approximating a time-continuous version of backprop, where plasticity acts at all times to reduce an output error induced by mismatch between different information streams in the network. The model is not only restricted to supervised learning, but can also be applied to unsupervised and reinforcement learning schemes, as demonstrated in simulations. This work has received funding from the European Union under grant agreements 720270, 785907 (HBP) and the Manfred Stark Foundation. Calculations were performed on UBELIX (University of Bern HPC) and bwHPC (state BaWu HPC, supported by the DFG through grant no INST 39/963-1 FUGG).

T-16. Learning in a multi-layer network of an electric fish

Salomon Z Muller SZM2106@COLUMBIA.EDU
Larry F Abbott LFA2103@COLUMBIA.EDU
Nathaniel B Sawtell NS2635@COLUMBIA.EDU

Columbia University

Learning that is distributed across multiple neurons and synapses in multi-layer networks is challenging both to implement and to understand. Combining experimental data and modeling, we have developed a detailed network-level description of learning in the electrosensory lobe (ELL) of mormyrid electric fish. These fish detect weak electric fields produced by their prey despite producing much stronger electric fields themselves (Bell 1981). Sensory responses to predictable self-generated fields are removed by the ELL circuitry to reveal weaker prey-generated fields (Enikolopov et al. 2018). This happens because the ELL creates a prediction of the responses to the self-generated field in the form of a negative image that is subtracted from the total sensory response. Past studies highlighted the role of anti-Hebbian plasticity in producing negative images (Bell et al. 1997; Kennedy et al. 2014), but left large parts of the ELL circuitry unexplained. We present a full circuit-level account of the ELL that, in addition to explaining ELL function, introduces a number of general principles. First, negative images are computed by a two-stage network. Second, neurons at the first stage of this network are two compartment, with one compartment controlling learning and the other computing the transformation from synaptic input to spiking output. Third, the circuit is organized anatomically on the basis of how the neurons learn about sensory input, not how they respond to it, a feature that has also been suggested regarding the cerebellum (Herzfeld et al. 2015) and may have broad implications for other brain regions as well, including hippocampus and cortex.

T-17. Optimal cortical plasticity in a model of perceptual learning

Haozhe Shan[1] SHANHAOZHE@GMAIL.COM
Haim Sompolinsky[2] HAIM@FIZ.HUJI.AC.IL

[1]Harvard University
[2]Hebrew University of Jerusalem

During perceptual learning, animals improve accuracy of discrimination between similar stimuli through training. The cortical plasticity associated with this process is poorly understood, particularly at the level of networks. We studied optimal plasticity patterns for perceptual leaning of discrimination between a pair of noisy stimuli, in a multi-layer, feedforward model of a sensory system. In particular, we asked how it can perform Maximum Likelihood Discrimination (MLD) while minimizing perturbation to its responses and performances for other ecologically relevant stimuli. When model neurons have identical receptive field filters, it is in general necessary to induce changes in the earliest processing stage. Without such changes, it is difficult for readout mechanisms, linear or non-linear, to recover MLD performance in latter stages. Indeed, a gradient descent learning causes a wide spread changes in all stages of the network, disrupting the responses of the network to untrained stimuli — a form of catastrophic forgetting. Widespread plasticity is particularly necessary if primary cortical neurons (the first processing layer) have narrow receptive fields. On the other hand, in a wide and shallow architecture, whereas the primary sensory stage consists of many diverse receptive field filters, MLD can be implemented by a linear readout without modifying these filters. These results suggest that in perceptual learning tasks with large input noise, modifications to early sensory areas are important for perceptual learning, and that diversity of receptive field types is advantageous for performing perceptual learning with less interference of other tasks.

T-18. What does a dopamine ramp mean?

HyungGoo Kim HKIM07@G.HARVARD.EDU
Naoshige Uchida UCHIDA@MCB.HARVARD.EDU

Harvard University

Recent studies have shown that the dopamine (DA) concentration in the ventral striatum gradually ramps up as the animal approaches a reward location. Such DA ramps have been proposed to encode a gradually increasing value as defined by an estimate of temporally discounted future reward, or 'state value' (Howe et al., 2013; Hamid et al., 2016). However, the notion that DA represents 'value' contradicts the canonical view that DA activity represents reward prediction error (RPE) or temporal difference (TD) error, which is, before obtaining a reward, approximately the derivative of the value function. Gershman (2013) showed theoretically that if the value function conforms a particular shape (a convex function), DA ramps may result from a derivative-like RPE computation (i.e. TD error). To distinguish these possibilities – state value versus RPE, we devised a set of experimental tests using visual virtual reality in mice. We monitored dopamine axon activity from the ventral striatum using fiber photometry while mice ran on a 1-dimensional virtual corridor to obtain a reward. Our results showed that DA ramped as the animal approached a reward location. In some trials, the animal was teleported from an intermediate location to a location closer to the reward. The state value hypothesis predicts a step-wise change in DA activity, whereas the RPE hypothesis predicts a phasic (transient) excitation at the time of teleport. Our results showed a phasic activation of DA axons. Furthermore, the magnitude of these DA transients was larger when a teleportation of the same distance occurred closer to a reward location, consistent with a convex value function. These results indicate that DA ramp occurs due to dynamic sensory inputs that indicate the proximity to reward, and DA ramps are consistent with TD error, i.e. a derivative-like computation over a convex value function across space.

T-19. Dopaminergic and frontal basis of decisions guided by sensory evidence and reward value

Armin Lak ARMINLAK@GMAIL.COM
Michael Okun M.OKUN@LEICESTER.AC.UK
Morgane Moss MORGANE.MOSS.17@UCL.AC.UK
Harsha Gurnani HARSHAGURNANI@GMAIL.COM
Karolina Farrell KAROLINA.FARRELL@UCL.AC.UK
Miles Wells MILES.WELLS@UCL.AC.UK
Charu Reddy C.REDDY@UCL.AC.UK
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini MATTEO@CORTEXLAB.NET

University College London

Efficient decisions require combining immediate sensory evidence with past reward values. It is not known how the brain performs this combination and learns from outcomes of such decisions. We trained mice in a decision task that requires combining sensory evidence and past rewards: a visual detection task in which reward size for correct decisions changed across blocks of trials. Mice behaved efficiently in this task, shifting their psychometric curves so that reward size had a larger influence in trials with low contrast stimuli, where sensory uncertainty is higher. Their decisions were well described by a reinforcement learning model that included an estimate of sensory uncertainty. The model computes two key internal variables that drive learning: the predicted value of the decision and the uncertainty-dependent prediction error. These variables depend on both reward size and trial-by-trial sensory uncertainty. We established a causal neural correlate of the first variable – predicted reward - by recording and manipulating the activity of neurons in prelimbic frontal cortex. These neurons responded mainly during choice execution, and reflected the predicted reward. Crucially, optogenetic manipulation of these neuronal responses did not influence ongoing decisions but caused learning: it affected subsequent decisions. We established a necessary and sufficient correlate of the second variable – uncertainty-dependent prediction error - by imaging and optogenetic manipulation of midbrain dopamine neurons. Fiber photometry revealed that these neurons' activity depended on reward size and sensory uncertainty, consistent with encoding of prediction errors. Further, consistent with prediction error encoding, optogenetic activation and suppression of dopaminergic neurons shifted the psychometric curves in opposite directions. These results provide a quantitative framework for decisions guided by reward value and sensory evidence, and reveal neuronal computations that enable the brain to learn and make efficient decisions when challenged with internal and environmental uncertainty.

T-20. Serotonin neurons track reward history and modulate flexible behavior

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

Johns Hopkins University

In order to survive, animals must navigate complex and dynamic environments. Doing so requires using previous experience to guide behavior and updating that behavior with each subsequent experience. The neural mechanisms of learning and decision making that underlie this adaptive behavior are largely unknown, but prior research points to a role for neuromodulatory neurons in the updating process. Among these populations, serotonin neurons are hypothesized to regulate flexible decision making after a change in action-outcome contingency. To test this hypothesis, we designed a dynamic foraging task for head-fixed mice that is amenable to extracellular electrophysiological recordings from optogenetically-identified serotonin neurons. In the task, mice exhibited flexible behavior, using reward history to select actions in a probabilistic and dynamic environment. Our recordings demonstrate a relationship between serotonin neuron activity and recent reward history, regardless of the actions taken.

To probe the meaning of this signal, we adapted a reinforcement learning model that tracks reward history in order to generate an expectation about average reward rate—an expectation which is used to modulate learning. In a sparsely rewarding environment in which an animal has learned to expect low reward rates, for example, changes in behavior should be more strongly driven by a reward than lack of reward. First, we found that the activity of serotonin neurons correlated with the model-generated expected average reward rate. Next, we ablated serotonin neurons, effectively removing expected average reward rate representations. Consistent with the prediction of the model, selective lesions of serotonin neurons in this task resulted in more perseverative behavior, characterized by an increase in learning from better-than-expected outcomes and a decrease in learning from worse-than-expected outcomes. These results suggest an important function for serotonin neuron activity in driving flexible behavior in a dynamic environment.

T-21. Inhibition stabilization is a widespread feature of mouse cortical networks

Alessandro Sanzeni[1] ALESSANDRO.SANZENI@GMAIL.COM
Bradley Akitake[1] BRADLEY.AKITAKE@NIH.GOV
Nicolas Brunel[2] NICOLAS.BRUNEL@DUKE.EDU
Mark Histed[3] MARK.HISTED@NIH.GOV

[1]NIMH/NIH
[2]Duke University
[3]National Institutes of Health

Strong recurrent synaptic coupling in networks like the cerebral cortex can provide several advantages for computation, allowing networks to track fast feedforward inputs, amplify weak signals, and perform noise-robust classification. Theoretical studies have shown that networks of excitatory and inhibitory cells with sufficiently strong recurrent excitation are unstable in the absence of inhibitory feedback, and networks in this general class are thus termed “inhibition-stabilized networks” (ISNs).

Cortical intracellular recordings show behavior consistent with ISN models in some conditions; however, in many cases, a disinhibitory mechanism without ISN regime could explain the observations. Moreover, modeling studies have suggested that inhibition stabilization, and thus a significant amount of recurrent input, may arise only in the presence of external or sensory activity [Ahmadian,2013]. Therefore, it is still unknown whether recurrent connections in cortical circuits are strong enough to produce inhibition-stabilization, and whether this property persists across areas and states, in particular in the absence of sensory stimulation.

Here we look for ISN behavior in several cortical areas and states by testing a key prediction of ISN models: inhibitory neurons should show a paradoxical suppression of firing when the inhibitory network is stimulated [Tsodyks,1997]. Pairing optogenetic stimulation with in vivo pharmacology to identify inhibitory cells in awake mice, we observe that several cortical networks - visual, auditory, somatosensory, and motor cortex - are stabilized by inhibition even at rest, i.e. when sensory stimuli are not being delivered. Further, we found inhibitory stabilization was robust to reductions of network activity during light anesthesia. A quantitative analysis of the network response shows that recurrent coupling in mouse cortex is moderately strong, suggesting mouse visual cortex operates above, but near, the transition from non-ISN to ISN.

Our results suggest that inhibition stabilization is a widespread phenomenon in mouse cortex across a range of activity levels and network states.

T-22. Low threshold layer-specific neural ensembles contributing to the ignition of perception

James Marshel JMARSHEL@STANFORD.EDU
Timothy Machado TAMACHADO@STANFORD.EDU
Sean Quirin SQUIRIN@STANFORD.EDU
Yoon Kim KYS8892@STANFORD.EDU
Brandon Benson BENSONB@STANFORD.EDU
Jonathan Kadmon KADMONJ@STANFORD.EDU
Cephra Raja CPRAJA@STANFORD.EDU
Adelaida Chibukhchyan ADELAIDA@STANFORD.EDU
Charu Ramakrishnan CHARUR@STANFORD.EDU
Masatoshi Inoue MINOUE2@STANFORD.EDU
Hideaki Kato HEKATO@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Karl Deisseroth DEISSERO@STANFORD.EDU

Stanford University

While neural activity patterns generate perceptual experience, it is unclear which specific patterns are minimally sufficient to causally generate individual percepts. Addressing this question requires simultaneous read-out and play-in of naturally occurring cellular patterns in behaving animals. To identify cellular determinants of perception, we developed an optical read-write system capable of kilohertz write speed, and 3D cortical access both across superficial to deep layers, and across large cortical surface areas. This neural interface required two innovations: (1) a new opsin with unprecedented properties including red-shifted light sensitivity, extremely large photocurrents, millisecond spike-timing fidelity, and compatibility with simultaneous two-photon imaging; (2) new custom-engineered spatial light modulators capable of stimulating 3D neuronal ensembles of hundreds of individually-specified neurons across a millimeter of cortex at 1 kHz temporal resolution.

We applied our tools to mouse primary visual cortex during a visual cue-guided behavioral task. We found: (1) stimulation of naturally-recruited ensembles containing as few as 20 neurons both selectively ignited widespread population responses among functionally-related neurons and ignited percepts sufficient to guide correct visual discriminations without external sensory input; (2) the threshold for ignition was lower in layer 5 versus 2/3, and ignition selectively propagated from layer 2/3 to 5; (3) these effects could be enhanced by, but did not require task learning; (4) multineuronal dynamics evoked by visual and optogenetic stimulation were strikingly similar; (5) these shared dynamics correlated with behavior across animals, with stronger neural ignition yielding better performance. Theoretical analysis could account for such data through neural models with excitatory subnetworks poised near dynamic instability, enabling small selective (but not random) neural ensembles to trigger ignition. Overall our results reveal a striking operating regime for visual cortex, and generally enable precise control over natural neural codes across broad 3D cortical space with fine spatiotemporal resolution.

T-23. A premotor circuit responsible for flexible decision making

Zheng Wu[1] ZW2308@COLUMBIA.EDU
Ashok Litwin-Kumar[1] AK3625@COLUMBIA.EDU
Philip Shamash[2] PHILIP.SHAMASH.17@UCL.AC.UK
Alexei Taylor[1] AT3175@CUMC.COLUMBIA.EDU
Richard Axel[1] RA27@COLUMBIA.EDU
Michael Shadlen[1] MS4497@CUMC.COLUMBIA.EDU

[1]Columbia University
[2]University College London

Cognitive flexibility allows us to connect evidence across time and make context-dependent decisions. We

adapted a delayed match to sample (DMS) task in mice to study flexible decision making. Mice must remember a sample odor over a delay, then compare it to a test odor to determine if it is a match or non-match, and indicate the decision by licking left or right, respectively. We ask 1) How is sample memory represented? and 2) How is the match/non-match decision made?

Recordings from piriform, orbitofrontal, and anterolateral premotor cortices (Pir, OFC, ALM) revealed a graded representation of odor identity and choice across these areas, with odor identity most prominent in Pir and choice dominant in ALM. Decoding analysis suggests that sufficient information is present in OFC and Pir to solve match/non-match, and this information may instruct ALM to plan licking.

If ALM is simply a readout of upstream areas, it should only be essential once test odor arrives. Surprisingly, we found that silencing ALM during sample and delay epochs dramatically impaired animal's performance. Importantly, silencing did not affect licking itself or control tasks that do not require sample memory. This suggests that ALM plays a critical role in processing the sample odor identity to plan licking.

Additional observations led us to favor a synaptic gating model over the classical attractor network model. First, silencing during the late delay produces less impairment than silencing during the sample+delay, suggesting that the sample memory is not fully explained by persistent spiking. Second, two-photon calcium imaging of ALM uncovered enrichment of neurons in the superficial layer-2 (L2) that exhibited selectivity for sample odor over the memory delay. These observations raise the possibility that sample identity configures the premotor circuit to establish the appropriate sensorimotor mapping of test odor to lick response.

T-24. Anterior cingulate cortex directs exploration of alternative strategies

Dougal Tervo[1] GOWANTERVO@GMAIL.COM
Elena Kuleshova[1] KULESHOVA.ELENA.P@GMAIL.COM
Max Manakov[1] MANAKOV.MAKSIM@GMAIL.COM
Mikhail y[1] M.PROSKURIN@GMAIL.COM
Mattias Karlsson[2] MKARLSSO@GMAIL.COM
Alla Karpova[1] ALLA.KARPOVA@GMAIL.COM

[1]HHMI Janelia Research Campus
[2]Spike Gadgets

The ability to adjust one's behavioral strategy without instruction in complex environments is at the core of higher cognition. It has been posed that the brain simplifies the problem of choosing appropriate behavioral strategies in uncertain situations by first re-evaluating previously learned strategies before deciding whether to construct new ones. Functional imaging studies in humans have supported this idea and uncovered activation of anterior cingulate cortex (ACC) during resurgence of previously learned strategies. Other findings in humans, non-human primates and rodents have implicated ACC in keeping track of alternative learned strategies and in learning higher order statistics about the environment that can inform internal estimates of a strategy's reliability. However, whether ACC plays an active role in strategy switching, or merely monitors behavioral performance remains unclear. We provide causal evidence that ACC actively arbitrates between persisting with the ongoing strategy and switching away from it temporarily to re-evaluate learned alternatives. Using a foraging task that isolates this point of commitment, we found a neurophysiological signature of a resurgent strategy in ACC and demonstrate that transient perturbation of ACC activity at the point of commitment significantly reduces resurgence of the alternative strategy. In contrast, perturbation during the presumed time when the ongoing strategy is re-evaluated leads to a significant increase in subsequent exploration of the alternative, suggesting that a separate computation within ACC actively suppresses the region's own ability to countermand the ongoing strategy. Furthermore, we establish that these two computations are segregated within ACC micro-circuitry by demonstrating that they can be individually manipulated by perturbing distinct output pathways through viral-mediated access. We suggest that when an internal estimate in the intra-telecenphalic sub-circuit calls for countermanding the ongoing strategy, activity in the pyramidal tract is released from inhibitory influence and facilitates rapid behavioral adaptation by sampling from the learned strategy set.

T-25. Amygdala ensembles encode behavioral states

Jan Grundemann[1] JAN.GRUNDEMANN@UNIBAS.CH
Yael Bitterman[2] YAEL.BITTERMAN@FMI.CH
Benjamin Grewe[3] BGREWE@ETHZ.CH
Tingjia Lu[2] TINGJIA.LU@FMI.CH
Sabine Krabbe[2] SABINE.KRABBE@FMI.CH
Mark J Schnitzer[4] MSCHNITZ@STANFORD.EDU
Andreas Luthi[2] ANDREAS.LUTHI@FMI.CH

[1]University of Basel
[2]Friedrich Miescher Institute for Biomedical Research
[3]ETH Zurich
[4]Stanford University

Internal states, including affective or homeostatic states, are important behavioral motivators. The amygdala is a key regulator of motivated behaviors, yet how distinct internal states are represented in amygdala circuits is unknown. Here, by longitudinally imaging neural calcium dynamics across different environments in freely moving mice, we identify changes in the activity levels of two major, non-overlapping populations of principal neurons in the basal amygdala (BA) that predict switches between exploratory and non-exploratory (defensive, anxiety-like) states. Moreover, the amygdala broadcasts state information via several output pathways to larger brain networks, and sensory responses in BA occur independently of behavioral state encoding. Thus, the brain processes external stimuli and internal states orthogonally, which may facilitate rapid and flexible selection of appropriate, state-dependent behavioral responses.

T-26. Exploration via disrupted sensorimotor control dynamics

Becket Ebitz[1] REBITZ@GMAIL.COM
Jon Cohen[2] JDC@PRINCETON.EDU
Tim Buschman[2] TBUSCHMA@PRINCETON.EDU
Tirin Moore[3] TIRIN@STANFORD.EDU
Benjamin Hayden[1] BENHAYDEN@GMAIL.COM

[1]University of Minnesota
[2]Princeton University
[3]Stanford University

In variable or uncertain environments, solely “exploiting” rewarding options is not the best strategy. Instead, intelligent decision-makers also “explore” alternatives: forgoing some immediate rewards in order to gather information. Fifty years of neurobiological research has focused on the neural basis of reward-maximizing exploit choices, but much less is known about the neural basis of exploratory decisions. Understanding these neural mechanisms could provide new insights into the computations used for flexibility and discovery in healthy brains. Here, we examined how decision-making differed across exploration and exploitation, using a combination of computational modelling, electrophysiology, and causal manipulations in rhesus macaques. The converging results from each of these approaches suggest that exploration occurs when the constraints on prefrontal activity are disrupted. During exploitation, prefrontal activity is tightly constrained and, as a result, it traces out repeated trajectories through neural state space. However, during exploitation, the forces that act to confine neural activity to these trajectories are reduced—producing noisier neural trajectories that are easier to perturb. Relaxing the constraints on prefrontal activity could be the mammalian mechanism for the behavioral noise necessary for discovery and learning. This view makes a testable prediction: if prefrontal control is disrupted during exploration, exploratory decisions should depend more on low-level, hardwired biases—such as our preference for high-contrast stimuli. Indeed, we found precisely this effect in psychophysical studies done in three monkeys and twenty-five Princeton undergraduates. Together, these results support the hypothesis that we explore, in part, via releasing prefrontal

control over behavior and suggest that adjustments in the strength of the constraints on prefrontal activity may be a critical mechanism for regulating prefrontal control.

T-27. Multiple maps of the same context stably coexist in the mouse hippocampus

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

Weizmann Institute of Science

Hippocampal place cells selectively fire when an animal traverses a particular location, and are considered a neural substrate of spatial memory. While place cells remap across spatial contexts by changing their firing rates and/or firing locations, within a fixed familiar context it is generally accepted that their spatial tuning is consistent over time. Here, we show spontaneous remapping between multiple stable spatial representations of the same familiar context, without any apparent changes in sensory input, experience, behavior, or motivational state. Alternations between representations occurred only between trials, and were coherent between the recorded population of cells, consistent with the auto-associative properties attributed to the hippocampus. The distinct representations were spatially informative and persistent over weeks, indicating they can be reliably stored and retrieved from long-term memory. We found components within the spatial code that were shared across spatial representations of the same context as well as different contexts, suggesting a reuse of preconfigured network states. Our results raise the question of whether a memory associated with a given spatial context could be supported by multiple neuronal representations, rather than just one.

T-28. Representation of 3D space in the entorhinal cortex of flying bats

Gily Ginosar[1] GILYGINOSAR@GMAIL.COM
Johnatan Aljadeff[2] J.ALJADEFF@IMPERIAL.AC.UK
Yoram Burak[3] YORAM.BURAK@ELSC.HUJI.AC.IL
Haim Sompolinsky[3] HAIM@FIZ.HUJI.AC.IL
Liora Las[1] LIORALAS@GMAIL.COM
Nachum Ulanovsky[1] NACHUM.ULANOVSKY@WEIZMANN.AC.IL

[1]Weizmann Institute of Science
[2]Imperial College London
[3]Hebrew University of Jerusalem

The medial entorhinal cortex (MEC) contains a variety of spatial cells, including grid cells and border cells. In 2D, grid cells fire when the animal passes near the vertices of a 2D spatial lattice (or grid), which is characterized by circular firing-fields separated by fixed distances, and 60^o local angles – resulting in a hexagonal structure. Although many animals navigate in 3D space, no studies have examined the 3D volumetric firing of MEC neurons. Here we addressed this by training Egyptian fruit bats to fly in a large room (5.8x4.6x2.7m), while we wirelessly recorded single neurons in MEC. We found 3D border cells and 3D head-direction cells, as well as many neurons with multiple spherical firing-fields. 20% of the multi-field neurons were 3D grid cells, exhibiting a narrow distribution of local distances between neighboring fields – but not a perfect 3D global lattice. The 3D grid cells formed a functional continuum with less structured multi-field neurons. Both 3D grid cells and multi-field cells exhibited an anatomical gradient of spatial scale along the dorso-ventral axis of MEC, with inter-field spacing increasing ventrally – similar to 2D grid cells in rodents. We modeled 3D grid cells and multi-field cells as emerging from

pairwise-interactions between fields, using an energy potential that induces repulsion at short distances and attraction at long distances. Our analysis shows that the model explains the data significantly better than a random arrangement of fields. Interestingly, simulating the exact same model in 2D yielded a hexagonal-like structure, akin to grid cells in rodents. Together, the experimental data and preliminary modeling suggest that the fundamental property of grid cells is multiple fields that repel each other with a characteristic distance-scale between adjacent fields – which in 2D yields a global hexagonal lattice while in 3D yields only local structure but no global lattice.

T-29. Detailing the network structure and downstream function of a Drosophila ring attractor

Dan Turner-Evans TURNEREVANSD@JANELIA.HHMI.ORG
Saba Ali ALIS@JANELIA.HHMI.ORG
Arlo Sheridan SHERIDANA@JANELIA.HHMI.ORG
Tyler Paterson PATERSONT@JANELIA.HHMI.ORG
Kris Jensen JENSENK@JANELIA.HHMI.ORG
Rob Ray RAYR@JANELIA.HHMI.ORG
Scott Lauritzen LAURITZENS@JANELIA.HHMI.ORG
Davi Bock BOCKD@JANELIA.HHMI.ORG
Vivek Jayaraman VIVEK@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

Bees, ants, and other insects display impressive path integration behavior when returning to their hives or nests after foraging. Flies too appear to rely on path integration when performing a local search near places where they have recently encountered food. During path integration, animals keep track of direction and distance to a goal as they move through their environment, an ability that is thought to rely on internal representations. Such representations have been described in mammals and —in the case of heading direction— in flies. While significant progress has been made in understanding how these activity patterns might be generated and updated, the network connections predicted by most theoretical models have not been tested experimentally. We sought to address this shortcoming by using electron microscopy and RNA sequencing to reconstruct a biologically implemented ring attractor for compass direction in the Drosophila central brain. We then explored the roles of the constituent cell classes with two-color, two-photon calcium imaging and genetic perturbation. The exposed attractor architecture is complex beyond existing theories. It contains local, reciprocal connections between its cell classes that may help to maintain the activity and two distinct longer-range feedback loops that interact to update the heading representation as the animal turns. The network also contains an interneuron that passes the animal's heading downstream by inhibiting partners everywhere except at the heading location. We performed calcium imaging on these downstream cell classes and discovered that activity in at least one class conjunctively encodes both heading and forward velocity. The population activity of this class could thus theoretically be integrated to trace the animal's path in space, linking direction to distance estimation as would be required for navigation in two-dimensions. Together, our findings indicate a path towards understanding the mechanistic implementation of a broadly relevant navigational computation.

T-30. Population activity of cerebellar granule cells in awake mice

Natasha Cayco Gajic[1] CAYCOGAJIC@GMAIL.COM
R Angus Silver[1] A.SILVER@UCL.AC.UK
Frederic Lanore[2] FREDERIC.LANORE@U-BORDEAUX.FR

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

The cerebellum is thought to learn sensorimotor relationships to coordinate movement. In the cerebellar input layer, granule cells integrate sensory and motor information from neocortex (via the pontine nucleus) and primary sensory areas. Classic work by Marr and Albus posited that the cerebellar cortex separates overlapping input patterns by projecting them into a high-dimensional granule cell population code. Recent theoretical work has also shown that the cerebellar circuitry is well-suited for pattern separation and associative learning based on Marr-Albus theory. However, how the granule cell population activity encodes sensory and motor information is poorly understood, and whether granule cell populations can support high-dimensional representations has not been tested in mammalian cerebellum, due to the technical challenges involved. To address this, we expressed GCaMP6f in granule cells in the Crus 1 region of the cerebellar hemisphere of mice. We took advantage of high-speed random-access 3D two-photon microscopy to simultaneously monitor the Ca2+ activity in hundreds of granule cell axons (parallel fibers, PFs) of awake animals for the first time, while recording locomotion and whisker pad movement with high-speed video cameras. To better understand how the cerebellar input layer controls active movement, we analyzed PF population activity to determine if locomotion or whisking activity recruits specific population of PF axons. We find that PF population activity transitions between separate, orthogonal coding spaces representing periods of quiet wakefulness verses active movement. Finally, we characterise the dimensionality of PF population activity.

T-31. Head direction and orienting-suppression signals in primary visual cortex of freely-moving rats

Grigori Guitchounts[1] GUITCHOUNTS@FAS.HARVARD.EDU
Javier Masis[1] JMASIS@FAS.HARVARD.EDU
Steffen Wolff[1] STEFFENWOLFF@FAS.HARVARD.EDU
David Cox[2] DAVID.D.COX@IBM.COM

[1]Harvard University
[2]IBM

Recent studies have brought into question the idea that information in sensory cortex is only sensory. Signals related to locomotion or navigation have been observed in several sensory cortical brain regions. Neurons in primary visual cortex (V1), for example, modulate their firing rates based on the animal's running speed, even in the absence of visual input. V1 activity has recently also been shown to be modulated by perceived position in an environment and passive head turning. Most experiments of this nature have been performed in head-fixed mice whose range of motion is limited, leaving open the question of how V1 is modulated by the variety of movements free animals make. We recorded neuronal activity continuously (24/7) for up to three weeks in rat V1 using tetrodes while the animals lived in their home cages. Coupling the recordings with measurements of behavior using a head-mounted accelerometer/gyroscope/magnetometer allowed us to probe the relationship between V1 activity and natural movement, while the continuous nature of the recordings allowed us to examine differences in activity at the two phases of the light cycle (light and dark). We find that V1 activity linearly encodes the animal's 3D head direction (yaw, roll, and pitch), even in complete darkness. Further, V1 dynamics are strongly modulated by turning (left, right, up, down, clockwise, counterclockwise). While activity is generally suppressed during turns in the dark, it is increased at the end of turns in the light. A linear classifier can distinguish different turn types. Bilateral lesions of secondary motor cortex (M2), a major source of cortical feedback to V1, greatly diminish turn-related activity, both in the light and in the dark, while preserving responses to visual stimulation and head direction encoding, suggesting that the latter signals originate elsewhere (e.g. retrosplenial cortex, another major input to V1). These results support predictive coding theories of cortical function and suggest that corollary discharge signals in mammalian brains contain detailed movement and head direction information.

T-32. Predicting natural behavior from brain-wide neural activity

Monika Scholz MSCHOLZ@PRINCETON.EDU
Ashley Linder ASHLEY.N.LINDER@GMAIL.COM
Francesco Randi FRANDI@PRINCETON.EDU
Anuj Sharma AKSHARMA@PRINCETON.EDU
Xinwei Yu XINWEIY@PRINCETON.EDU
Joshua Shaevitz SHAEVITZ@PRINCETON.EDU
Andrew Leifer LEIFER@PRINCETON.EDU

Princeton University

We record calcium activity from the majority of head neurons in freely moving C. elegans to reveal where and how natural behavior is encoded in a compact brain. We find that a sparse subset of neurons distributed throughout the head encode locomotion. A linear combination of these neurons' activity predicts the animal's velocity and body curvature and is sufficient to infer its posture as it crawls. This sparse linear model outperforms single neuron or PCA-based models at predicting behavior. Among neurons important for the prediction are well-known locomotory neurons, such as AVA, as well as neurons not traditionally associated with locomotion. We compare neural activity of the same animal during unrestrained movement and during immobilization and find large differences between brain-wide neural dynamics during real and fictive locomotion. To our knowledge, this is the first work to provide brain-wide neural decoding of behavior in an unrestrained animal. By leveraging the unique capabilities of C. elegans we address questions about the extent, location and coding of behavior-related signals in the brain.

T-33. Dissociating task acquisition from expression during learning reveals latent knowledge

Tom Hindmarsh Sten[1] THINDMARSH@MAIL.ROCKEFELLER.EDU
Eleni Papadoyannis[2] ESP304@NYU.EDU
Kelly Fogelson[3] KFOGELS1@JHMI.EDU
Sarah Elnozahy[3] SELNOZA1@JHU.EDU
Celine Drieu[3] CDRIEU1@JHU.EDU
Rupesh Kumar Chillale[4] RUPESH.KUMAR.CHILLALE@ENS.FR
Yves Boubenec[4] YVES.BOUBENEC@ENS.FR
Robert Froemke[2] ROBERT.FROEMKE@MED.NYU.EDU
Peter Holland[3] PCH@JHU.EDU
Srdjan Ostojic[4] SRDJAN.OSTOJIC@ENS.FR
Kishore Kuchibhotla[3] KKUCHIB1@JHU.EDU

[1]Rockefeller University
[2]New York University
[3]Johns Hopkins University
[4]Ecole Normale Superieure

Performance on cognitive tasks during learning is used to measure knowledge, yet it remains controversial since such testing is susceptible to context. To what extent does performance during learning depend on the testing context, rather than underlying knowledge (1-3)? We trained mice, rats and ferrets to examine how testing context impacts the acquisition of knowledge versus its expression. We interleaved reinforced trials with "probe" trials in which we omitted reinforcement. Across tasks, all animal species performed remarkably better in probe trials during learning and inter-animal variability was strikingly reduced. Reinforcement feedback is thus critical for learning-related plasticity but, paradoxically, hides the expression of underlying knowledge. Probing learning by omitting reinforcement uncovers this latent knowledge and identifies context, and not knowledge, as the source of individual variability. To identify computational mechanisms that might underlie these observations, we built a network model in which we hypothesized that learning occurs during reinforced trials while the changes between

contexts modulate only the read-out parameters in the output layer. This minimal model provided a parsimonious description of our large and diverse behavioral dataset. Preliminary neural data from the auditory cortex support this dissociation and point to distinct neural mechanisms for task acquisition and expression.

T-34. A value-based explanation for lapses in perceptual decisions

Sashank Pisupati SPISUPAT@CSHL.EDU
Lital Chartarifsky-Lynn LCHARTAR@CSHL.EDU
Anup Khanal KHANAL@CSHL.EDU
Anne Churchland CHURCHLAND@CSHL.EDU

Cold Spring Harbor Laboratory

During perceptual decisions, even well-trained subjects can display a constant rate of errors independent of evidence strength, assumed to arise from inattention or motor errors. These are referred to as “lapses” and their proper treatment is crucial for accurately estimating perceptual parameters (Prins 2012) and interpreting inactivations (Erlich 2015), however the factors influencing them remain poorly understood. Here, we propose uncertainty-guided exploration as an underlying cause for lapses. We demonstrate that perceptual uncertainty modulates the probability of lapses both within and across modalities on a multisensory discrimination task in rats. These effects cannot be accounted for by inattention or motor error, however they are concisely explained by a normative model of uncertainty-dependent exploration (Gershman 2018) We tested the predictions of the exploration model by changing the reward magnitude associated with one of the actions. As predicted by the model, this affected the lapses asymmetrically, mimicked the effect of the reward history associated with that action, and did not affect “sure-bet” decisions where the animal always exploited. Moreover, reversing the reward contingency leads to a gradual increase and subsequent decrease in lapses as predicted. Finally, we demonstrate that unilateral muscimol inactivations of secondary motor cortex and posterior striatum affect lapses asymmetrically across modalities. The inactivations do not affect "sure-bet" decisions and are captured by a reduction in contralateral action value in the model, suggesting that these areas encode the expected value of actions based on multisensory evidence. Together, our results suggest a value-based explanation for lapses, and that far from being a nuisance, lapses are informative about individual animals’ exploration-exploitation tradeoff.

T-35. Striatum encodes task adaptative expectations and mediates their impact on perceptual decisions

Ainhoa Hermoso Mendizabal[1] AINHOAHEM@GMAIL.COM
Alexandre Hyafil[2] ALEXANDRE.HYAFIL@GMAIL.COM
Pavel Ernesto Rueda Orozco[3] PAVEL.RUEDA@GMAIL.COM
David Robbe[4] DAVID.ROBBE@INSERM.FR
Jaime de la Rocha[1] JROCHAV@CLINIC.CAT

[1]IDIBAPS
[2]Universitat Pompeu Fabra
[3]Universidad Nacional Autonoma de Mexico
[4]Institut de Neurobiologie de la Mediterranee

Prior experiences shape the way we perceive the world by creating expectations, a reference frame for future decisions. Little is known however about where in the brain these expectations are represented and which are the neural dynamics underlying their potential contribution in decisions. We trained rats in two-alternative forced choice auditory task, where the probability to repeat the previous stimulus category was varied in trial blocks. All rats capitalized on the serial correlations of the stimulus sequence by consistently exploiting a transition bias: a tendency to repeat or alternate their previous response using an internal trial-by-trial estimate of the sequence

repeating probability. Surprisingly, the transition bias was null in trials immediately following an error. The internal estimate however was not reset and became effective again causing a bias after the next correct response. Thus, rats used behavioral outcomes to flexibly modulate how expectations influenced their decisions. We performed bilateral pharmacological inactivations in the dorsomedial striatum (DMS). The inactivation diminished significantly the transition bias and the stimulus impact on choice suggesting a necessary role of DMS in the expectation biases generation. Neural population recordings showed that DMS neurons carried task relevant information across trials. In particular, neurons were not only selective to the present response but also to the present transition (Repetition vs Alternation). Importantly, this information was carried out to the following trial, but only when the current trial was rewarded. The representation of unrewarded responses and transitions was in contrast extinguished after revealing the negative trial outcome. Thus, as found in the behavior, response outcome gated the information flow about these actions towards the future possibly reflecting the flexible modulation of history biases in upcoming choices. Together, our results suggest that DMS dynamically encodes the relevant variables that are necessary to build the expectation biases exhibited by our animals.

T-36. Measuring and modeling the weight dynamics of many synapses onto diverse cell-types in vivo

*Aran Nayebi[1] ANAYEBI@STANFORD.EDU
*Joshua Melander[2,1] MELANDER@STANFORD.EDU
Bart Jongbloets[2] JONGBLOE@OHSU.EDU
**Tianyi Mao[2] MAO@OHSU.EDU
**Haining Zhong[2] ZHONG@OHSU.EDU
**Surya Ganguli[1] SGANGULI@STANFORD.EDU

[1]Stanford University
[2]Oregon Health and Science University

Dendritic spines have been studied intensely for over a century, partially because such spines, corresponding to excitatory synapses onto pyramidal cells, can be easily visualized with traditional light microscopy. However, many synapses cannot be identified on the basis of morphology alone, either because they are present on aspiny neurons or occur on spines that project along the axial plane of the microscope light path. Here, we directly tag a synaptic protein (PSD-95) and visualize cortical excitatory synapses onto spiny pyramidal neurons, as well as onto aspiny inhibitory interneurons, in vivo. We performed, for the first time, month-long longitudinal two-photon imaging of thousands of synapses onto three cell-types: L2/3 pyramidal (PYR), PV+, and VIP+ neurons. We found robust differences in synapse stability between cell types. In addition to observing previously unseen spines and shaft synapses, we demonstrated that PSD-95 fluorescence strongly correlates with synaptic strength (measured via patch-clamp physiology and two-photon glutamate uncaging), thereby allowing us to track, in an unprecedented fashion, the simultaneous dynamics of thousands of individual synaptic strengths, on diverse cell-types, over a month. Using PSD-95 fluorescence as an indicator of synaptic strength, we found qualitatively different synaptic weight distributions onto each of the three cell-types.

This dataset enables us to quantitatively model synaptic strength dynamics through a discrete-time KxK Markov transition matrix M that describes transitions between K different binned synaptic strengths, including zero strength (absence of a synapse). Our learned transition matrices (from a semi-automatically scored dataset) accurately predicted the observed cell-type specific stationary distributions of synaptic strength. Intriguingly, while the Markov transition matrices describing synaptic dynamics are asymmetric, they obey detailed balance for all three cell-types. This provides preliminary evidence that across cell-types, synaptic strength changes obey time-reversible dynamics that could be the result of stochastic descent on a nontrivial energy landscape.

T-37. Is synaptic potentiation required for associative learning in adult fruit fly?

Richard Lonsdale RICHARD.LONSDALE3@BTINTERNET.COM
Pedro F Jacob JOAQUIM.FARIAJACOB@CNCB.OX.AC.UK
Bayar Menzat BAYAR.MENZAT@CNCB.OX.AC.UK
Scott Waddell SCOTT.WADDELL@CNCB.OX.AC.UK
Tim P Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Drosophila melanogaster (fruit fly) learns to associate a specific odour with an unpleasant experience, and avoid that odour in the future. These memories reside in the fly mushroom body (MB), where odour-signalling Kenyon cells (KCs) excite mushroom body output neurons (MBONs) that drive approach behaviour. Dopaminergic neurons (DANs) modulate the KC-MBON connections. These DANs respond to an unpleasant experience (e.g. electric shock), releasing dopamine that depresses KC-MBON synapses for recently activated KCs. The associated odour drive to approach MBONs is reduced, thus increasing avoidance of that odour. Synaptic depression is crucial for associative learning in drosophila. The role of synaptic potentiation is less clear, although recent evidence shows potentiation of KC-MBON synapses following shocks without an accompanying odour. We have constructed a computational model of part of the MB to investigate how synaptic depression and potentiation interact to change fly behaviour. The model comprises a rate-based representation of the main neuron types (KC, DAN & MBON), with elevated DAN activity driving KC-MBON synaptic plasticity such that synapses from recently activated KCs are depressed, while inactive synapses are potentiated. Applying the model to reversal learning (aversive training applied twice, with odours reversed on second application), shows that potentiation enables the fly to acquire and prioritise new memories whilst retaining older memories. We also identify a potential mechanism for relief learning (shock followed by odour). We suggest that DANs modulating avoidance-driving MBONs may be inhibited during a shock. A subsequent rebound in DAN activity, paired with odour exposure, could establish a relief memory driving approach to that odour. Finally, we compare behavioural experiments on the odour preferences of flies after 10 different aversive training protocols, showing that our simple model of plasticity can explain a surprising number of behavioural results.

T-38. Flexible mapping of visual scenes to a heading representation via inhibitory Hebbian plasticity

Sung Soo Kim[1] SUNGSOO@JANELIA.HHMI.ORG
Ann Hermundstad[1] HERMUNDSTADA@JANELIA.HHMI.ORG
Larry Abbott[2] LFABBOTT@COLUMBIA.EDU
Vivek Jayaraman[1] VIVEK@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]Columbia University

Many animals can acquire and maintain their bearings in a variety of different surroundings, flexibly mapping sensory scenes onto internal representations of heading and place. In Drosophila melanogaster, heading is represented by the angular position of a localized bump of population activity of E-PG neurons, each arborizing in a single sector of the torus-shaped ellipsoid body (EB). A population of visual-feature-selective ring neurons each make inhibitory synapses onto the entire population of E-PG neurons around the EB, providing the ideal substrate for a flexible mapping of visual features to heading. Indeed, the E-PG bump position in the same visual environment varies across flies, and sometimes even in the same fly over time, strongly suggesting that the pinning 'offset' between landmarks and the angular position of the EB bump is not stereotyped, but develops over time with experience. We suggest that this variability of the pinning offset is the natural consequence of a circuit that maps arbitrary visual scenes to a stable heading representation, and hypothesize that it is realized via competitive and inhibitory Hebbian plasticity between ring neurons and E-PG neurons. To test our hypothesis that

synaptic weights between these neurons are depressed when they are co-active, we put tethered flies in a virtual reality arena and used two-photon calcium imaging and optogenetics to enforce an artificial offset between the E-PG bump position and landmarks for 3-5 minutes. We found that the E-PG offset to landmark cues was shifted towards the artificially enforced value, consistent with our hypothesis. Our results provide direct physiological evidence of population-level Hebbian plasticity in the fly compass system. Given that E-PGs receive both visual and angular velocity inputs, this plasticity may allow animals to quickly reconcile and register self-motion- and landmark-based cues to stably maintain and update their heading in different visual environments.

T-39. Orientation preference maps in the mouse lemur confirm common design principles in primate visual cortex

Chun Lum Andy Ho[1] CHUN.HO@UNIGE.CH
Robert Zimmermann[1] ROBERT.ZIMMERMANN@UNIGE.CH
Juan Weidinger[2] CHEPE@NLD.DS.MPG.DE
Banty Tia[3] BANTY.TIA@MNHN.FR
Fabien Pifferi[3] FABIEN.PIFFERI@MNHN.FR
Fabienne Aujard[4] FABIENNE.AUJARD@MNHN.FR
Fred Wolf[2] FRED@NLD.DS.MPG.DE
Daniel Huber[1] DANIEL.HUBER@UNIGE.CH

[1]University of Geneva
[2]Max Planck Institute for Dynamics and Self-Organization, Gottingen
[3]Museum national d'histoire naturelle, Paris
[4]Musee National d'Histoire Naturelle, Paris

In mammalian primary visual cortex (V1) two major modes of representing orientation information have been described. Whereas in rodents V1 neurons with different orientation selectivities are found intermingled in a salt-and-pepper fashion, they are organized in orientation pinwheel maps in all primates studied so far. Is this difference in representation due to constraints of the smaller brain size in rodents or can pinwheels also be found in primates with rodent-sized brains? To address this question we performed intrinsic optical signal imaging in the mouse lemur (Microcebus murinus), one of the smallest primates (60g), with a brain size comparable to the rat. Our results show that mouse lemur V1 contains robust orientation pinwheel maps with properties comparable to other primates and adhering to a common design framework. Surprisingly, the spacing of pinwheels was found to be of similar size compared to other primates being up to two orders of magnitude larger in size. The seemingly incompressible nature of orientation pinwheels might also be related to the relative size of V1, especially in light of the animal's visual acuity. As a pinwheel column can be considered as a functional unit in early visual processing, the number of pinwheels found in V1 can give a good estimation of the visual acuity and vice versa. Our behavioural data on mouse lemur visual acuity suggests that they require a relatively expanded V1 area given their small brain size in order harbor the appropriate number of pinwheels. Indeed, immunohistochemical analysis revealed that the mouse lemur has one of the highest V1 to neocortex ratios found in primates (20%). Taken together, this study not only revealed the existence of pinwheels in the mouse lemur V1, but the seemingly incompressibility of this basic functional unit might explain the over proportional V1 size found among small primates.

T-40. Sustained and dynamic representations reduce interference in short-term memory

Alexandra Libby ALEXGLIBBY@GMAIL.COM
Tim Buschman TBUSCHMA@PRINCETON.EDU

Princeton University

Nothing in the environment is experienced in isolation. Previous events, memories and associations all impact sensory processing. For example, animals and humans use the memory of recent stimuli to predict future stimuli, allowing them to respond faster and more accurately to a stimulus. Despite extensive study, the neural mechanisms that maintain representations in short-term memory are not well understood, particularly in service of predictions. To address this, we explored how the brain holds sensory representations in short-term memory during an implicit learning paradigm. Naive mice were exposed to sound sequences consisting of four chords (i.e. ABCD and XYC*D) over four days. The sequences were constructed such that the first two stimuli (A-B or X-Y; the 'context') predicted the identity of the third stimulus (C or C*, respectively). Thus, the animal had to use the short-term memory of the context information to predict the upcoming, third stimulus. Multiple-electrode recordings in primary sensory cortex found neurons learned the predictive association between A/X and C/C* by forming a common representation (i.e. neurons responded to A and C). However, this association also led to a 'postdiction' response: the sensory representation of the third stimulus interfered with the sensory representation of the previously experienced context stimulus. Despite the sensory representation of context being lost, we found neurons in primary sensory cortex maintained a separate mnemonic representation of context. Critically, the mnemonic representation was orthogonal to the sensory representation, reducing interference with new, incoming stimuli. Furthermore, we found that orthogonalization of the mnemonic representation was the result of a combination of both sustained and dynamic responses to the contextual stimulus in auditory cortex. These results shed light on recent debates on whether working memory representations are sustained or dynamic, showing that their combination can protect short-term memory from interference.

T-41. Laminar-specific cortico-cortical loops in mouse visual cortex

Hedi Young[1,2] HEDI.YOUNG@NEURO.FCHAMPALIMAUD.ORG
Beatriz Belbut[2] BEATRIZ.BELBUT@RESEARCH.FCHAMPALIMAUD.ORG
Leopoldo Petreanu[1,2] LEOPOLDO.PETREANU@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Champalimaud Centre for the Unknown

The mammalian neocortex is hierarchically organized such that connections between areas may be divided into 2 types: feedforward (FF) connections from lower to higher areas, and feedback (FB) connections in the opposite direction. Finding fundamental rules of connectivity for FF and FB projections is necessary to constrain theories of hierarchical cortical computation. Some theories require looped FF<–>FB interactions while others do not. However, it is not known whether cortico-cortical (CC) connections target cortical neurons with specific projections. Here we measured the connectivity of FF and FB connections with the 3 major classes of projection neurons (PT, CT & IT) using subcellular channelrhodopsin-2-assisted circuit mapping (sCRACM) in combination with retrograde tracers. We recorded from pairs of neighboring neurons in the same cortical layer (L) in primary visual cortex (V1) or the higher order lateromedial (LM) visual area in acute brain slices containing FB or FF ChR2-expressing axons, respectively. For each pair, one cell projected to the source of the ChR2-expressing FF or FB inputs, and one cell projected to a different cortical or subcortical area. FF and FB innervated projection neurons in L2/3, L5 and L6. Both FF and FB inputs made strong connections with IT neurons projecting back to the source while avoiding adjacent PT L5 and CT L6 neurons. Moreover, both FF and FB inputs selectively formed loop connections with L6 neurons projecting back to the source area. FB inputs formed loops with L5 neurons projecting back to the source by selectively targeting their apical dendrites in layer 1. However, neither FF nor FB inputs preferentially target L2/3 neurons projecting back. The wiring specificity of CC connections supports models of hierarchical processing requiring loop connectivity but suggests that only L6 and the distal tufts of the apical dendrites of L5 neurons are involved in this type of computations.

T-42. Learning a latent manifold of odor representations in piriform cortex

Anqi Wu[1] ANQIW@PRINCETON.EDU
Stan Pashkovski[2] PASHKOVS@HMS.HARVARD.EDU
Sandeep Datta[2] SRDATTA@HMS.HARVARD.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Harvard University

A major difficulty in studying the neural mechanisms underlying olfactory perception is the lack of obvious structure in the relationship between odorants and the neural activity patterns they elicit. Here we use odor-evoked responses in piriform cortex to identify a latent manifold specifying latent distance relationships between olfactory stimuli. Our approach is based on the Gaussian process latent variable model, and seeks to map odorants to points in a low-dimensional embedding space, where distances between points in the embedding space relate to the similarity of population responses they elicit. The model is specified by an explicit continuous mapping from a latent embedding space to the space of high-dimensional neural population firing rates via nonlinear tuning curves, each parametrized by a Gaussian process. Population responses are then generated by the addition of correlated, odor-dependent Gaussian noise. We fit this model to large-scale calcium fluorescence imaging measurements of population activity in layers 2 and 3 of mouse piriform cortex following the presentation of a diverse set of odorants. The model identifies a low-dimensional embedding of each odor, and a smooth tuning curve over the latent embedding space that accurately captures each neuron's response to different odorants. The model captures both signal and noise correlations across more than 500 neurons. We validate the model using a cross-validation analysis known as co-smoothing to show that the model can accurately predict the responses of a population of held-out neurons to test odorants.

T-43. Inferring brain-wide neuronal dynamics on single trials via neural stitching of many Neuropixels recordings

Daniel O'Shea[1] DJOSHEA@STANFORD.EDU
William Allen[1] WALLEN@STANFORD.EDU
Liqun Luo[1] LLUO@STANFORD.EDU
Karl Deisseroth[1] DEISSERO@STANFORD.EDU
David Sussillo[2,1] SUSSILLO@GOOGLE.COM

[1]Stanford University
[2]Google AI

An emerging goal of systems neuroscience is to understand how complex neuronal dynamics enable the brain's ability to select, plan, and execute adaptive behaviors. While recent progress has been made in understanding neuronal population dynamics in one or a few recorded brain regions, naturalistic behaviors require the coordination of many interconnected regions throughout the brain. Fortunately, it is now becoming possible to record densely throughout the brain with new electrophysiology technology, although not yet fully simultaneously. We begin to address the challenge of modeling global brain dynamics by developing an extension to LFADS (latent factor analysis via dynamical systems). As in LFADS, our method infers the latent, single trial dynamics of neural spiking data and performs neural dynamical stitching across separate recording sessions. Our extension leverages the brain-region location of each recorded neuron to construct region-specific latent dynamics as a subset of the global brain latent dynamics in the model.

We apply this method to a dataset comprising 23,881 neurons identified within 34 brain regions across 87 Neuropixels recording sessions in 21 mice performing an olfactory-decision task. We find that our modified, dynamically-stitched LFADS model can indeed model this large dataset at the single trial level. We find that (1) the model successfully recapitulates gating of the odor-response dynamics by a specific, persistent, and brain-wide pattern of population activity that encodes motivational state at the single trial level, and (2) the satiety state of the animal is reliably encoded in the latent initial condition, and the odor identity in the model's inferred inputs. In summary, this modeling tool enables researchers to leverage non-overlapping electrophysiology datasets from across the brain to extract better estimates of the global latent dynamics that drive behavior. This approach further enables comparison of the contributions of individual brain regions to the global computation.

T-44. High-order network interactions in neural population representations revealed using vine copulas

Houman Safaai[1,2] HOUMAN_SAFAAI@HMS.HARVARD.EDU
Selmaan Chettih[1] SELMAAN.CHETTIH@GMAIL.COM
Stefano Panzeri[2] STEFANO.PANZERI@IIT.IT
Christopher Harvey[1] HARVEY@HMS.HARVARD.EDU

[1]Harvard University
[2]Istituto Italiano di Tecnologia

Neural processes in cortex, including sensory perception, decision-making, and movement planning and execution require the coordinated function of populations of neurons. A major open question in systems neuroscience is to understand the emergent properties in a neural population, beyond what can be inferred from single neurons. To address this question, an ideal model should capture the general interaction structure between large populations of simultaneously recorded neurons and provide measures of the information in representations at any order and scale of the network. Probabilistic and information theoretic approaches to address this question, such as maximum entropy, have been proposed. In these models, the statistical dependencies are estimated by defining constraints on the interaction moments of neural activity up to a few orders of interaction. These methods are largely limited by the fact that the space of possible neural responses becomes exponentially large when considering higher-order interactions, especially for general response distributions and when considering stimulus dependencies. Here, we propose a new approach in which we model any general population dependency structure between neurons as an expert mixture of vine copulas. Using nonparametric copulas, we decompose the full density function between the neural responses into a set of bivariate structures that can be computed in a sequential fashion, each of which is equivalent to defining constraints of all possible orders in a maximum entropy model. We show that, corresponding to the vine sequential structure, the total information represented in the network can be decomposed into individual components related to each interaction order and, because of the bivariate structures, the approach does not suffer from the curse of dimensionality. We used vine copula modelling to study population representations in posterior parietal and visual cortical neural populations during a navigation-based choice task and investigated the contributions of high-order interactions.

I-1. An astroglial-neuromodulatory system integrates unsuccessful actions to trigger giving-up behavior

Yu Mu[1] AUMUYU@GMAIL.COM
Davis Bennett[1] BENNETTD@JANELIA.HHMI.ORG
Mikail Rubinov[2] MIKARUBI@GMAIL.COM
Sujatha Narayan[1] NARAYANS@JANELIA.HHMI.ORG
Chao-Tsung Yang[3] CHAOTSYN@GMAIL.COM
Masashi Tanimoto[1] TANIMOTOM@JANELIA.HHMI.ORG
Brett Mensh[1] BMENSH@GMAIL.COM
Loren Looger[1] LOOGERL@JANELIA.HHMI.ORG
Misha Ahrens[1] AHRENSM@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]Vanderbilt University
[3]Fred Hutchinson Cancer Research Center

Animals initiate actions to respond to incoming stimuli and changes in their environment, and reach goals. Equally importantly, animals can terminate actions when behaviors consistently fail to achieve their goals. The neural substrates of such “giving up” behavior are largely unknown. How does the brain detect behavioral failures, accumulate repeated failures, and trigger a switch in behavior? By imaging whole-brain neuronal activity in larval zebrafish in a paradigm where they “give up” after becoming unable to swim forward, and simultaneously monitoring the behavior under “successful” and “failed” swim bouts in a virtual reality system, we identified a circuit performing the computation and mediating the behavioral state switches (Figures 1 and 2). Noradrenergic neurons in the medulla oblongata encode the failure signal, which is based on the amount of motor effort and extent of sensory motor mismatch. The failure information transfers to and accumulates in astroglia in lateral hindbrain, through projections from noradrenergic neurons to the gliapil and neuropil region. Perturbation experiments show that calcium elevation in noradrenergic neurons and lateral-hindbrain astroglia is necessary and sufficient for triggering the switch from the active to the passive state (Figure 3). Thus, noradrenergic neurons and astroglia together perform a computation critical for behavior: they encode and accumulate evidence that current behavior is ineffective and consequently drive changes in brain and behavioral states (Figures 4 and 5).

I-2. Critical behavior and hysteresis in a cortical model with a reservoir of spatiotemporal patterns

Silvia Scarpetta[1] SILVIASKA@GMAIL.COM
Antonio De Candia[2] ADECANDIA@GMAIL.COM
Ilenia Apicella[3] ILENIA.APICELLA@HOTMAIL.IT
Ludovico Minati LMINATI@IEEE.ORG
Alessandro Iannone[4] ALESSANDRO.IANNONE93@GMAIL.COM

[1]University of Salerno
[2]University of Napoli
[3]University of Padova
[4]University of Pisa

Many experimental results, both in vivo and in vitro, support the idea that the brain cortex operates near a critical point and at the same time works as a reservoir of precise spatiotemporal patterns. However, the mechanism at the basis of these observations is still not clear. In this work we introduce a model which combines both these features, showing that scale-free avalanches are the signature of a system posed near the spinodal line of a first-order transition, with many spatiotemporal patterns stored as dynamical metastable attractors. Specifically, we studied a network of leaky integrate-and-fire neurons whose connections are the result of the learning of multiple spatiotemporal dynamical patterns, each with a randomly chosen ordering of the neurons.We found that

the network shows a first-order transition between a low-spiking-rate disordered state (down), and a high-rate state characterized by the emergence of collective activity and the replay of one of the stored patterns (up). The transition is characterized by hysteresis, or alternation of up and down states, depending on the lifetime of the metastable states. In both cases, critical features and neural avalanches are observed.

I-3. Inhibitory microcircuits for top-down plasticity of sensory representations

Katharina Anna Wilmes KATHARINA.WILMES@GMAIL.COM
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK

Imperial College London

Animals learn better when it matters to them. For example, they learn to discriminate sensory stimuli when they receive a reward. As a result of learning, neural responses to sensory stimuli are adjusted even in the first processing stages of sensory areas (Goltstein et al. 2013, Khan et al. 2018, Poort et al. 2015). It is thought that behaviourally relevant contexts, such as rewards, trigger an internal top-down signal available to these early sensory circuits. This could be mediated by cholinergic inputs from the basal forebrain for example (Letzkus et al. 2011). One challenge remains: contextual signals are typically present for a short time only, but synaptic changes require time to be expressed. How can these time scales be bridged? We hypothesise that interneuron circuits, which recently emerged as key players during learning and memory, bridge the timescales. We investigate how temporary top-down modulation by rewards can interact with local excitatory and inhibitory plasticity to induce long-lasting changes in sensory circuitry. We propose that learning can happen in two stages: 1) Unspecific top-down signals rapidly induce an inhibitory connectivity structure between different interneuron types. 2) The inhibitory structure induces changes in sensory representations by guiding excitatory plasticity between pyramidal cells. Using a computational model of layer 2/3 primary visual cortex, we demonstrate how inhibitory microcircuits could store information about the rewarded stimulus to guide long-term changes in excitatory connectivity in the absence of further reward. We make specific testable predictions in terms of activity of different neuron types. The suggested two-stage plasticity mechanism in canonical cortical microcircuits could be conserved across different modalities.

I-4. Genetically driven wiring of the C. elegans connectome

Daniel Barabasi[1] BARABASI@FAS.HARVARD.EDU
Istvan Kovacs[2] KOVACS.PISTI@GMAIL.COM
Albert-Laszlo Barabasi[3] BARABASI@GMAIL.COM

[1]Harvard University
[2]Northeastern University
[3]Northeastern University, Harvard University, Central European University

An important challenge of modern neuroscience is to unveil the mechanisms shaping the wiring of the connectome, answering the difficult question of how the brain wires itself. Lacking well established principles, current models assume random wiring patterns with biologically plausible spatial constraints. Yet, connectomes display a high degree of wiring reproducibility, so that multiple circuits and architectural features appear to be identical within a species, invariants that the current models are unable to explain. It is increasingly established that encoding the connectome requires both transcriptionally encoded neuronal identity and genetically driven mechanisms that encode the formation of specific synapses and gap junctions. Here we show that such transcriptionally encoded mechanisms can be described by unique mathematical operators, leading to imprints in the connectome, as a detectable biclique motif. We validate the proposed model in C. elegans, by identifying the underlying bicliques, and showing that they are absent in graphs of comparable size and network characteristics, indicating their biological roots. The model predicts that neurons in each biclique must share expression patterns and morphological features, helping gain further direct evidence of its predictive power and helping us to identify the specific genes that potentially determine the local wiring patterns. Evidence of bicliques in higher organisms suggests that the model captures the general features of the genetic mechanisms governing brain wiring. Exploring specific bicliques reveals genetically deterministic functional connectivity motifs.

I-5. A multiplication rule for integrating corticostriatal and thalamostriatal signals during reward-guided movements

Kwang Lee[1] KWANGLEE@MEDNET.UCLA.EDU
Konstantin Bakhurin[2] KONSTANTIN.BAKHURIN@DUKE.EDU
Leslie Claar[3] LDCLAAR@GMAIL.COM
Sotiris Masmanidis[1] SMASMANIDIS@UCLA.EDU

[1]University of California, Los Angeles
[2]Duke University
[3]Allen Institute for Brain Science

An essential aspect of signal processing in the brain is the transformation of synaptic input into neuronal output. Cortical and thalamic projections provide excitatory drive to the striatum, but little is known about how these inputs, either individually or collectively, regulate striatal dynamics during behavior. The lateral striatum is involved in reward-guided licking, and receives overlapping input from the secondary motor cortex (M2) as well as parafascicular thalamic nucleus (PF). Here we investigated how behaviorally evoked neural dynamics in the lateral striatum are shaped by three inputs corresponding to the ipsilateral and contralateral M2, and ipsilateral PF. Mice were trained to perform anticipatory licking movements in response to reward-associated cues. Using neural recordings and optogenetic terminal inhibition we examined the contribution of M2 and PF projections on medium spiny projection neuron (MSN) activity. By comparing the effect of suppressing one, two, or three projections on firing rate, we found that MSNs appear to combine cortical and thalamic input signals through multiplication. In support of a multiplicative effect, the gain of the neuronal response varied with the number of simultaneously suppressed inputs. These results reveal a simple yet computationally powerful signal processing function of the striatum.

I-6. The causal role of unsupervised temporal learning in the development of V1 complex cells

Giulio Matteucci GMATTEUC@SISSA.IT
Davide Zoccolan ZOCCOLAN@SISSA.IT

International School for Advanced Studies (SISSA)

Visual perception relies on cortical representations of visual objects that remain relatively stable with respect to the tremendous variation in object appearance typically experienced during natural vision (e.g., because of position, size or viewpoint changes). Such stability, known as transformation tolerance, is built incrementally along the cortical hierarchy devoted to shape processing. Already in primary visual cortex (V1), early instances of position tolerance can be found, with neurons known as “complex cells” responding to their preferred feature across slightly offset positions. To date, the origin of transformation tolerance, in V1 as well as in higher-order visual areas, remains poorly understood. One of the leading theories, known as “unsupervised temporal learning”, postulates that visual neurons exploit the temporal continuity of visual experience (i.e., the natural tendency of different object views to occur nearby in time) to associatively link temporally-contiguous stimuli, so as to factor out object identity from other faster-varying, lower-level visual attributes (e.g., object positions, size, etc.). To causally test this hypothesis, we reared newborn rats in visually controlled environments, where the animals were exposed to either natural movies (control group) or to their frame-scrambled versions (resulting in temporally unstructured visual input; experimental group) for the whole duration of the critical period of visual cortical development. Following this controlled rearing phase, we performed multi-electrode extracellular recordings from V1 of each animal, finding a reduction of the proportion of complex cells in the experimental group and a concomitant decrease in the ability of such neurons to code stimulus orientation in a position-invariant way. These findings causally demonstrate that temporal continuity of the visual input plays an important role in establishing the tolerance properties of visual neurons during postnatal development. This empirically validates a class of unsupervised learning principles postulated by prominent theories of visual cortical processing and development.

I-7. Adversarial examples influence human visual perception

Gamaleldin Elsayed[1] GAMALELDIN.ELSAYED@GMAIL.COM
Shreya Shankar[2] SS.SHANKAR505@GMAIL.COM
Brian Cheung[3] BCHEUNG5@GMAIL.COM
Nicolas Papernot[4] NICOLAS@PAPERNOT.FR
Alexey Kurakin[1] KURAKIN@GOOGLE.COM
Ian Goodfellow[1] GOODFELLOW@GOOGLE.COM
Jascha Sohl-Dickstein[1] JASCHASD@GOOGLE.COM

[1]Google Brain
[2]Stanford University
[3]University of California, Berkeley
[4]Pennsylvania State University

Machine learning models are vulnerable to adversarial examples: small changes to images can cause computer vision models to make mistakes such as identifying a school bus as an ostrich. However, it is still an open question whether humans are prone to similar mistakes. One interesting phenomenon is that adversarial examples often transfer from one model to another, making it possible to attack models that an attacker has no access to. This naturally raises the question of whether adversarial examples similarly transfer to humans. Clearly, humans are prone to many cognitive biases and optical illusions, but these generally do not resemble small perturbations, nor are they generated by optimization of a machine learning loss function. Thus, susceptibility to adversarial examples has been widely assumed – in the absence of experimental evidence – to be a property of machine learning classifiers, but not of human judgement. A rigorous investigation of the above question creates an opportunity both for machine learning and neuroscience. If we knew conclusively that the human brain could resist a certain class of adversarial examples, this would provide an existence proof for a similar mechanism in machine learning security. If we knew conclusively that the brain can be fooled by adversarial examples, then machine learning security research should perhaps shift its focus from designing models that are robust to adversarial examples to designing systems that are secure despite including non-robust components. Likewise, if adversarial examples developed for computer vision affect the brain, this phenomenon discovered in the context of machine learning could lead to a better understanding of brain function. Here, we investigate this question by studying the effect of adversarial examples on human visual perception. We find that adversarial examples that strongly transfer across computer vision models influence the classifications made by time-limited human observers.

I-8. Efficient phase coding in hippocampal place cells

Pavithraa Seenivasan PAVITHRAA@IISC.AC.IN
Rishikesh Narayanan RISHI@IISC.AC.IN

Indian Institute of Science

Hippocampal place cells encode space using both firing rates and the precession of spike phase with reference

to an extracellular theta oscillation. Phase precession is considered as a network-driven phenomenon within the temporal sequence compression (TSC) framework, emerging primarily through a speeding in the intracellular theta with respect to its extracellular counterpart. Consequently, the significance of neuronal intrinsic properties in defining and regulating this phase code has largely remained unexplored. Here, we built a conductance-based model for phase precession within the TSC framework to investigate the role of neuronal excitability in regulating encoding efficiency. We defined efficiency of the phase code within a place field using an information maximization framework. We recruited an unbiased stochastic search strategy to generate thousands of models that span a multi-dimensional parametric space to assess the physiological constraints on the concomitant emergence of encoding efficiency and excitability robustness. We presented the same afferent inputs to models with distinct intrinsic properties and compared their phase codes to find that the efficiency of the phase code was critically reliant on neuronal intrinsic properties. Strikingly, despite this critical reliance on intrinsic properties, we found that disparate parametric combinations could elicit similar efficiency in phase coding and similar excitability properties, pointing to a significant explosion in the degrees of freedom available for an encoding system to achieve efficient encoding and robust homeostasis. Finally, to assess experience dependence of place fields, we presented asymmetric place field inputs to these model neurons. This resulted in asymmetry of the firing rate profile and a leftward shift of the phase precession curve. Our study defines efficiency of phase coding in place cells and unveils a critical role of neuronal intrinsic properties in achieving such efficiency, while also suggesting degeneracy as a framework for solving the twin goals of efficient encoding and robust homeostasis.

I-9. Arousal modulates retinal inputs to superior colliculus

Sylvia Schroeder[1] SYLVIA.SCHROEDER@UCL.AC.UK
Nicholas A Steinmetz[1] N.STEINMETZ@UCL.AC.UK
Michael Krumin[1] M.KRUMIN@UCL.AC.UK
Matteo Rizzi[1] M.RIZZI@UCL.AC.UK
Marius Pachitariu[2] MARIUS10P@GMAIL.COM
Leon Lagnado[3] L.LAGNADO@SUSSEX.AC.UK
Kenneth D Harris[1] KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini[1] M.CARANDINI@UCL.AC.UK

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

The activity of multiple regions in the early visual system, including primary visual cortex (V1), is modulated by non-visual factors such as locomotion and arousal. Here we show that this modulation is present even earlier, at the activity of retinal ganglion cells (RGCs). Using two-photon imaging, we recorded neuronal activity in superficial superior colliculus (sSC) in response to drifting gratings, uniform gray screens or during darkness while mice were head-fixed and free to run on a treadmill. To track arousal level, we measured running speed and pupil diameter. About 40% of sSC neurons showed significant positive and negative correlations with running and pupil diameter both during spontaneous and visually driven activity. Effects of arousal on direction tuning curves were heterogeneous across neurons, and could be described with a linear model. To establish whether these behavioral effects were inherited from V1, we sued Neuropixels probes to record from sSC while inactivating ipsilateral V1 optogenetically. V1 inactivation decreased average visual responses in sSC, but did not significantly decrease behavioral modulation. We then imaged calcium activity in synapses of RGCs onto sSC neurons. To our surprise, synaptic activity was correlated with locomotion and arousal. This occurred even during darkness, excluding the possibility that changing light levels reaching retina due to pupil dilation caused the change in responses. Direction tuning was modulated similarly to sSC neurons. To test whether modulation of retinal synapses originates in the modulation of RGC firing rates, rather than synaptic neuromodulation, we recorded from retinal axons in the optic tract, in darkness. In some axons, firing rates were modulated by running. We conclude that behavioral modulation of activity in the early visual system may be due to behavioral modulation of retinal inputs. This modulation may occur both at the retinofugal synapses and in the activity of retinal ganglion cells.

I-10. In search of optimal exploration

Lior Fox LIOR.FOX@MAIL.HUJI.AC.IL
Yonatan Loewenstein YONATAN.LOEWENSTEIN@MAIL.HUJI.AC.IL

Hebrew University of Jerusalem

Balancing exploration and exploitation is essential when learning operant tasks. Exploitation is defined as choosing the action associated with the highest value (expected future accumulated rewards). Reinforcement learning (RL) algorithms can estimate these values, and finding their neural correlates has dominated the field of neuroeconomics. Exploration, by contrast, while necessary for these algorithms, has received relatively little attention. Even its goal has remained poorly defined. In practice, exploration is implemented by adding stochasticity to the action-selection, or in more sophisticated learning, by utilizing visit counters. Both approaches are clearly suboptimal, particularly because they do not take into account the future accumulated knowledge gained from specific actions. Here we postulate that an optimal exploration policy is one that maximizes the entropy of the visitation distribution over the state-actions. In other words, a policy that results in visits of all state-action pairs as uniformly as possible. We show that given the model of the environment, formulated as a Markov Decision Process, finding such optimal exploration policy can be reduced to the standard optimization problem of finding a maximum entropy distribution under linear constraints. This is analogous to the finding of an optimal policy (solving the Bellman optimality equation) using dynamic programming. In practice, however, an agent does not have access to the full model of the environment. Therefore, we developed a model-free algorithm that learns to explore using value-like exploration variables. When combined with standard RL approaches, our algorithm has been shown to improve the performance in large-scale problems. We show numerically in several examples that this algorithm approximates the maximum-entropy distribution of state-action visitations. These exploration variables and their updates can be used when seeking the neural correlates of exploration representation, similar to the way that action-values and TD-error have been used to identify the neural correlates of operant learning.

I-11. Circuit dynamics and plasticity after eye opening and activity deprivation in visual cortex of freely behaving rodents

Yue Wu[1] YUE.WU@BRAIN.MPG.DE
Keith B. Hengen[2,3] KHENGEN@WUSTL.EDU
Gina G Turrigiano[2] TURRIGIANO@BRANDEIS.EDU
Julijana Gjorgjieva[1] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Brandeis University
[3]Washington University in St. Louis

Neural circuits utilize various homeostatic plasticity mechanisms to maintain stability, while executing precise behavioral actions, and flexibility, while undergoing learning and experience-dependent plasticity. An appropriate paradigm to study these processes is the visual cortex of rodents during the classical critical period (4-6 weeks after birth). The critical period is characterized by a high degree of plasticity that can be induced by manipulating visual experience, for instance, suturing the lids of one eye, a process known as monocular deprivation (MD). We examined extensive datasets of electrophysiological recordings over 9 days of the collective activity of multiple cells in the monocular region of V1 in freely behaving rodents during normal development 2-3 weeks after eye opening and after MD. Previous studies have shown that while firing rates of individual neurons remain relatively stable during this period, MD induces an initial drop in firing rates followed by their homeostatic recovery despite long-lasting deprivation (Hengen et al. Neuron 2013, Hengen et al. Cell 2016). We investigated higher-order network properties during normal development and prolonged MD. In contrast to firing rates, pairwise correlations increased during the period of 2-3 weeks after eye opening, while MD induced considerable weakening of pairwise correlations followed by their homeostatic recovery. To understand how the network exploits diverse homeostatic mechanisms to return firing rates and correlations to baseline after prolonged MD, we tackled the problem by a plastic spiking recurrent neural network model of leaky integrate-and-fire neurons. We examined plastic and homeostatic synaptic changes that dynamically shape cortical networks during sensory deprivation consistent with the experimentally measured firing rate and correlation changes. This approach enables us to explore the individual functionality of distinct plasticity mechanisms following activity deprivation.

I-12. Evidence accumulation through cyclic AMP synchronizes network activity to estimate elapsed time

Stephen Thornquist[1,2] THORNQUIST@FAS.HARVARD.EDU
Maximilian Pitsch[3,4] MAXI-PITSCH@WEB.DE
Charlotte Auth[3,4] CHARLOTTE.AUTH@STUD-MAIL.UNI-WUERZBURG.DE
Michael Crickmore[3,4] MICHAEL.CRICKMORE@CHILDRENS.HARVARD.EDU

[1]Crickmore Lab
[2]FM Kirby Neurobiology Center
[3]Boston Children's Hospital
[4]Harvard University

Actions, perceptions, and decisions often hinge on evidence accumulated over long timescales and from multiple sources. In many cases, these timescales are millions of times longer than the time constant of membrane potentials, posing a challenge for their representations by purely electrical means. We show that during a Drosophila mating four male-specific Corazonin neurons collect evidence about the passage of time. At the onset of mating, high CaMKII activity within each neuron blocks cyclic AMP elevation that would otherwise result from recurrent Corazonin-network activity. As time passes, CaMKII activity declines, allowing cyclic AMP to accumulate, driving Protein Kinase A to enhance network activity. The network eventually achieves criticality, causing a sustained eruption in calcium levels that is synchronized across all four neurons. The eruption signals the conclusion of a six-minute time interval, driving sperm transfer and a simultaneous shift in motivation. The many highly conserved and broadly expressed pathway components involved suggest this ancient module may allow diverse circuitries to achieve and report consensus decisions using conflicting information distributed across neuronal networks. Neurons likely use these evolutionarily ancient strategies to retain and process much more sophisticated information than their immediate electrical input provides, and importantly such representations are invisible to typical measurements of neuronal computation, such as calcium levels and spiking activity

I-13. Event-to-event variability in autoassociative spiking dynamics during ripples

Xinyi Deng[1] XD2188@COLUMBIA.EDU
Shizhe Chen[2] SZDCHEN@UCDAVIS.EDU
Xuexin Wei[1] XW2481@COLUMBIA.EDU
Liam Paninski[1] LMP2107@COLUMBIA.EDU
Loren Frank[3] LOREN@PHY.UCSF.EDU

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

The ability to flexibly remember experiences at different levels of specificity is crucial to how we learn and how we make decisions. However, the underlying mechanisms involved in flexibly storing and retrieving memories in varying degrees of detail remain elusive. Current theories suggest that memories of past experiences are stored when specific patterns of neural activity cause changes in the connections among neurons, and they are retrieved when these patterns are reinstated. For example, when an animal moves through its environment, spiking activity

of hippocampal place cells is paced by an underlying "clock" with a constant rate at theta frequency. When the animal slows down or pauses, place cell population is sequentially reactivated during sharp wave-ripples (SWRs), which often represents a replay of past trajectory on a compressed time scale. Is the hippocampal clock also constant during the SWR state? Here we report that rhythmicity underlying the SWR-associated organization of population spiking exhibits high event-to-event variability on a continuous scale of 6-50 Hz. This continuous scale is observed across laps, familiarity of spatial environment, and types of replay content and is approximately lognormal. Decoding analyses using clusterless methods further suggest that during awake immobility, same spatial experience are replayed in multiple SWR events, each time with a different clock and the rate of this clock is randomly sampled from a lognormal distribution. We propose that such changing clock might constitute a general mechanism for flexible memory consolidation.

I-14. Neural straightening of natural videos in macaque primary visual cortex

Olivier Henaff[1] HENAFF@CNS.NYU.EDU
Yoon Bai[2] YOONBAI@UTEXAS.EDU
Julie Charlton[2] JCHARLTON@UTEXAS.EDU
Ian Nauhaus[2] INAUHAUS@GMAIL.COM
Eero Simoncelli[1] EERO.SIMONCELLI@NYU.EDU
Robbe Goris[2] ROBBE.GORIS@UTEXAS.EDU

[1]New York University
[2]University of Texas at Austin

Many behaviors rely on predictions derived from recent sensory input. Making visual predictions is challenging because naturally occurring light patterns evolve in a complex, nonlinear manner. We hypothesize that the visual system transforms its inputs so as to straighten their temporal trajectories, thereby enabling prediction through linear extrapolation. Previously, we provided psychophysical support for this theory (CoSyNe 2018). Here, we test the temporal straightening hypothesis in primate visual cortex. We first compared the straightness of natural videos with the straightness of neural population activity elicited by these videos. We presented random sequences of static frames taken from 10 short videos and used multi-electrode arrays to record V1 population activity in anesthetized macaque monkeys. We obtained temporal trajectories of population activity by arranging neural responses in the videos' natural order. Estimating the average curvature of a temporal trajectory is straightforward in the (deterministic) image domain, but challenging in the (noisy) neural domain. As such, we developed a data-efficient and unbiased estimator of neural curvature. We found that V1 systematically straightens natural videos. Straightening occurred over multiple timescales (30–100 ms), and appears to be specific to natural videos. We created synthetic videos that fade from the first to the last frame of each of the natural videos. These movies are straight in the image-domain, but elicited neural trajectories that were much more curved than expected from chance. We then asked which computations give rise to neural straightening. We simulated responses of a hierarchical model that captures the primary nonlinear transformations performed by the early visual system (retina–LGN–V1). The model's V1 response trajectories mimicked our observations. In contrast, deep neural networks trained for object recognition do not straighten natural videos. Temporal straightening may thus be an objective that specifically shapes the computations of the primate visual system.

I-15. Temporal expectation signal in the basal forebrain predicts decision speed

Shih-Chieh Lin[1] SHIHCHIEH.LIN@GMAIL.COM
Ksenia Vlasov[2] KSENIA.VLASOV@GMAIL.COM
Hachi Manzur[2] HEMANZUR@GMAIL.COM
Alessandro Scaglione[2] ALESSANDRO.SCAGLIONE@GMAIL.COM

[1]National Yang-Ming University
[2]National Institute on Aging, NIH

The speed of decision making in response to a stimulus is determined not only by the motivational salience of the stimulus but also by internal expectations about when the stimulus will occur, the latter of which is referred to as temporal expectation. Previous studies have found that the influence of motivational salience on decision speed is mediated by the phasic bursting response of noncholinergic basal forebrain (BF) neurons. However, it remains unknown whether the influence of temporal expectation on decision speed is mediated by the same neuronal circuit. Here we show that the same BF neurons also encode temporal expectation and their neuronal activity predicts decision speed. In rats trained to perform simple reaction time (RT) tasks under different foreperiod distributions, stronger temporal expectations of stimulus onset were coupled with stronger inhibitions of BF neuronal activity during the pre-stimulus foreperiod. Moreover, the decrease of BF neuronal activity during the foreperiod was quantitatively coupled with a rebound excitation after stimulus onset. The rebound excitation occurred shortly after the initial phasic bursting response to stimulus onset, and its strength was coupled with faster RTs. Together, these results support the conclusion that noncholinergic BF neurons integrate both motivational salience and temporal expectation, and serve as a common gain control mechanism to regulate decision speed.

I-16. Credit-assignment to state-independent task representations

Nitzan Shahar[1,2] SHAHAR.NITZAN@GMAIL.COM
Tobias Hauser[1,2] T.HAUSER@UCL.AC.UK
Rani Moran[1] RANI.MORAN@GMAIL.COM
Rogier Kievit[3] ROGIER.KIEVIT@MRC-CBU.CAM.AC.UK
Ray Dolan[1] R.DOLAN@UCL.AC.UK

[1]University College London
[2]Max Planck UCL Centre for Computational Psychiatry and Ageing Research
[3]University of Cambridge

Model-free learning enables an agent to make better decisions based on previous experience, while holding only minimal knowledge about an environment's structure. It is generally assumed that model-free state representations are based on outcome-relevant features of the environment. We challenge this assumption, suggesting that the model-free system inevitably assign credit to task representations, even when those are completely irrelevant to the outcome. Here, we examined data from 765 individuals performing a two-step reward decision task where only stimulus-identity, but not spatial-motor aspects of the task predicted reward. We found that participants assigned value to spatial-motor representations despite it being outcome-irrelevant. Strikingly, spatial-motor value associations affected behavior across task-states and stages, suggesting the cognitive system assign credit to low-level task representations in a state-independent manner. Finally, the impact of spatial-motor value associations was attenuated in individuals with greater deployment of goal-directed (model-based) strategies. Taken together, these findings suggest that the cognitive system assign credit to low-level task representations regardless of the environment structure and highlight a need for a broader understating of how model-free representations are formed and regulated according to mental representations of the environment.

I-17. Photoreceptor synapses are tuned to environmental statistics in the larval zebrafish

Cornelius Schroder[1] CORNELIUS.SCHROEDER@UNI-TUEBINGEN.DE
Philipp Berens[1] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Takeshi Yoshimatsu[2] T.YOSHIMATSU@SUSSEX.AC.UK
Tom Baden[2] T.BADEN@SUSSEX.AC.UK

[1]University of Tubingen
[2]University of Sussex

It has long been postulated that the visual system is tuned to the statistics of the environment. Studying the retina of larval zebrafish, we here show that such an adaptation can already occur at the very first synapse of the visual system, in the photoreceptor. In particular, we found that UV-sensitive cones are non-homogeneously distributed across the retina, showing the highest density in the so-called strike zone (SZ), with which zebrafish larvae hunt their prey. We further performed two-photon imaging of Ca and glutamate release in UV cones using transgenic lines. In some of the experiments, we used double color imaging, allowing us to assess pre- and postsynaptic dynamics simultaneously. Interestingly, the synaptic dynamics resulted in higher information rates for transmitted information from UV cones in the SZ. To understand the mechanism underlying the environmental adaptation, we performed Bayesian inference for the parameters of a biophysical model of the photoreceptor synapse (ribbon synapse) based on the imaging data. To this end, we used a likelihood free inference approach, allowing us to compare different retinal regions in a principled manner. We found that the presynaptic calcium level and the size of vesicle pools varied between zones, yielding more efficient UV-cones in the SZ. We are currently testing these predictions experimentally. Our work contributes to the growing body of literature demonstrating how strikingly sensory systems are adapted to their natural environment or behavioral demands, already at the very first synapse.

I-18. Cortical-like dynamics in recurrent E-I networks optimized for fast probabilistic inference

Rodrigo Echeveste RODRIGO.ECHEVESTE@ENG.CAM.AC.UK
Guillaume Hennequin G.HENNEQUIN@ENG.CAM.AC.UK
Mate Lengyel M.LENGYEL@ENG.CAM.AC.UK

University of Cambridge

The dynamics of sensory cortices show a suite of basic, ubiquitous features that have so far escaped a common, principled theoretical account. These include strong, inhibition-dominated transients at stimulus onset, gamma oscillations, and noise variability - all stimulus-dependent. We present a unifying model in which all these dynamical phenomena emerge as a consequence of the efficient implementation of the same computational function: fast probabilistic inference. For this, we used a novel approach and trained a recurrent E/I neural circuit model of a V1 hypercolumn. The network was required to modulate not only the mean (as conventional) but also the variability of its stationary response distributions in order to match the corresponding input-dependent posterior distributions inferred by an ideal observer. The optimized neural circuit featured a number of remarkable computational and biological properties. Computationally, after training on a reduced stimulus set, it exhibited strong forms of generalization by producing near-optimal response distributions to novel inputs which required qualitatively different responses. Furthermore, the network discovered non-equilibrium dynamics, a state-of-the-art machine learning strategy to speed up inferences. The circuit also exhibited realistic biological properties for which it was not trained directly. It achieved divisive normalization and displayed marked transients at stimulus onset, as well as strong gamma oscillations, both scaling with stimulus contrast. Crucially, these dynamical phenomena did not emerge in a control network trained to match mean responses only (without modulating variability). Further analyses of transients and oscillations in the optimized network revealed distinct functional roles for them in speeding up inferences and made predictions that we confirmed in novel analyses of awake monkey V1 recordings. Our results offer a principled theoretical account of the basic motifs of cortical dynamics and predict further properties of these motifs that can be tested in future experiments.

I-19. Cortical credit assignment by Hebbian, neuromodulatory and inhibitory plasticity

Johnatan Aljadeff[1] J.ALJADEFF@IMPERIAL.AC.UK
Claudia Clopath[1] C.CLOPATH@IMPERIAL.AC.UK
James D'amour[2] JIM.D'AMOUR@NIH.GOV
Rachel Field[3] REF2129@COLUMBIA.EDU
Robert Froemke[4] ROBERT.FROEMKE@MED.NYU.EDU

[1]Imperial College London
[2]National Institutes of Health
[3]Columbia University
[4]New York University

The cortex solves the credit assignment problem. Synaptic weights adjust over time despite the lack of a clear relationship between strengths of individual synapses and the behavioural output. We propose that a combination of multiple mechanisms- Hebbian, acetylcholine (Ach) and noradrenaline (NE) dependent excitatory plasticity, together with inhibitory plasticity restoring detailed E/I balance, can effectively translate the partial error information available to each synapse into learning rules that solve the credit assignment problem. To show that, we studied a binary neuron learning random input-output associations (i.e., a simple perceptron), where plasticity of both excitatory and inhibitory synapses was subject to plasticity mechanisms taken from the experimental literature. Analysis of the model allowed us to find conditions on the plasticity rules which guarantee that the neuron can robustly learn a number of associations of the same order as the theoretical capacity of the perceptron. We used these conditions to make two key predictions regarding acetylcholine-dependent and inhibitory plasticity. We tested and confirmed the latter prediction by reanalysing a published in vitro dataset (D'amour & Froemke, 2015). Furthermore, we identified distinct functional roles for each of the plasticity mechanisms in our model. In line with previous literature (Yu & Dayan, 2005), Ach plays a role in making appropriate adjustments to weights when the output is close but not equal to the desired one; whereas NE plays a role in 'wholesale' changes to the output which is important when the learning process is 'stuck' far from the desired state. Importantly, our analysis also suggests a novel role for tight E/I balance: it reduces the dimensionality of the learning problem the neuron faces, allowing imperfect 'supervision' by the neuromodulatory systems to guide learning effectively.

I-20. Comparison of deep networks to biological hearing across many psychophysical and neural experiments

Alexander Kell ALEXKELL@MIT.EDU
Erica Shook ERSHOOK@MIT.EDU
Josh McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

Despite the recent renaissance of deep neural networks in neuroscience, the extent to which these networks are viable models of biological sensory systems remains unclear. Apparent discrepancies with biological perceptual systems (such as susceptibility to "adversarial" stimuli and the dependence on huge labeled training corpuses) suggest that contemporary deep neural networks are at best incomplete models of sensory systems. To probe potential similarities and differences with the human auditory system, we simulated a large set of psychophysical and fMRI experiments on a deep neural network trained to recognize speech and music. The network generalized in a human-like manner for several classes of unnatural speech manipulations (e.g., inharmonic speech). It failed to generalize like humans for a number of other speech alterations, most of which occur in everyday environments and thus are likely familiar to humans but were absent from the network's training data (e.g., whispering). However, in many of these cases human-like generalization could be largely achieved by simply learning a linear transformation of the network's features, indicating that the network had learned somewhat general and perceptually relevant representations. Additionally, we found that the averaged network layer activations reproduced the results of several fMRI experiments. The similarities between network and human suggest that deep neural networks can replicate aspects of human audition, and the discrepancies revealed here suggest targets for future modeling efforts.

I-21. Speed-information-performance tradeoffs in efficient sensory adaptation

Wiktor Mlynarski[1] WIKTOR.MLYNARSKI@IST.AC.AT
Ann Hermundstad[2] HERMUNDSTADA@JANELIA.HHMI.ORG

[1]Institute of Science and Technology
[2]HHMI Janelia Research Campus

Adaptation is a pervasive phenomenon across the nervous system. Sensory neurons, from periphery to cortex, adapt to changes in environment statistics. A dominant hypothesis is that adaptation enables a dynamic re-allocation of finite metabolic resources to support the performance of downstream tasks. During periods of adaptation, however, task performance and resource expenditure could cross levels critical for survival. To avoid such transient periods of vulnerability, adaptive sensory codes should balance task performance with the ability to detect and react to changes in the environment.

In this work, we propose a theoretical framework, based on Bayesian inference and rate-distortion theory, for analyzing and optimizing dynamics of sensory adaptation. We use this framework to study a model sensory system that operates in a dynamic environment and whose task is to accurately reconstruct incoming stimuli. The system encodes stimuli in stochastic spiking activity via an adaptable nonlinearity. We find that encodings optimized for task performance hinder the ability of the system to detect changes in the stimulus distribution. This results in a slow transient period of poor task performance when the environment changes. When optimized for inferring these changes, the system can adapt more quickly at the cost of degraded performance when the environment is stable. We derive encoding schemes that interpolate between these objectives, yielding diverse experimental predictions.

To demonstrate the generality of our approach, we apply it to a well-known model of power-law adaptation. For the first time, we show that power-law adaptation enables better information transmission and faster responses to changes in the stimulus distribution compared to exponential and perfect adaptation. The behavior holds for stimuli with 1/f statistics, but reverses for Gaussian stimuli. This application demonstrates that our approach can potentially reveal previously unappreciated aspects of adaptive neural coding.

I-22. Causal relationship between neural activity in orbitofrontal cortex and economic decisions

Weikang Shi[1] SHI.WEIKANG@WUSTL.EDU
Sebastien Ballesta[2] BALLESTA@UNISTRA.FR
Katherine E Conen[3] KCONEN@ORN.MPG.DE
Camillo Padoa-Schioppa[4] CAMILLO@WUSTL.EDU

[1]Washington University
[2]Centre de Primatologie de l'Universite de Strasbourg
[3]Max Planck Institute for Ornithology

[4]Washington University in St. Louis

Multiple lines of evidence link economic choice behavior to the orbitofrontal cortex (OFC). Most notably, studies in which monkeys chose between different juices showed that neurons in OFC encode the identities and subjective values of offered and chosen goods. This and corroborating results suggest that decisions are generated within this area. However, a causal relationship between neural activity in OFC and economic choices has not been established. In principle, such relationship is demonstrated if choices are predictably biased by electrical stimulation. Building on this concept, classic work in area MT found that low-current stimulation facilitates, while high-current stimulation disrupts perceptual decisions. One challenge in applying this concept to economic choices is the lack of columnar organization in OFC. In other words, neurons associated with two goods available for choice are physically intermixed in cortex. We developed two experimental paradigms to circumvent this challenge. In Exp.1, monkeys chose between two juices offered sequentially (pseudo-random order). High-current stimulation ($\geq 100\mu A$) delivered during offer1 or during offer2 consistently biased choices against the juice offered during stimulation. Furthermore, high-current stimulation during offer2 (when values were compared) increased choice variability (i.e., disrupted decisions). In Exp.2, we took advantage of the fact that offer value cells in OFC undergo range adaptation. Consequently, for any given juice, the increase in offer value induced by low-current stimulation should be proportional to the value range. If two juices are offered, low-current stimulation should increase both values, but ultimately bias choices in favor of the juice with the larger value range. In the experiments, monkeys chose between two juices offered simultaneously, and we delivered electric current ($50\mu A$). Confirming predictions, the stimulation induced a choice bias strongly correlated with the difference in value ranges. Taken together, our results demonstrate that economic decisions are causally related to neural activity in OFC.

I-23. Bayesian computation through cortical latent dynamics

Hansem Sohn[1] HANSEM@MIT.EDU
Devika Narain[1,2] DNARAIN@MIT.EDU
Nicolas Meirhaeghe[1] NMRGHE@MIT.EDU
Mehrdad Jazayeri[1] MJAZ@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Erasmus Medical Center

Statistical regularities in the environment create prior beliefs that we rely on to optimize our behavior when sensory information is uncertain. Bayesian theory formalizes how prior beliefs can be leveraged, and has had a major impact on models of perception, sensorimotor function, and cognition. However, it is not known how recurrent interactions among neurons mediate Bayesian integration. Using a time interval reproduction task in monkeys, we found that prior statistics warp the underlying structure of population activity in the frontal cortex, allowing the mapping of sensory inputs to motor outputs to be biased in accordance with Bayesian inference. Analysis of neural network models performing the task revealed that this warping was implemented through a low-dimensional curved manifold, which allowed us to further probe the potential causal underpinnings of this computational strategy. These results uncover a simple and general principle whereby prior beliefs exert their influence on behavior by sculpting latent cortical dynamics.

I-24. Multiscale calcium imaging of auditory cortex in awake marmoset monkeys

Xindong Song SONGXINDONG@GMAIL.COM
Xiaoqin Wang XIAOQIN.WANG@JHU.EDU
Yueqi Guo YGUO29@JHU.EDU

Johns Hopkins University

The common marmoset (Callithrix jacchus), a highly vocal New World monkey species, has emerged in recent years as a promising non-human primate model for neuroscience research. Two-photon imaging in the auditory cortex in marmosets is technically challenging since a mechanically vibrating laser scanner can generate sounds that are audible to the animal and thus interferes with the experimental design. In the current study, we developed a flexible, agile, yet silent two-photon microscope based on an acousto-optical deflector (AOD) scanner. An optical window was implanted over the auditory cortex of awake marmosets. A clear tonotopic structure can be observed through intrinsic signal imaging at both green and blue wavelengths. Multiple virus injections carrying GCaMP were made through the silicone-based optical window. A dual-virus strategy was used to separate controls over expression specificity and expression level. A clear macroscopic wide-field calcium response was observed starting 10 days after virus injections, at sound levels as low as 10 dB lower than hearing thresholds. The tonal response recorded from wide-field fluorescence imaging was consistent with intrinsic signal imaging results and was also limited mainly within the primary auditory cortex. By using complex sound stimuli, such as music or marmoset vocalization recordings, strong and widespread cortical responses ($\Delta F/F{>}10\%$) can be evoked in both primary and secondary cortical areas. The customized silent two-photon microscope was used to measure the responses of individual neurons at a microscopic scale. The general response patterns within each two-photon field-of-view were consistent with wide-field imaging results. However, individual neurons' responses can be heterogeneous even for close-by neurons. The multiscale calcium imaging approach reported here thus provides a new experimental paradigm for functional mapping of the marmoset auditory cortex in the awake condition in a high throughput way over conventionally electrophysiology methods.

I-25. Integration of external and internal steering signals by neurons that control heading in Drosophila

Sasha Rayshubskiy RAYSHUBSKIY@FAS.HARVARD.EDU
Rachel Wilson RACHEL_WILSON@HMS.HARVARD.EDU

Harvard University

Although locomotion per se does not require a brain, the brain is required for flexibility during locomotion. The brain allows locomotion to be guided by multimodal sensory guidance cues, and it integrates these sensory cues with computed latent state variables (e.g., heading direction representations). There are many studies of multimodal sensory integration and latent variables in the brain, but we know much less about how the outcomes of brain computations are actually communicated to the spinal cord during locomotion. To address this issue, we investigated the signals that the Drosophila brain sends to leg-control networks in the ventral nerve cord (the spinal cord analog) during walking behavior. Specifically, we focused on two anatomically-unique descending neurons ("upper motor neurons"), DNa01 and DNa02. Each of these cells is present in two copies in each brain, with one copy per hemisphere. We found these cells predict spontaneous and stimulus-evoked changes in heading direction, and unilateral activation of one cell is sufficient to bias heading. These cells respond to multimodal sensory cues, and their responses predict the side of the fly's turn, regardless of the side of sensory stimulation. In other words, these neurons operate in action-centric coordinates ("turn left" or "turn right"); they do not represent the side of the body where the sensory cue is localized. In addition, we found these cells participate in spontaneous transitions in behavioral arousal: they are hyperpolarized when the fly rests. Finally, we found these cells also predict the outcome of internally-generated, navigational motor commands that occur in the absence of sensory stimuli and result from rotation of heading direction representations in the brain via microstimulation. In short, these cells integrate multimodal sensory cues with internal steering drives to control locomotion, providing a link between higher brain functions and lower-level central pattern generators.

I-26. Dynamic n-back reinforcement learning representations in ventromedial prefrontal cortex

Rong Guo[1] RONG.GUO@TU-BERLIN.DE
Jan Glascher[2] GLAESCHER@UKE.DE
Felix Blankenburg[3] FELIX.BLANKENBURG@FU-BERLIN.DE
Klaus Obermayer[1] KLAUS.OBERMAYER@TU-BERLIN.DE

[1]Technische Universitat Berlin
[2]University Medical Center Hamburg-Eppendorf
[3]Freie Universitat Berlin

Neuroimaging research on statistical sequence learning has proceeded largely independently of the research on computational reinforcement learning because the focus is primarily on learning the temporal dependencies of environmental stimuli rather than on optimizing reward-based performance. Here, we investigated whether and to what extend computational and neural mechanisms of reinforcement learning can also be employed in statistical sequence learning. We developed an n-back Markov decision task and combined it with model-based fMRI analysis. 34 participants were instructed to predict the occurrence of a visual reward on either the left or the right side of the screen with the goal of maximizing their rewards. Reward appeared equally often on the left and right screen side (zero-order reward probabilities = 0.5) while the temporal dependences of reward were precisely controlled in two sets of conditions: (a) first-order conditions with 1-back conditional reward probability of 0.2 or 0.8 and (b) second-order conditions with 2-back conditional reward probability of 0.2 or 0.8. Crucially, the average rewards for the left and right options were equal but the reward received was conditioned on a specific temporal order of events, thus permitting us to disentangle the different orders of learning processes. We found that subjects were able to benefit from the n-back conditional probabilities for maximizing reward. We modeled the choice behavior with different orders of reinforcement learning models. We found a co-existence of 1-back and 2-back expected values in the ventromedial prefrontal cortex, suggesting that this region dynamically responds to the computational reinforcement-learning mechanism for estimating higher-order sequential dependencies in the same manner as it estimates average rewards. We further compared the model-based results with other computational strategies, i.e., Bayesian sequence learning and the win-stay-loss-shift strategy. In summary, our study highlighted the dynamic recruitment of an n-back reinforcement learning mechanism for guiding choice behavior.

I-27. Multi-area recordings reveal distinct functional ensembles spanning mouse visual cortex

Xiaoxuan Jia XIAOXUANJ@ALLENINSTITUTE.ORG
Joshua Siegle JOSHS@ALLENINSTITUTE.ORG
Severine Durand SEVERINED@ALLENINSTITUTE.ORG
Greggory Heller GREGGH@ALLENINSTITUTE.ORG
Tamina Ramirez TAMINAR@ALLENINSTITUTE.ORG
Yazan Billeh YAZANB@ALLENINSTITUTE.ORG
Shawn Olsen SHAWNO@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

The mammalian visual cortex is organized hierarchically, with feedforward, feedback, and lateral connections. While the anatomical structure is relatively fixed, functional networks must reconfigure rapidly in response to changes in external contexts and internal states. Therefore, understanding how information is propagated dynamically in a network in both feedforward and feedback directions is fundamental for deciphering the computational basis for intelligence.

Here, we studied this question by simultaneously measuring neurons across all layers of 6 visual cortical areas (VISp, VISl, VISrl, VISal, VISam, and VISpm). We developed a standardized platform that utilizes multiple Neuropixels probes to record spikes from an average of 597+/-32 neurons (n = 10 mice) distributed across these areas simultaneously. We computed directed connectivity matrices for all neurons using cross-correlograms (CCG) and Granger causality analysis. To explore network substructures, we clustered these connectivity matrices and found two distinct subnetworks: one ensemble was dominated by neurons that drive activity in the network, while another ensemble was more driven by the network. Interestingly, neurons in the 'driver' ensemble showed earlier and more transient responses compared to those in the 'driven' ensemble. 'Driver' neurons were also more temporally precise and showed higher synchrony within the ensemble. The spatial distribution of these two groups of neurons varied across cortical layers, with 'driver' neurons residing in the middle and superficial layers, while 'driven' neurons were biased toward deep layers. To explore the link between temporal dynamics and information propagation between areas, we built a multi-area rate model (constrained by the adjacency matrix derived from CCG analysis) that captured the measured neural dynamics. Our study is the first to measure single-neuron functional connectivity with high temporal resolution in the multi-area networks of mouse visual cortex. Our findings on neuronal dynamics and network information flow could have important implications for understanding cortical computation.

I-28. A novel approach to the empirical characterization of learning in biological systems

Sara A Solla[1] SOLLA@NORTHWESTERN.EDU
Yarden Cohen[2] YARDENC@BU.EDU
Predrag Cvitanovic[3] PREDRAG.CVITANOVIC@PHYSICS.GATECH.EDU

[1]Northwestern University
[2]Boston University
[3]Georgia Institute of Technology

Learning to execute precise yet complex motor actions through practice is a trait shared by most organisms. Here we develop a novel experimental approach for the comprehensive investigation and characterization of the learning dynamics of practiced motion. Our approach, based on the dynamical systems framework, considers a high-dimensional behavioral space in which a trial-by-trial sequence of motor outputs defines a trajectory that converges to a fixed point - the desired motor output. In this scenario, details of the internal dynamics and the trial-by-trial learning mechanism cannot be disentangled from behavioral noise for nonlinear systems or even well estimated for linear systems with many parameters. To overcome this problem, we introduce a novel approach: the sporadic application of systematic target perturbations that span the behavioral space and allow us to estimate the linearized dynamics in the vicinity of the fixed point. The steady-state Lyapunov equation then allows us to separately identify both the noise covariance and the linearized dynamics. We illustrate the method by analyzing sequence-generating neural networks with either intrinsic or extrinsic noise, at time resolutions that span from spike timing to spiking rates. We demonstrate the utility of our approach in experimentally plausible and realizable settings, and show that this method can fully characterize the linearized across trials learning dynamics as well as extract meaningful internal properties of the unknown mechanism that drives the motor output within each trial. We then illustrate how the approach can be extended to nonlinear learning dynamics through a flexible choice of the basis and the magnitude of perturbations.

I-29. Verbal free recall exhibits grid-like modulation of entorhinal theta in a 2D semantic-temporal space

Ethan Solomon ESOLO@PENNMEDICINE.UPENN.EDU
Nora Herweg NHERWEG@SAS.UPENN.EDU
Michael Kahana KAHANA@PSYCH.UPENN.EDU

University of Pennsylvania

The medial temporal lobe (MTL) of humans and animals supports navigation by forming neural representations of spatial layouts. Famously, individual neurons within the entorhinal cortex (EC) show specificity for the vertices of a hexagonal grid spanning the environment, and such grid-sensitive representations have recently been found in mesoscale analyses of entorhinal fMRI BOLD signal and theta power in the LFP. However, it remains an open question as to whether hexadirectional modulation of EC extends to explicitly non-spatial tasks, suggesting a fundamental, MTL-based neural mechanism for the exploration of any arbitrary representational space. Indeed, it is well-known that the MTL is crucial to the encoding and retrieval of episodic memory in general. Here we set out to ask whether verbal free recall is mechanistically similar to exploration of physical space. We (1) constructed 2D representations of "verbal space" using axes of temporal and semantic relations for words that occurred in 12-item lists, and (2) asked whether established neural signatures of spatial navigation – including hexadirectional modulation of EC theta – generalize to free exploration of these constructed verbal spaces. In a dataset of 252 neurosurgical patients with indwelling electrodes, we found that, as in recall of spatial information, the distance between recalled words in a temporal-semantic space was correlated with increases in pre-recall hippocampal theta. In a subset of 23 subjects with electrodes in the EC, we further found hexadirectionally-modulated theta power in high-performing subjects, indicating a correlation between theta and the direction of traversal through verbal space. The effect was specific to 6-fold symmetry in the theta band, and specific to the EC, mirroring longstanding findings from the spatial literature. This result provides powerful evidence that the MTL constructs arbitrary relational spaces to support memory in a general sense, converging on the same neural mechanisms to support spatial and verbal tasks.

I-30. Mapping the human navigation network with 7T-fMRI: A voxel-wise encoding model

Matthias Nau[1] MATTHIAS.NAU@NTNU.NO
Tobias Navarro Schroder[1] T.NAVARROSCHROEDER@NTNU.NO
Markus Frey[1] MARKUS.FREY@NTNU.NO
Christian F Doeller[2] DOELLER@CBS.MPG.DE

[1]Kavli Institute for Systems Neuroscience and Centre for Neural Computation (NTNU)
[2]Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig

The hippocampal formation derives cognitive maps from perceptual inputs, implicating visual stream and mediotemporal (MTL) regions in transforming sensory information into world-centered mnemonic codes. Inspired by the success of encoding models in capturing the tuning and topography of population activity underlying vision and working-memory, we develop voxel-wise and multivariate encoding models to study neural systems supporting spatial orienting during navigation. We used 7T-fMRI to monitor brain activity at submillimeter resolution in seventeen participants performing a spatial memory task, whilst navigating through a virtual environment (Fig. 1). We combined their virtual view direction (vVD) with a variety of circular-gaussian vVD-basis functions that differed in kernel spacing and width. For each parameter set, we then estimated model weights using an iterative ridge-regression training procedure that maximized predictability to find the optimal parameters explaining each voxel's time course. Using the resulting model weights, we then predicted held-out data to show that vVD not only predicts activity in visual regions involved in self-motion- and scene processing, but also in MTL regions maintaining environment-anchored spatial representations, like the entorhinal cortex (EC) (Fig. 2). Notably, each region was best predicted with distinct vVD-tuning width that increased anteriorly from retrosplenial and parahip-

pocampal to entorhinal cortices. Further, we explored inverted encoding modeling in this context and successfully reconstructed vVD also from regions with high-level mnemonic function, like the hippocampus, and revealed the functional topography of vVD-tuning within the EC. By leveraging virtual reality, high-resolution 7T-fMRI and encoding modeling, we characterize and map the intricate view-tuning of visual stream and MTL regions that contribute to our sense of head-direction. Our approach not only elucidates neural computations supporting complex naturalistic behaviors such as navigation, but can also be extended to new variables, techniques and species to study population tuning on many levels also in other cognitive domains.

I-31. Neural mechanisms underlying economic decisions under sequential offers

Sebastien Ballesta[1] SEBASTIEN.BALLESTA@GMAIL.COM
Camillo Padoa-Schioppa[2] CAMILLO@WUSTL.EDU

[1]University of Strasbourg
[2]Washington University in St. Louis

Economic choices involve two mental stages – values are assigned to the available options and a decision is made by comparing values. Evidence from lesion studies, functional imaging, neurophysiology and computational modeling suggests that economic decisions between goods are formed within the orbitofrontal cortex (OFC). Importantly, current notions emerged almost exclusively from studies where two offers were presented simultaneously. Yet, in many real-life decisions, offers available for choice appear or are examined sequentially. Previous work on choices under sequential offers suggested that choices in the two modalities (sequential and simultaneous offers) rely on fundamentally different mechanisms. Unfortunately, discrepancies in data analyses make it difficult to compare findings from different lines of work. To shed light on this critical issue, we recorded from the OFC of rhesus monkeys choosing between two juices offered sequentially and in variable amounts. An analysis of neuronal responses across time windows revealed the presence of different groups of neurons encoding the value of individual juices, the binary choice outcome and the chosen value. These groups of cells closely resemble those identified under simultaneous offers 1, suggesting that decisions in the two modalities are formed in the same neural circuit. We then examined four hypotheses on the mechanisms underlying value comparison (i.e., the decision). Our data clearly confuted the notion of a single neuronal pool. They were also inconsistent with mutual inhibition between pools of offer value cells. Interestingly, we did not find any sustained representation of the first offer value in OFC. Instead, decisions appeared to involve mechanisms of circuit inhibition whereby each offer value inhibited the population of neurons representing the opposite choice outcome. Our results reconcile seemingly divergent findings on the neural mechanisms underlying good-based decisions. They also shed light onto the role played by shifts of attention or mental focus in economic choices.

I-32. A recurrent network model for the neural mechanism of speed, accuracy and urgency signal in decision making

Tianming Yang[1] ZHZHEWEI@LIVE.COM
Zhewei Zhang[1] ZWZHANG@ION.AC.CN
Huzi Cheng[2] HUZCHENG@IU.EDU

[1]Chinese Academy of Sciences
[2]Indiana University Bloomington

Decision making is often a process that takes consideration of both accuracy and speed. Currently, a great number of popular models to explain the accuracy and speed in decision making are based on variations of drift diffusion models (DDM), in which one accumulates evidence toward decision bounds. Attractor-based recurrent neural networks have been proposed to explain the underlying neural mechanism, which may account for many

experimental findings. Yet, some key predictions are inconsistent with experimental observations. Here, we propose an alternative neural network modeling approach. Our network model is trained to learn the statistical structure of temporal sequences of sensory events, action events, and reward events. We demonstrate its learning with a reaction version of the weather prediction task previously studied in monkey experiments, in which both the animals' behavior and the neuronal responses were consistent with the DDM. The network model's performance is able to reflect the accuracy and reaction time pattern of the animals' choice behavior. Crucially, the analyses of the unit responses in the network reveal that they match important experimental findings. Notably, we find the units encoding the accumulated evidence and the urgency signal. We further identify two groups of units based on their connection weights to the choice output units. Simulated lesions of the output connections of each group of units produce double-dissociated effects on the network's choice and reaction time behavior. Graph analyses reveal that these two groups of units belong to one highly inter-connected sub-network. Our work shows that a network model that is simply trained to learn the statistical structure of temporal event sequences may reflect many experimental findings of choice and urgency signal in the cortex and basal ganglia. It has interesting theoretical implications for the understanding of the computational principles in the brain.

I-33. Dissociation between postrhinal cortex and other parahippocampal regions in the representation of egocentric boundaries

Xenia Gofman[1] XEN0712@GMAIL.COM
Gilad Tocker[2] GILADTOCKER@GMAIL.COM
Shahaf Weiss[3] SHAHAF.ME@GMAIL.COM
Charlotte N Boccara[4] C.BOCCARA@GMAIL.COM
Li Lu[5] LI.LU@BCM.EDU
May-Britt Moser[6] MAY-BRITT.MOSER@NTNU.NO
Edvard I Moser[6] EDVARD.MOSER@NTNU.NO
Dori Derdikman[1] DERDIK@TECHNION.AC.IL

[1]Technion
[2]Bar Ilan University, Technion
[3]Max Planck Institute for Brain Research
[4]King's College London
[5]Norwegian University of Science and Technology
[6]Kavli Institute for Systems Neuroscience and Centre for Neural Computation (NTNU)

Navigation requires the integration of many sensory inputs to form a multi-modal, cognitive map of the environment, which is believed to reside in the hippocampal region via spatially tuned cells. While the cognitive map is represented in allocentric coordinates, the environment, is sensed through diverse sensory organs, most of them situated in the animal's head, and therefore, perceived in head-centered egocentric coordinates. Here we show a new functional type of cell, egocentric boundary cell (EBC) recorded mainly in the postrhinal cortex (POR). EBCs receptive fields were computed by “egocentric wall rate maps”, in which the firing rate of the cell was calculated as a function of the distance of the rat from the wall in different egocentric angles. Significant EBCs were found to encode the boundaries of the environment in an egocentric reference frame. We checked the significantly of EBCs by two different and unbiased measures: 1) shuffling procedure- cells that passed the egocentric angular tuning criteria were classified as EBCs. 2) a generalized linear model (GLM). Both analyses revealed that the majority of 'pure' EBCs were in the POR, while cells with conjunction representation of EBC and Head direction (EBC x HD) were mostly found in our data in the parasubiculum (PaS) and medial entorhinal cortex (MEC). Previous theoretical models have suggested that HD cells can participate in a generation of an egocentric-allocentric transformation. Following this, we suggest that egocentric sensory input, entering the system, feeds into EBCs in the POR, then, combined with the HD signal to form a conjunctive representation in EBC x HD cells in the PaS and the MEC. Our observation suggests a new role of HD cells: they are not only used as an internal compass but also play a pivotal role in the transformation of egocentric signals into the allocentric cognitive map.

I-34. The interaction between decision-making and attention signals in the lateral intraparietal area

Chechang Nie CCNIE@ION.AC.CN
Yilin Chen YILINCHEN@ION.AC.CN
Tianming Yang TYANG@ION.AC.CN

Chinese Academy of Sciences

The lateral intraparietal area (LIP) has been shown to play an important role in visuo-motor information transformation. It has been proposed that there is a representation of the salience map in the LIP, where the neural activity indicates the behavior relevance of objects in the visual space. On the other hand, many studies have shown that the LIP activity reflects the process of evidence accumulation during decision-making related to eye movements. Although the two roles that the LIP has been indicated to play are potentially compatible, there have been no studies so far to check whether they really are. Here, we studied the representation of decision variables in the LIP to investigate if manipulations of bottom-up attention would affect decision making. We trained a macaque monkey to perform a random dot motion direction discrimination task. In some trials during the motion presentation period, we randomly flashed one of the choice targets, generating a bottom-up attention signal that elicited responses in the LIP. The visual perturbation was not associated with the direction discrimination and the monkey was supposed to ignore it. By measuring the behavior effects of the perturbations and recording single unit activities from the LIP, we found that the perturbation biased the monkey's choice toward the flashed target and the LIP neurons encoded the location of the perturbation transiently. Further analyses revealed how the bottom-up attention signal affected the representation of decision making process in the LIP. The findings provide important insights on the roles of the LIP in visuomotor transformation and help us to understand the neural circuitry underlying decision-making.

I-35. Can neuromodulation explain why slow waves have different effects on memory in sleep and anesthesia?

Trang-Anh Nghiem[1] TRANGANH.NGHIEM@GMAIL.COM
Nuria Tort-Colet[2] NURIA.TORT-COLET@UNIC.CNRS-GIF.FR
Tomasz Gorski[2] TOMASZ.GORSKI.14@GMAIL.COM
Ulisse Ferrari[3] ULISSE.FERRARI@GMAIL.COM
Shayan Moghimyfiroozabad[2] SHAYANMOGHIMY@GMAIL.COM
Jennifer Goldman[2] JENNIFER.GOLDMAN.MCGILL@GMAIL.COM
Bartosz Telenczuk[2] TELENCZUK@UNIC.CNRS-GIF.FR
Cristiano Capone[4] CRISTIANO0CAPONE@GMAIL.COM
Thierry Bal[2] BAL@UNIC.CNRS-GIF.FR
Matteo di Volo[2] M.DIVOLO@GMAIL.COM
Alain Destexhe[2] DESTEXHE@UNIC.CNRS-GIF.FR

[1]Ecole Normale Superieure
[2]Centre National de la Recherche Scientifique
[3]Sorbonne Universite
[4]Istituto Nazionale di Fisica Nucleare

During both anesthesia and sleep, neurons display collective activity characterized by slow ($<$1 Hz) oscillations, emerging from the alternation between transients of high activity (UP states) and transients of near silence (DOWN states). However, memory is reinforced exclusively in sleep. Here, we uncover distinct slow oscillatory temporal patterns in sleep versus anesthesia: in sleep, long DOWN states are systematically followed by short UP states, unlike in anesthesia. This difference was found to be robust across species, brain regions, and anesthetics. A spiking neuron network model suggests a possible mechanism based on depressed acetylcholine levels under anesthesia, via an interplay between noise and spike-frequency adaptation. The predictions are confirmed by

cholinergic neuromodulation *in vitro*, inducing a switch between the two types of slow waves. The results show that strong adaptation found in anesthesia, but not in sleep, can produce a different type of slow oscillations, less sensitive to external inputs, that may prevent memory formation. As acetylcholine levels are also affected during Alzheimer's disease, we predict that slow oscillations during sleep in Alzheimer's patients become anesthesia-type. This can explain both sleep and memory symptoms and guide future treatments.

I-36. Minimal-dimensionality implementation of behavioral tasks in recurrent networks

Alexis Dubreuil[1] ALEXIS.DUBREUIL@ENS.FR
Francesca Mastrogiuseppe[2] FRAN.MASTROGIUSEPPE@GMAIL.COM
Srdjan Ostojic[1] SRDJAN.OSTOJIC@ENS.FR

[1]Ecole Normale Superieure
[2]Gatsby Computational Neuroscience Unit, University College London

Population recordings have revealed that cortical activity lies in spaces of low dimensionality when animals perform a variety of behavioral tasks. The ubiquity of this observation raises outstanding questions on the relationship between the dimensionality of neural dynamics and the computations they implement. In particular, what is the minimal possible dimensionality of the activity in a network performing a given task? Do low-dimensional implementations of a task have a computational advantage with respect to higher-dimensional ones? To address these questions, we exploit the recently introduced framework of low-rank recurrent networks [Mastrogiuseppe & Ostojic 2018]. In these networks, low-rank connectivity matrices generate low-dimensional dynamics at the level of the network population. We develop a new variation of the backpropagation algorithm which allows us to train recurrent networks on stereotypical tasks by constraining the rank of the final connectivity matrix. We specifically focus on the network implementations characterized by minimum rank and we use a mean-field approach to reverse-engineer them, thus connecting the low-rank connectivity patterns with the network dynamics implementing the task. We applied this approach to a variety of systems neuroscience tasks which involve different cognitive components: working memory, sensory integration and context-dependent switching. For each task, the rank of the minimal implementation obtained via training provides an effective quantification of its complexity. We use our theoretical framework to identify how different task components are implemented in trained networks, unveiling simple 'computational primitives' defined in terms of low-rank connectivity components. We show that these primitives can be flexibly combined together to construct network models which implement more complex or multiple tasks. Our findings suggest that low-rank implementations may allow a network to maximize the number of tasks it can perform.

I-37. Switching network synchronization not via connectivity but single-cell properties

Susanne Schreiber[1] S.SCHREIBER@HU-BERLIN.DE
Susana Andrea Contreras[1] SUSANA.CONTRERAS@HU-BERLIN.DE
Janina Hesse[1] JANINA.HESSE@BCCN-BERLIN.DE
Allan Gulledge[2] ALLAN.GULLEDGE@DARTMOUTH.EDU
Nikolaus Maier[3] NIKOLAUS.MAIER@CHARITE.DE
Dietmar Schmitz[3] DIETMAR.SCHMITZ@CHARITE.DE
Jan-Hendrik Schleimer[4] JH.SCHLEIMER@HU-BERLIN.DE

[1]Humboldt-University Berlin
[2]Dartmouth College
[3]Charite - Universitaetsmedizin Berlin
[4]Humboldt-University Berlin, Germany

Synchronous activity of neuronal networks characterizes brain states in health and disease. Connectivity within the network is considered the key player for rhythm generation. Yet, topology cannot explain how small parameter changes affecting cellular voltage dynamics, like a rise in extracellular potassium, can evoke strongly synchronized states as they are observed during epileptic seizures. Based on mathematical methods, we here uncover a novel mechanism that boosts network synchronization via an exclusive alteration of neuron-intrinsic properties. The mechanism relies on a critical transition in cellular dynamics that has so far been overlooked for its potential to induce drastic changes in synchronization: spike generation via the saddle homoclinic orbit bifurcation. Interestingly, when transitioning into homoclinic spike generation, small changes in biophysical parameters suffice to drastically increase synchronization – to a much stronger extent than if spike generation was induced via the typically considered type I (saddle-node-on-invariant-cycle) or type II (Hopf) bifurcations. We find that extracellular potassium accumulation, temperature, capacitance, or leak can induce the critical transition in single-cell dynamics and indeed boosts network synchronization. Interestingly, a number of these parameters have been implicated in the generation of epileptic seizures and the critical transition here presents itself as a novel candidate mechanism for seizure induction, for example in response to high extracellular potassium concentrations. We here predict hallmarks of the synchronizing transition and provide in-vitro evidence that homoclinic spiking is indeed observed in the brain. Taken together, the mechanism relies on cellular properties but unfolds its functional consequences at the macroscopic network level – bridging from cells to networks. It unveils a generic principle that provides an attractive explanation for physiological and pathological switches between synchronous and asynchronous states. These are produced by state-dependent cellular parameters in the brain without requiring changes in connection strengths.

I-38. Population codes of prior knowledge learned through environmental regularities

Silvan Quax S.C.QUAX@GMAIL.COM
Sander Bosch S.BOSCH@DONDERS.RU.NL
Marcel van Gerven M.VANGERVEN@DONDERS.RU.NL

Radboud University

The process of perception is faced with the difficult problem of inferring the underlying cause of the perceptual input that reaches the senses. An efficient neural machinery is required to perform this inference despite the noise and ambiguity in our observations. An important element in this inference process is the prior knowledge about the occurrence of certain causes that is acquired over time. Humans are able to learn and incorporate this kind of knowledge in an efficient, Bayes optimal, way in many situations. How the brain learns to represent this prior knowledge and to incorporate it in its sensory inference, is still an open question. A tempting hypothesis is that the brain learns to represent prior probabilities about the environment through the frequency with which an object occurs. Here, we simulated some commonly used tasks in the experimental literature where the prior probability of observing a stimulus from one class could be different from the probability of observing a stimulus from another class. We show that a model population of neurons learns to correctly represent and incorporate prior knowledge, by only receiving feedback about the accuracy of their inference and without probabilistic feedback. When manipulating both input noise and prior probabilities we found that, while neurons fire sparser when there is less input noise, they fire less sparse when a certain class is more likely. This counter-intuitive finding could be well explained by a template matching code, where less noise makes it less likely that a neurons receptive field is accidentally activated, while a higher prior probability biases the neurons activity such that it is activated more easily. These results pose some clear and testable predictions about the neural representations of prior knowledge in commonly used experimental paradigms that can guide new experiments.

I-39. Top-down feedback controls subjective contour responses in mouse visual cortex

Alexandr Pak PAK6@PURDUE.EDU
Esther Ryu ERYU@PURDUE.EDU
Alexander Chubykin CHUBYKIN@PURDUE.EDU
Purdue University

Our perception of the world is not always aligned with an actual sensory input from the environment. One of the famous examples is Kanizsa figures with illusory contours. Our visual system recognizes the subjective figure despite the lack of physical bases for them such as color or luminance gradients or lines forming the figure. These illusory shapes offer a unique opportunity to gain insight into the neural basis of the subjective nature of animal perception. Prior studies in primates have shown that when the subjective contour intersects the neuron's receptive field (RF) in the primary visual cortex (V1), it induces a similar response to a real line. Furthermore, the secondary visual area V2 had an earlier response to the illusory contour compared to V1. These observations were intriguing given that no bottom-up input was present. They led to the hypothesis that higher visual areas might supply the missing information via top-down feedback. However, no direct neurophysiological evidence has been demonstrated. Here, we used in vivo extracellular silicon probe recordings along with optogenetics to study neural dynamics of subjective contours in awake mice. We show, for the first time to our knowledge, that mouse V1 responds to subjective contours, Kanizsa figures both up- and downregulate responses of distinct neuronal groups, and inactivation of the top-down feedback from lateromedial (LM) area suppresses subjective contour responses in V1. These results suggest that one of the ways top-down feedback influences sensory processing at lower levels is by filling-in the missing information. Our findings contribute to the emerging view of the perception being the result of an interplay between feedforward and feedback activity.

I-40. Panoptic recording of cross-cortical neuronal dynamics during working memory

Isaac Kauvar IKAUVAR@GMAIL.COM
Timothy Machado TAMACHADO@STANFORD.EDU
Elle Yuen EYUEN1@STANFORD.EDU
John Kochalka KOCHALKA@STANFORD.EDU
William Allen WALLEN@STANFORD.EDU
Cephra Raja CPRAJA@STANFORD.EDU
Nandini Pichamoorthy NANDINIP@STANFORD.EDU
Gordon Wetzstein GORDON.WETZSTEIN@STANFORD.EDU
Karl Deisseroth DEISSERO@STANFORD.EDU
Stanford University

Mammalian neocortex plays a central role in cognition. Dynamics evolve over interconnected networks of neurons spanning centimeters in space–at millisecond temporal scales. To determine how information is shared and propagated across the cortical network, it is essential to observe neuronal activity across many cortical regions at high speed. Unfortunately, this is not feasible with existing approaches. Here, we present an optical system for simultaneously recording from over a thousand neural sources across the entirety of mouse dorsal cortex at over 25 Hz. Our approach, Cortical Observation by Synchronous Multi-focal Optical Sampling (COSMOS), uses a novel single-sensor multi-focal widefield macroscope to measure Ca2+ dynamics from neurons spanning the curved cortical surface with high signal-to-noise. By quantifying the orientation tuning of neuronal sources in primary visual cortex, we provide evidence that a subset of sources are single neurons while the rest are mixtures of a few neurons. We applied this system to record from mice performing a three-spout olfactory decision-making behavior. From our neural data, we are able to decode individual licks to different spouts and demonstrate that moment-by-moment information about behavioral experience is distributed across dorsal cortex. Surprisingly, we

also discover a widespread representation of working memory. The spout to which the mouse is planning to lick can be predicted during the variable length intertrial interval, even when no motor action is visible, and without any cue indicating to the mouse where it should lick. These results suggest a model of cortical computation as concurrently widespread, with shared information distributed across multiple regions in an overlapping but not entirely redundant manner. In sum, COSMOS provides a unique means of dissecting the global nature of the cortical computations underlying cognition and behavior.

I-41. Attention suppresses spontaneous activity of fast-spiking in visual cortex

Elizabeth Sachse ESACHSE@U.ROCHESTER.EDU
Adam Snyder ADAM.SNYDER@ROCHESTER.EDU
University of Rochester

Multi-unit recordings have helped identify different features of neuronal populations and how they are modulated by attention. Such cellular and circuitry features are important clues to the mechanisms of spatial attention. Information coding in V4 is influenced by complex organization in cortical columns that contain several cell types. Fast-spiking parvalbumin interneurons make up the majority of inhibitory neurons in the cortex and have dense connections to local excitatory neurons. Inactivation of these neurons in mouse V1 leads to decreased correlation among excitatory cortical neurons both during spontaneous and stimulus-evoked activity, indicating these inhibitory neurons likely act to regulate activity of local excitatory populations across space and time through changes in neural synchrony [1]. For example, low synchrony may be beneficial during spontaneous activity, as it indicates reduced dependence on internal signals and greater receptivity to stimulation. On the other hand, higher synchrony may help effectively transmit information downstream during stimulus processing. To investigate this, we compared attention dynamics of V4 fast-spiking and regular-spiking neurons by recording population activity while monkeys performed an attention task. We measured dynamics of attention modulation for each neuron during spontaneous activity and stimulus-processing and classified neurons as fast-spiking (FS) or regular-spiking (RS) based on their action potential waveforms. We found that patterns of attention dynamics co-segregated with cell classes. This suggests that FS and RS cell classes play distinct functional roles in the dynamic allocation of attention across sensory stimulation contexts.

I-42. A Bayesian approach to unsupervised learning in the olfactory system

Naoki Hiratani N.HIRATANI@UCL.AC.UK
Peter Latham PEL@GATSBY.UCL.AC.UK
Gatsby Computational Neuroscience Unit, University College London

One of the most important questions in neuroscience is: how can animals make sense of the outside world without explicit feedback, and using only local, biologically plausible, learning rules? We address this question in the context of olfaction, where it is known that animals can learn novel odors based on only a few presentations. A candidate mechanism for this data-efficient learning is Bayesian learning, in which learning is treated as an inference problem. Recent studies suggest that single neurons can implement approximate Bayesian learning by plausible synaptic plasticity mechanisms given some supervising signals [Aitchitson, Pouget & Latham, arXiv, 2017; Hiratani & Fukai, PNAS, 2018]. Here we extend these approaches to unsupervised learning by networks of neurons rather than supervised learning by single neurons. We formulate olfactory learning as an integrated Bayesian inference problem, then derive a set of synaptic plasticity rules and neural dynamics that enables near-optimal learning of near-optimal odor identification. The proposed model enables faster and more scalable learning of olfactory stimuli compared to previous approaches. Moreover, our theory reveals the theoretical basis of the nonlinear f-I curve, and the functional role of changes in neuronal excitability and learning rate during development.

The model also predicts a non-trivial relationship between the learning rate and the lifetime sparseness of the postsynaptic firing activity. We extend the framework to reward-based olfactory learning, and show that the circuit is able to rapidly learn odor-reward association under a plausible neural architecture. Although we have focused on olfaction, our approach can be applied to other sensory modalities as well.

I-43. Cell-type specific modulation of information transfer by dopamine

Niccolo Calcini N.CALCINI@NEUROPHYSIOLOGY.NL
Ate Bijlsma A.BIJLSMA@NEUROHYSIOLOGY.NL
Yiping Zhang Y.ZHANG@NEUROPHYSIOLOGY.NL
Angelica da Silva Lantyer A.LANTYER@NEUROPHYSIOLOGY.NL
Koen Kole K.KOLE@NEUROPHYSIOLOGY.NL
Tansu Celikel CELIKEL@NEUROPHYSIOLOGY.NL
Fleur Zeldenrust F.ZELDENRUST@NEUROPHYSIOLOGY.NL

Radboud University

Neurons in the sensory cortices form representations of the world in a cell-type specific manner. The responses of regular spiking (pyramidal) cells and fast-spiking (inhibitory) neurons in cortical networks are primarily shaped by each neuron's place in the network (connectivity of the network) and its biophysical properties (ion channel expression and density, Azarfar et al. 2018), and are subject to top-down regulatory control performed by neuromodulatory neurotransmitters.

We studied the effects of D1 dopamine receptor (D1R) activation on neural computations of regular and fast-spiking neurons in L2/3 of barrel cortex. Using a recently developed ex-vivo method (Zeldenrust et al. 2017), we estimated the information transfer and fitted computational models to a large number of in vitro recordings (Lantyer et al., 2018). D1R activation resulted in cell-type specific regulation: while it hyperpolarized the spike-threshold and increased the firing rate in fast-spiking neurons, it inhibited stimulus-evoked spiking in regular spiking neurons. This resulted in increased information transfer in fast-spiking but not regular spiking neurons. These differences in neural responses were accompanied by faster decision-making during tactile object localization in freely behaving mice (gap-crossing task, (Celikel & Sakmann, 2007)). Tonic D1R activation during task execution promoted faster decision making with less sensory evidence. We hypothesize that D1R modulation improves the integration of tactile evidence to reach a perceptual decision by modulating the cell-type specific gain modulation to enhance fast-spiking neurons' contribution to sensory coding.

I-44. Cognitive strategies in a color matching puzzle for dogs

Philip Meier PHILIP@CLEVER.PET
Leo Trottier LEO@CLEVER.PET
Dan Knudsen DAN@CLEVER.PET

CleverPet, Inc.

What motivates behavioral change when reward payouts are identical? We propose that reinforcement learning in biological organisms naturally transitions to lower-dimensional behavioral strategies for repeated tasks that have multiple paths to a successful outcome. Here we analyze a novel symmetric color-matching task performed by dogs. This study leverages data generated by a consumer electronics product (the CleverPet Hub^{TM}) that engages dogs with puzzle games via operant conditioning. We analyze a task that requires dogs to match the colors displayed on three LED touchpads (e.g. so that they are either all yellow, all blue, or all white). The game is similar to a low-dimensional Rubik's cube. Repeatedly performing the same action will cycle back to an earlier puzzle state. The topology of the game allows alternative sequences of actions to yield equally valid solutions. While a Rubik's cube has $> 10^{19}$ states and one solution, the three-color matching puzzle has 24 unsolved states and three solutions. Dogs are trained by rewarding solutions that require progressively fewer actions. We find that high-performing dogs do not converge on an "optimal effort" solution (e.g. the least number of steps to get a reward), but do find sufficient local minima that yield rewards by respecting the reinforced 6-action limit. Notably, the solutions discovered involve a simpler learning rule that generalizes across puzzle starting states, such as "avoid pressing blue." We analyze the strategies employed by different dogs and find multiple attractors, such that some dogs constantly "avoid blue" and others "avoid yellow." For each initial state of the puzzle, behavioral entropy reduces over time, even when the fraction of error trials does not go down.

I-45. Characterizing V1 neurons using convolutional neural networks and project pursuit models with interpretable sparse coding kernels

Ziniu Wu[1] ZINIUW@CS.CMU.EDU
Yimeng Zhang[1] YIMENGZH@ANDREW.CMU.EDU
Shiming Tang[2] TANGSHM@PKU.EDU.CN
Tai Sing Lee[1] TAISLEE@ANDREW.CMU.EDU

[1]Carnegie Mellon University
[2]Peking University

Recent statistical modeling techniques, such as pursuit models and convolutional neural network, intended to recover visual neurons' receptive fields or preferred features by fitting neuron's responses to a large number of images. These models have decent performance in terms of prediction but tend to yield kernels with fairly noisy and uninterpretable features. This might due to the fact that the problem is under-constrained, and there are many possible local minima solutions. We hypothesize that features learned by unsupervised learning based on sparse coding principle (Olshausen and Field, 1996) might provide more interpretable kernels to build projection pursuit or CNN models for predicting neurons' responses. Our experiments show that this approach yielded more interpretable feature kernels and at the same time produced better prediction performance than CNN models (Zhang et al. 2018, J. Computational Neuroscience) with the same number or more parameters. Our experiments also suggest that our model is more robust against noise than the CNN models. Having a set of interpretable feature detectors provide a new approach for us to model and classify V1 neurons with complex tunings (Tang et al. 2017, Current Biology) based on their feature preference.

I-46. Hierarchical motion structure is employed by humans during visual perception

Johannes Bill JOHANNES_BILL@HMS.HARVARD.EDU
Hrag Pailian PAILIAN@FAS.HARVARD.EDU
Samuel Gershman GERSHMAN@FAS.HARVARD.EDU
Jan Drugowitsch JAN_DRUGOWITSCH@HMS.HARVARD.EDU

Harvard University

Making sense of the hierarchical arrangement of form and motion is central to visual scene perception. For example, while driving, other vehicles' locations must be anticipated from the traffic flow even if they are temporarily occluded. Despite its ubiquity in everyday reasoning, surprisingly little is known about how exactly humans and animals employ motion structure knowledge when perceiving dynamic scenes. To investigate this question, we propose a formal framework for characterizing structured motion and generating structured motion-stimuli, which supports a wide range of hierarchically arranged real-world motion relations among stimulus features. A key benefit is that the joint distribution of generated stimulus trajectories is analytically tractable, which allowed us to compare human performance to ideal observers. To do so, we first introduced structured motion in the well-established multiple object tracking task. We found that humans performed better in conditions with struc-

tured than independent object motion, indicating that they benefitted from structured motion. A Bayesian observer model furthermore revealed that the observed performance gain is not due to the stimulus itself becoming simpler, but due to active use of motion structure knowledge during inference. A second experiment, in which trajectories of occluded objects had to be predicted from the remaining visible objects, provided a fine-grained insight into which exact structure human predictions relied on in the face of uncertainty: Bayesian model comparison suggests that humans employed the correct or close-to-correct motion structure, even for deep motion hierarchies. Overall, we demonstrated – to our knowledge – for the first time that humans can make use of hierarchical motion structure when perceiving dynamic scenes, and flexibly employ close-to-optimal motion priors. Our proposed formal framework is compatible with existing neural network models of visual tracking, and can thus facilitate theory-driven designs of electrophysiology experiments on motion representation along the visual pathway.

I-47. Breathing coordinates network dynamics underlying memory consolidation

Nikolaos Karalis[1] NIKOLASKARALIS@GMAIL.COM
Anton Sirota[2] SIROTA@BIOLOGIE.UNI-MUENCHEN.DE

[1]Friedrich Miescher Institute for Biomedical Research
[2]Ludwig-Maximilians University Munich

During offline states, cortical and hippocampal networks are dominated by nonlinear dynamics, such as hippocampal ripples and neocortical DOWN/UP states. The coordination of these network dynamics underlies memory reactivation and information transfer between these regions and is believed to be critical for the consolidation of memories, giving rise to prominent models of systems memory consolidation. However, the mechanisms underlying such persistent coordination of the dynamics across global brain networks are poorly understood and the existing frameworks based on pair-wise interactions cannot explain the large-scale brain synchrony during offline states. From a theoretical perspective, a global pacemaker has been postulated as an effective solution to the coupling of distinct nonlinear network dynamics, but the neural implementation of such a mechanism remains elusive. Here we address the hypothesis that respiration acts as an oscillatory master clock, persistently coupling distributed brain circuit dynamics. Using large-scale recordings from the prefrontal cortex, hippocampus, amygdala, nucleus accumbens, and thalamus in behaving mice, we demonstrate that respiration entrains the neuronal activity in all these regions, giving rise to global synchronization. In parallel, it mediates the interaction of local nonlinear network dynamics, including hippocampal ripples and cortical DOWN/UP states, as well as gamma oscillations, during quiescence and sleep, effectively providing the substrate for coherent systems memory consolidation across distributed brain structures. Further, using pharmacological manipulations and analytical methods, we identify a novel joint circuit mechanism underlying the respiratory entrainment of the limbic circuits, in the form of an intracerebral respiratory corollary discharge (RCD) and a respiratory olfactory reafference (ROR), suggesting a distributed predictive signaling mechanism across cortical and subcortical networks. These results highlight breathing, a perennial brain rhythm, as an oscillatory scaffold for the functional coordination of the limbic circuit, enabling the segregation and integration of information flow across neuronal networks and coordinating memory consolidation processes.

I-48. Biologically plausible mechanism for noise resilience in continuous parameter working memory

Nimrod Shaham[1] NSHAHAM@FAS.HARVARD.EDU
Alon Salhov[2] ALON.SALHOV@MAIL.HUJI.AC.IL
Yoram Burak[3] YORAM.BURAK@ELSC.HUJI.AC.IL

[1]Harvard University
[2]Hebrew university of Jerusalem

[3]Hebrew University of Jerusalem

How does the brain store information in working memory is one of the most fundamental questions in neuroscience. Specifically, the ability of neural circuits in the brain to store values of a continuous parameter in working memory raises intriguing theoretical questions, which are highlighted when considering continuous attractor models. These models are known to be sensitive to mistuning and to noisy dynamics, which can drive diffusion along the attractor. In this work we show that when taking into account the fact that different feedback loops in the brain have different time scales, we can get automatically a reduction of the noise driven diffusivity along the attractor. This is inspired by the work of Lim and Goldman (2013), where the authors showed that negative derivative feedback can reduce sensitivity to some forms of parameter mistuning. Here we study a continuous attractor neural network model with several different synaptic time scales, and noisy spiking neurons. We calculate the diffusion coefficient along the attractor explicitly, and show that the difference in timescales between negative and positive feedback loops, which gives rise to negative derivative feedback, can reduce the diffusivity dramatically. This result offers a biologically plausible explanation for the way our brain manages to store information about continuous parameters using noisy neurons. In addition, our results raise new questions about the relationship between information stored in a system and diffusivity due to noise.

I-49. Union of Intersections (UoI) for interpretable data driven discovery and prediction in neuroscience

Kris Bouchard[1,2] KRISTOFER.BOUCHARD@GMAIL.COM
Pratik Sachdeva[2] PRATIK.SACHDEVA@BERKELEY.EDU
Sharmodeep Bhattacharyya[3] BHATTASH@SCIENCE.OREGONSTATE.EDU
Mahesh Balasubramanian[4] MBALASU2@ASU.EDU
Shashanka Ubaru[5] UBARU001@UMN.EDU

[1]Lawrence Berkeley National Lab
[2]University of California, Berkeley
[3]Oregon State University
[4]Arizona State University
[5]IBM

Interpretable and predictive statistical data analysis methods can provide insight into the biological processes that generate neuroscience data. Neuroscience applications include network formation, receptive field estimation, determination of the genetic control of behaviors, and the calculation of meaningful low-dimensional representations. However, commonly used statistical inference procedures generally fail to identify the correct features, and further introduce consequential bias in the estimates. To address these issues, we developed Union of Intersections (UoI)1, a flexible, modular, and scalable framework for enhanced statistical feature selection and estimation. Methods (e.g., regression, classification, dimensionality reduction) based on UoI perform feature selection and feature estimation through intersection and union operations, respectively. In the context of linear regression (specifically UoILASSO), we summarize formal statistical proofs and extensive numerical investigation on synthetic data to demonstrate tight control of false-positives and false-negatives in feature selection with low-bias and low-variance estimates of selected parameters, while maintaining high-quality prediction accuracy (in the sense of cross-validation). In neuroscience data, we demonstrate: (i) the extraction of sparse, predictive, and interpretable functional networks from human ECoG during speech production and non-human primate single-unit recordings; (ii) increased model parsimony of spatio-temporal receptive fields from retinal ganglion cells; (iii) sparse genetic control of complex behavioral phenotypes; and (iv) parts-based decomposition (NMF) of neural spectrograms. Through optimization of single-node performance as well as multi-node scaling, we demonstrate the application of UoI to 4TB sized data sets in 45 minutes on a supercomputer2. To highlight the generality of the UoI framework, we show (with UoILogistic, UoICUR, and UoIVAR variants of the framework) improved prediction parsimony for classification, accurate matrix factorization, and dynamic modeling on diverse neuroscientific datasets1,3. These results demonstrate that UoI improves interpretation and prediction across diverse neuroscience applications.

I-50. Constructing models of heterogeneous dopamine activity

Ashok Litwin-Kumar AK3625@COLUMBIA.EDU
Columbia University

The Drosophila mushroom body exhibits dopamine (DA) dependent synaptic plasticity that underlies the acquisition and retrieval of associative memories. Recordings of DA activity in this system have identified signals related to external reinforcement such as reward and punishment. However, recent studies have also found that other factors including locomotion, novelty, reward expectation, and internal state also differentially modulate DA neurons. Such mixed signaling is also observed in the mammalian midbrain dopaminergic system. This heterogeneous activity is at odds with typical modeling approaches in which DA neurons are assumed to encode a global, scalar error signal. How can DA signals support appropriate synaptic plasticity in the presence of this heterogeneity? I develop a modeling approach that infers DA activity that is sufficient to solve a given set of behavioral tasks given architectural constraints informed by knowledge of mushroom body circuitry. Model DA neurons exhibit diverse tuning to task parameters while nonetheless producing coherent learned behaviors. These results provide a functional explanation for the heterogeneity of DA signals observed in associative learning.

I-51. Sensorimotor strategies and neuronal representations of whisker-based object recognition

Chris Rodgers CCR2137@COLUMBIA.EDU
Ramon Nogueira RN2446@COLUMBIA.EDU
Stefano Fusi SF2237@COLUMBIA.EDU
Randy Bruno RB2604@COLUMBIA.EDU
Columbia University

Humans and other animals can identify objects by active touch – coordinated exploratory motion and tactile sensation. Mice recognize objects by scanning them with their whiskers, just as we do with our fingertips. The brain must integrate exploratory motor actions with the resulting tactile signals in order to form a holistic representation of object identity. To elucidate the behavioral and neural mechanisms underlying this ability, we have developed a behavioral task for head-fixed mice – curvature discrimination – that challenges them to discriminate concave from convex shapes. In order to identify both efficient strategies and the mice’s actual strategy, we fit models to predict either stimulus or behavioral choice, respectively, using the motion of the whiskers and the resulting contacts. This analysis suggested that some whiskers preferentially contact convex shapes, and other whiskers concave. This is not a simple consequence of the shape geometry but a result of the active whisking strategies employed by the mice. Next, we recorded populations of neurons in the barrel cortex, which processes whisker input, to identify how it contributes to transforming fine-scale representations of sensory events into high-level representations of object identity. We fit each neuron’s firing rate as a function of both sensorimotor (e.g., whisker motion and touch) and cognitive (e.g., reward history) variables. Neurons responded strongly to whisker touch and, perhaps more surprisingly for a sensory area, to whisker motion. We also observed widespread and unexpected encoding of reward history. We suggest that barrel cortex neurons integrate sensory, motor, and reward signals in order to compute object shape. Taken together, these studies will shed light on the algorithm and implementation of the sensorimotor computations mediating object recognition.

I-52. A brain-machine interface for locomotion driven by subspace dynamics

Karen Schroeder KARSCH@GMAIL.COM
Sean Perkins SP3222@COLUMBIA.EDU
Qi Wang QW2161@COLUMBIA.EDU
Mark Churchland MC3502@COLUMBIA.EDU

Columbia University

Brain-machine interfaces (BMIs) for reaching have enjoyed continued performance improvements, allowing remarkable 2D cursor control. Yet there remains significant clinical need for locomotor (e.g., wheelchair control) BMIs, which could benefit a larger patient population. Fewer studies have addressed this need, and the best strategy for doing so remains undetermined. Here we demonstrate an approach based upon rhythmic neural activity. We leverage a behavioral task wherein monkeys cycle a hand-held pedal, forward or backward, to advance along a virtual track, pausing on targets for reward. This task does not emulate natural locomotion, but rather provides a view of cortical activity during learned, voluntary, rhythmic movement. Such activity is robust and was recently characterized (Russo et al. 2018), affording opportunities to develop novel decode algorithms and test them in an online setting.

We constructed a decoder that decoded virtual self-motion, based on recordings from 192 electrodes implanted in motor cortex. Unlike algorithms for cursor control, we did not directly map neural states to commanded velocity or position. Instead, we leveraged the most robust aspects of response structure: an overall shift in neural state when moving versus stationary, and rotations of the neural state while cycling. We used these features to decode when the subject was moving, and decoded direction based on the finding that neural-state rotations occur in different planes during forward versus backward cycling. Perhaps because the subject need not learn a novel mapping to control the BMI, performance was high even during the first few sessions of brain control. An additional performance gain was obtained by leveraging a neural dimension that reflected cycling direction at movement initiation. Resulting brain-control success rates were very close to those achieved under arm-control. Thus, rhythmic neural activity provides a robust substrate for BMI control, but requires different decode strategies than have been employed previously.

I-53. A theory of structured noise correlations in peripheral and higher order brain areas and their significance

Jesse Livezey[1] JLIVEZEY@LBL.GOV
Maximilian Dougherty[1] MAXDOUGHERTY@LBL.GOV
Pratik Sachdeva[2] PRATIK.SACHDEVA@BERKELEY.EDU
Sylvia Madhow[2] SMADHOW@LBL.GOV
Kristofer Bouchard[1] KEBOUCHARD@LBL.GOV

[1]Lawrence Berkeley National Laboratory
[2]University of California, Berkeley

Neural computations take place in the presence of noise. Across repeated presentations of a stimulus, neural measurements exhibit correlated variability (noise correlations). Measuring and understanding noise correlations are important for determining fundamental limits on the fidelity of neural representations. We address two outstanding issues in the field. The first is whether predictions about the structure of noise correlations and discriminability should differ in peripheral versus higher order sensory (or motor) areas in the brain. The second is how to quantify the significance of hypotheses about the optimality of observed correlations for different measures of discriminability. Our first contribution is a prediction, based on theory, that noise correlations measured in areas of the brain that represent a set of stimuli as categorical, e.g., phonemes in STG (Chang, 2010), will have a different structure compared to those measured in areas of the brain that represent a set of stimuli as continuous, e.g., rotations of moving bars in retina (Franke, 2016, Zylberberg, 2016). Our second contribution is a new null model for testing whether observed noise correlations are optimal. Specifically, this null model tests whether there is op-

timal alignment between the principal axes of the observed noise correlations and the mean stimulus-responses for a given discriminability measure. By contrast, previously proposed null models only test how often a model with no inter-neuron correlations but equal per-neuron statistics generates the observed discriminability. When compared across diverse real and synthetic datasets, we observe a profound difference in statistical significance, indicating that current null models cannot be used to test the optimality of observed correlations. Together, these results provide an experimentally testable prediction that the nature of neural computation (continuous versus categorical) should determine the structure of noise correlations across the neural hierarchy and improved methods for testing hypotheses about the optimality of observed noise correlations.

I-54. Point process latent variable models of larval zebrafish behavior

Scott Linderman[1] SCOTT.LINDERMAN@COLUMBIA.EDU
Anuj Sharma[1] 13ANUJS@GMAIL.COM
Robert Johnson[2] ROBERTEVANJOHNSON@FAS.HARVARD.EDU
Florian Engert[2] FLORIAN@MCB.HARVARD.EDU

[1]Columbia University
[2]Harvard University

A fundamental goal of systems neuroscience is to understand how neural activity gives rise to natural behavior. In order to achieve this goal, we must first build comprehensive models that offer quantitative descriptions of behavior. We develop new probabilistic models to tackle this challenge in the study of larval zebrafish, an important model organism for neuroscience. Larval zebrafish locomote via sequences of punctate swim bouts—brief flicks of the tail—which are naturally modeled as a marked point process. However, these sequences of swim bouts belie a set of discrete and continuous internal states, latent variables that are not captured by standard point process models. We incorporate these variables as latent marks of a point process and explore various models for their dynamics. To infer the latent variables and fit the parameters of these models, we develop an amortized variational inference algorithm that targets the collapsed posterior distribution, analytically marginalizing out the discrete latent variables. With a dataset of over 120,000 swim bouts, we show that our models reveal interpretable discrete classes of swim bouts and continuous internal states like hunger that modulate their dynamics. These models are an important step toward understanding the natural behavioral program of the larval zebrafish and, ultimately, its neural underpinnings.

I-55. Persistent gamma spiking in non-sensory fast-spiking cells predicts perceptual success

Hyeyoung Shin HYEYOUNG_SHIN@BROWN.EDU
Christopher Moore CHRISTOPHER_MOORE@BROWN.EDU

Brown University

Gamma oscillations (30 – 55 Hz) are hypothesized to temporally coordinate stimulus encoding, enabling perception. This hypothesis poses a conundrum: How can gamma serve as a persistent template when the stimulus itself drives its mediators, presumably perturbing its maintenance? Specifically, fast-spiking interneurons (FS), a key gamma generator, can be highly sensory responsive. Further, the gamma-band local field potential (LFP) shows properties inconsistent with temporal coordination. Combining tetrode recordings with controlled psychophysics, we found a subtype of non-sensory responsive FS that had a tendency to spike regularly at gamma-range intervals (gamma non-sensory FS; γnsFS). Successful detection was predicted by a further enhancement in the regularity of γnsFS gamma-interval spiking, persisting from before to after sensory onset. In contrast, the LFP gamma power negatively predicted detection, and was negatively correlated with gamma-band spiking of γnsFS. These results suggest that a distinct FS subgroup, not 'distracted' by sensory input, mediates perceptually relevant

gamma oscillations independent of the LFP.

I-56. Systematic changes of neural population activity during curl force field adaptation and generalization

Xulu Sun XULU2@STANFORD.EDU
Daniel O'Shea DJOSHEA@STANFORD.EDU
Eric Trautmann ETRAUT@STANFORD.EDU
Matthew D Golub MGOLUB@STANFORD.EDU
Tucker Fisher TGFISHER@STANFORD.EDU
Stephen Ryu SEOULMANMD@GMAIL.COM
Krishna Shenoy SHENOY@STANFORD.EDU

Stanford University

A hallmark of the motor system is its ability to execute different skilled movements as the situation warrants, thanks to the flexibility of motor learning. Much recent work has explored how the dynamics of neural populations drive movement preparation and execution, and a natural extension of this idea is to explore how changes in neural population dynamics facilitate motor learning. Previous neurophysiological studies of motor learning, using force-field adaptation paradigms, have characterized single neuron tuning properties and discovered that many individual neurons adapt their responses to compensate for the force field. However, a large portion of neurons showed heterogeneous responses to force-field learning that are challenging to understand when we look at one neuron at a time. A different approach is to understand motor learning in a neural population dynamics framework. Yet it remains unclear how the computation for motor learning operates in the neural population. Here, we explore the changes in neural population activity that occur during curl force field adaptation as well as the spatial generalization of adapted behaviors. We recorded neural activity in PMd and M1 when monkeys performed a curl field learning and generalization task, and found two patterns of systematic changes in population activity patterns before movement onset and during movement execution: 1) in a 2D neural state space, pre-movement neural population states rotate towards that of the adjacent target opposite the direction of the curl field, and the rotatory changes reflect spatial generalization of curl field learning, and 2) neural states for all reaching targets shift in a third dimension away from the baseline states, which indicates that the neural population explores a different repertoire of activity patterns after curl field learning. Overall, these findings are a first step towards understanding how neural population activity changes systematically to facilitate adaptation to novel arm dynamics.

I-57. Linking enhanced ability of detecting vocalization in noise with altered receptive fields

Natsumi Homma NATSUMI.HOMMA@UCSF.EDU
Christoph Schreiner CHRIS@PHY.UCSF.EDU

University of California, San Francisco

During critical periods, receptive fields (RFs) develop to adapt to the sensory environment by forming neural circuits that optimally process relevant stimuli. In humans, this developmental stage is essential for acquiring language. It has been shown that sound exposure during a critical period dynamically altered the RF properties in the auditory cortex such as frequency tuning, tuning bandwidth, or temporal resolution (de Villers-Sadani and Merzenich, 2011). However, the relationship between this altered neural coding and perceptual abilities is yet largely unknown. In this study, we tested the hypothesis that exposure to moderate levels of structured background noise during the critical period enhances the ability of adult animals to process vocalization in noise. We raised rat pups in moderate noises (60 dB SPL) of different spectro-temporal statistics during their auditory critical period (P6-45). Once these animals reached adulthood, we trained them to detect vocalizations presented in

these noises and compared to unexposed animals. The noise exposure enhanced their behavioral performance of detecting rat vocalizations in background noise. Improvement depended on stimulus statistics used for noise exposure. In addition, cortical signal encoding of vocalizations was improved in noise-exposed animals accompanied by specific shifts in RF properties compared to unexposed animals. In order to capture more complex aspects of altered neural coding, we used a multi-filter linear-nonlinear approach (Sharpee et al., 2004) and estimated their RFs. Preliminary data showed that noise exposure i) increased mutual information captured by the first filter, which is thought to act as a feature detector, and ii) enhanced synergy between the first and weaker, secondary filters. These findings support the idea that maturational noise exposure can alter cortical RF properties to enhance information extraction in noisy environment thus reducing the impact of background noise masking and helping the animals to perceptually segregate signals from noise background.

I-58. Automatically uncovering preferred stimuli of visual neurons using generative neural networks

Will Xiao[1] XIAOW@FAS.HARVARD.EDU
Peter Schade[1] PETER_SCHADE@HMS.HARVARD.EDU
Carlos Ponce[2] CRPONCE@WUSTL.EDU
Till Hartmann[1] TILL@HMS.HARVARD.EDU
Gabriel Kreiman[3] GABRIEL.KREIMAN@CHILDRENS.HARVARD.EDU
Margaret Livingstone[1] MARGARET_LIVINGSTONE@HMS.HARVARD.EDU

[1]Harvard University
[2]Washington University
[3]Boston Children's Hospital

What information does the brain represent, and what coding scheme does it use? These are central questions in neuroscience. Ever since Hubel & Wiesel (1962) discovered that neurons in V1 respond preferentially to bars of certain orientations, investigators have searched for preferred stimuli to reveal information encoded by neurons, leading to the discovery of cortical neurons that respond to specific motion directions (Hubel, 1959), color (Michael, 1978), binocular disparity (Barlow et al., 1967), curvature (Pasupathy & Connor, 1999), complex shapes such as hands or faces (Desimone et al., 1984; Gross et al., 1972), and even variations across faces (Chang & Tsao, 2017).

However, the classic approach for defining preferred stimuli depends on using a set of hand-picked stimuli, limiting possible answers to stimulus properties chosen by the investigator. Instead, we wanted to develop a method that is as general and uncorrelated with investigator bias as possible. To that end, we used a generative deep neural network as a vast and diverse hypothesis space. A genetic algorithm guided by neuronal preferences searched this space for stimuli. As a proof-of-concept, we evolved images to maximize firing rates of neurons in macaque inferior temporal cortex. The evolved images often evoked higher firing rates than the best of thousands of natural images. Furthermore, evolved images revealed neuronal selective properties that were sometimes consistent with existing theories but sometimes also unexpected. This method is well-positioned to corroborate, contrast with, and extend our understanding of neural coding in the visual cortex. It can be used to test existing theories of representation in the visual cortex, explore information encoded in coding schemes other than single-neuron spike rates—for example population response patterns, population correlations, local field potentials, etc.—and investigate other sensory modalities that can be modeled by neural networks.

I-59. A latent factor model for sensory coding and spontaneous activity in calcium imaging data

Marcus Triplett MARCUS.TRIPLETT@UQ.EDU.AU
Zac Pujic Z.PUJIC@UQ.EDU.AU
Biao Sun B.SUN1@UQ.EDU.AU
Lilach Avitan L.AVITAN@UQ.EDU.AU
Geoffrey Goodhill G.GOODHILL@UQ.EDU.AU

University of Queensland

The pattern of neural activity evoked by a stimulus can be substantially affected by ongoing spontaneous activity. For calcium imaging data there is currently a lack of statistically principled methods for estimating the relative contributions of evoked and spontaneous activity at each moment, making it difficult to, for instance, accurately estimate neural tuning curves. Here we develop a probabilistic latent variable model that decouples the components of calcium imaging data that are due to evoked activity from those driven by low dimensional spontaneous activity. We use variational inference to compute an approximate posterior distribution over latent sources of shared variability, identifying low dimensional structure underlying spontaneous activity and allowing for the estimation of stimulus tuning properties that take this variability into account. After validating the model on surrogate data we then apply it to the larval zebrafish optic tectum, revealing the extent to which neurons are driven by visual stimuli, latent sources of spontaneous activity, and their interaction. This model is broadly applicable to calcium imaging data, brings new insight into population-level neural activity, and can refine our understanding of the role of spontaneous activity in neural computation.

I-60. An egocentric view of neuronal noise correlation

Wen-Hao Zhang[1] WENHAO.ZHANG@PITT.EDU
Si Wu[2] SIWU@PKU.EDU.CN
Krešimir Josić[3] JOSIC@MATH.UH.EDU
Brent Doiron[1] BDOIRON@PITT.EDU

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

Cortical activity is noisy and correlated across neurons. Neural correlations can be the result of shared variability in feedforward inputs, generated internally through recurrent interactions, or driven by feedback from other cortical areas. However it remains unclear how correlated patterns in neural responses support the processing and interpretation of stimuli, and whether their origin determines their role. To address this question, we propose a hierarchical generative model describing sensory transduction, and propose a mechanism for the inference of latent aspects of stimuli through iterative sampling. We show how such sampling can be implemented in a stochastic recurrent neural network. The recurrent inputs in this network support the structured variability that is needed for sampling, and provide additional information about the stimulus that can be obtained from the context. The population activity combines the feedforward and recurrent inputs and represent the brain's estimate of an external stimulus via a probabilistic population code. We find that the neuronal stimulus estimate can be optimally decoded from population activity using a linear decoder which is blind to the overall correlation in the network. This decoupling of stimulus estimation and population correlations allows an increase in recurrent synaptic strength to drive higher gain responses without the typical penalty induced by neuronal correlations. Our model suggests that noise correlations can be byproducts of the recurrent computations underlying the estimation, and reflect the fluctuation in an ongoing estimation computation.

I-61. Long time-scales in primate amygdala neurons support aversive learning

Aryeh Taub ARYEH.TAUB@WEIZMANN.AC.IL
Yosef Shohat YOSEF.SHOHAT@WEIZMANN.AC.IL
Rony Paz RONY.PAZ@WEIZMANN.AC.IL
Weizmann Institute of Science

Associative learning forms when there is temporal relationship between a stimulus and a reinforcer, yet the inter-trial-interval (ITI), which is usually much longer than the stimulus-reinforcer-interval, contributes to learning-rate and memory strength. The neural mechanisms that enable maintenance of time between trials remain unknown, and it is unclear if the amygdala can support time scales at the order of dozens of seconds. In the present study, we preformed in monkeys daily sessions of tone CS – aversive odor US with ITI that varies between 20 – 60 sec. We analyzed both local activity in the dorsal-anterior-cingulate-cortex (ACC) and BLA and long-range communication between these structures during the ITI. We further computed learning strength and rate and correlated neuronal dynamics during the ITI with behavioral change on a trial-by-trial basis. We show that the ITI indeed modulates rate and strength of aversive-learning , and that single-units in the primate amygdala and ACC signal confined periods within the ITI, strengthen this coding during acquisition of aversive-associations, and diminish during extinction. Additionally, pairs of amygdala-cingulate neurons synchronize during specific periods suggesting a new role for this network in maintaining timescales to build the temporal and/or statistical structure of the environment, and suggests a mechanism to how longer intervals can promote learning and memory. In turn, it can also explain why under circumstances of multiple subsequent experiences, aversive memories are formed faster and exhibit strong-to-abnormal responses. Hence, in addition to existing learning models that implicate this specific network in anxiety and fear-disorders, our findings suggest a new model for how deviations in representation and computation in this circuitry can lead to maladaptive and exaggerated behaviors.

I-62. Neural network models for closed-loop musculoskeletal arm control

Francis Willett WILLETT2@GMAIL.COM
Krishna Shenoy SHENOY@STANFORD.EDU
Saurabh Vyas SMVYAS@STANFORD.EDU
Jonathan Michaels JMICHAE@STANFORD.EDU
Jaimie Henderson HENDERJ@STANFORD.EDU
Stanford University

Neural network models of reaching can aid neuroscientific investigations by generating hypotheses about neural mechanisms that may underlie existing data or by making predictions for what neural activity might look like during a novel task. To be most useful, neural network models should be capable of skilled motor control, with the ability to use visual and proprioceptive feedback to correct for errors during a reach. Unfortunately, neural network models that can engage in accurate and robust feedback control of a musculoskeletal arm are difficult to build. They require both a computationally efficient biomechanical model of an arm and a learning method powerful enough to teach the network to deal with the properties of muscles, rigid body dynamics, feedback delays and task requirements. Here we develop a new approach for building such flexible neural network models. Our approach is centered around building a differentiable approximation of a biomechanical arm model using a recurrent neural network (RNN) to speed up training. To demonstrate our approach, we use it to train a multi-compartment recurrent neural network controller to make instructed delay point-to-point reaches to arbitrary locations within a 20 x 20 cm workspace in the presence of visual distortions and perturbations. We show that the controller can successfully reach towards targets while correcting for cursor perturbations and adapting to visuomotor rotations, and that the model's neural activity in an instructed delay reaching task bears important similarities to nonhuman primate (NHP) neural data. Our approach is the first to our knowledge that enables fast and flexible training of recurrent neural network reaching models that are powerful enough to control complex, 3D musculoskeletal arms

in closed-loop. We aim to accelerate progress towards a community-maintained “working model” of motor cortical reaching by releasing trained arm approximations and software to the scientific community.

I-63. The emergence of multiple retinal cell types through efficient coding of natural movies

Samuel Ocko SOCKO@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Stephane Deny SDENY@STANFORD.EDU
Jack Lindsey LINDSEY6@STANFORD.EDU
Stanford University

One of the most striking aspects of early visual processing in the retina is the immediate parcellation of visual information into multiple parallel pathways, formed by different retinal ganglion cell types each tiling the entire visual field. Existing theories of efficient coding have been unable to account for the functional advantages of such cell-type diversity in encoding natural scenes. Here we go beyond previous theories to analyze how a simple linear retinal encoding model with different convolutional cell types efficiently encodes naturalistic spatiotemporal movies given a fixed firing rate budget. We find that optimizing the receptive fields and cell densities of two cell types makes them match the properties of the two main cell types in the primate retina, midget and parasol cells, in terms of spatial and temporal sensitivity, cell spacing, and their relative ratio. Moreover, our theory gives a precise account of how the ratio of midget to parasol cells decreases with retinal eccentricity. Also, we train a nonlinear encoding model with a rectifying nonlinearity to efficiently encode naturalistic movies, and again find emergent receptive fields resembling those of midget and parasol cells that are now further subdivided into ON and OFF types. Thus our work provides a theoretical justification, based on the efficient coding of natural movies, for the existence of the four most dominant cell types in the primate retina that together comprise 70% of all ganglion cells.

I-64. The effects of neural resource constraints on early visual representations

Jack Lindsey LINDSEY6@STANFORD.EDU
Samuel Ocko SOCKO@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Stephane Deny SDENY@STANFORD.EDU
Stanford University

Neural representations vary drastically across the first stages of visual processing. At the output of the retina, ganglion cell receptive fields (RFs) exhibit a clear antagonistic center-surround structure, whereas in the primary visual cortex (V1), typical RFs are sharply tuned to a precise orientation. There is currently no theory explaining these differences in representations across layers. Here, using a deep convolutional neural network trained on image recognition as a model of the visual system, we show that such differences in representation can emerge as a direct consequence of different neural resource constraints on the retinal and cortical networks, and for the first time we find a single model from which both geometries spontaneously emerge at the appropriate stages of visual processing. The key constraint is a reduced number of neurons at the retinal output, consistent with the anatomy of the optic nerve as a stringent bottleneck. Second, we find that, for simple downstream cortical networks, visual representations at the retinal output emerge as non-linear and lossy feature detectors, whereas they emerge as linear and faithful encoders of the visual scene for complex brains. This result predicts that the retinas of small vertebrates (e.g. salamander, frog) should perform sophisticated nonlinear computations, extracting features directly relevant to behavior, whereas retinas of large animals such as primates should mostly encode

the visual scene linearly and respond to a much broader range of stimuli. These predictions could reconcile the two seemingly incompatible views of the retina as either performing feature extraction or efficient coding of natural scenes, by suggesting that all vertebrates lie on a spectrum between these two objectives, depending on the degree of neural resources allocated to their visual system.

I-65. The dynamically balanced critical state in motor cortex

Moritz Helias M.HELIAS@FZ-JUELICH.DE
David Dahmen D.DAHMEN@FZ-JUELICH.DE
Sonja Grun S.GRUEN@FZ-JUELICH.DE
Markus Diesmann DIESMANN@FZ-JUELICH.DE

Research Centre Julich

The current work uncovers heterogeneity of connections in cortical networks as a source of a qualitatively new type of critical dynamics. Analyzing massively parallel spiking data from 155 neurons in macaque motor cortex we find a wide distribution of pairwise correlations. We quantitatively explain this observation by developing a finite-size theory of heterogeneous correlations (Dahmen et al., 2018); it combines linear response theory, field theory (Chow and Buice, 2015) and methods from disordered systems (Sompolinsky and Zippelius, 1982). Theory and data unequivocally imply a novel type of network dynamics: multiple dynamical modes are close to the breakdown of linear stability - a state of critical dynamics that unfolds on a low dimensional manifold. This work reconciles two apparently contradictory dynamical regimes of cortical operation: Criticality is typically associated with avalanche-like activation of large populations of neurons (Beggs and Plenz, 2003); it requires a tight adjustment between excitation and inhibition and often features high computational performance. In contrast, dynamical balance (van Vreeswijk & Sompolinsky, 1996) robustly results from excess inhibition which leads to decorrelation of neural activity (Renart et al., Science, 2010; Ecker et al., 2010), as widely observed in mammalian cortex. Our work turns the view from criticality in the one-dimensional population activity to criticality that unfolds in a multitude of directions in the space of all neurons. We thus hypothesize that heterogeneity of connections is the structural substrate for a rich dynamical repertoire of motor cortex that is essential for high performance in such concepts as reservoir computing (Maass et al., 2002; Jaeger & Haas, 2004).

I-66. Feature selective encoding of substrate vibrations in the forelimb somatosensory cortex

Daniel Huber DANIEL.HUBER@UNIGE.CH
Mario Prsa MARIO.PRSA@UNIGE.CH
Karin Morandell KARIN.MORANDELL@UNIGE.CH
Geraldine Cuenu ANGIE.CUENU@UNIGE.CH

University of Geneva

How features of vibrotactile stimuli are cortically encoded has primarily been investigated using active sensory channels such as the primate hand or rodent whiskers. High frequency ($>$100 Hz) vibrations delivered to monkey fingertips or rat whiskers produce temporally entrained spiking and frequency independent firing rates in primary somatosensory cortex (S1). In such active systems the coding scheme for frequency recognition seems therefore to be based on identifying the dominant vibration phase in the cyclically entrained spike trains rather than on differences in rates of spike discharge. Here we elucidate how substrate vibrations applied to the mouse forelimb, a fundamentally passive sensory channel, are perceived and represented in S1. In contrast to vibrotactile stimulation of the primate hand, we found that mouse forelimb S1 (fS1) neurons rely on a different coding scheme: their spike rates are tuned to a preferred stimulus feature and not temporally entrained. Information theory analysis revealed that maximum information is conveyed at the tuning curve peaks, suggesting that these neurons operate as feature detectors. The encoded feature was found to be the product of vibration frequency and a power function of amplitude. We identified the same quantity as the perceptually relevant stimulus variable when mice performed a frequency discrimination task. Surprisingly, in half of responsive neurons, an ideal observer could predict behavioral choice from neural activity evoked by identical stimuli. Taken together, our findings suggest that the stimulus feature encoded by fS1 neurons at their peak firing rates is also what the mice perceive upon forelimb vibration and, unlike neurons in primate S1, their responses are not essentially sensory. Finally, histology, deafferentation and optogenetic tagging revealed that these responses originate from subcutaneous mechanoreceptors located on forelimb bones. We conclude that mouse limbs are best adapted for passive “listening” of substrate vibrations, rather than for active vibrotactile exploration.

I-67. Inferred network model spontaneously generates complex spatio-temporal dynamics matching spike data

Guido Gigante[1,2] GUIDO.GIGANTE@GMAIL.COM
Cristiano Capone[3] CRISTIANO0CAPONE@GMAIL.COM
Paolo Del Giudice[1] PAOLO.DELGIUDICE@ROMA1.INFN.IT

[1]Istituto Superiore di Sanita
[2]PRORA
[3]Istituto Nazionale di Fisica Nucleare

Statistical inference methods are widely used in neuroscience to recover effective models from observed data. However, few studies attempted to investigate the spiking dynamics of inferred models, and none, to our knowledge, at the network level. We introduce a principled modification of a widely used generalized linear model (GLM), and learn its structural and dynamic parameters from in-vitro spike data. The spontaneous activity of the new model captures prominent features of the non-stationary dynamics displayed by the biological network, where the reference GLM largely fails: the model in fact autonomously generates a wide spectrum of network bursts, with inter-burst interval and event size statistics that fairly match the ones observed in the data. The inference procedure consistently reveals two well-defined time-scales, one at around 100 ms and one at few seconds, for the introduced spike-frequency adaptation mechanism; such mechanism explains significant temporal correlations observed in the bursting activity of the real network. At a more fine-grained level, the model is able to reproduce the order of activation of the individual neurons inside bursts: the role of the inferred synaptic efficacies thus goes beyond accounting for correlations in the firing of different neurons, towards the uncovering of effective causal dynamical interactions. Two ingredients were key for the success of the proposed model in comparison with more standard GLMs. The first is a saturating transfer function: beyond its biological plausibility, it limits the neuron’s information transfer, improving robustness against endogenous and external noise. The second is a super-Poisson spikes generative mechanism, that allows the model neuron to flexibly incorporate the observed activity fluctuations; we hypothesize that such mechanism emerges from the undersampling of the network. To test this hypothesis, we present preliminary results, using high density (4096-electrode) recordings, for the scaling of the unobserved fluctuations with the subsampling of the network.

I-68. Learning-induced involvement of cortical areas during a whisker-based short-term memory task

Vahid Esmaeili VAHID.ESMAEILI@EPFL.CH
Marta Boscaglia BOSCAGLIAMARTA@GMAIL.COM
Eloise Charriere ELOISE.CHARRIERE@EPFL.CH
Anastasiia Oryshchuk ANASTASIIA.ORYSHCHUK@EPFL.CH
Sylvain Crochet SYLVAIN.CROCHET@EPFL.CH
Carl Petersen CARL.PETERSEN@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Execution of goal-directed behaviors involves coordination of a distributed network of cortical regions. How different brain areas encode different aspects of behavior and how these representations change upon learning remains to be elucidated. To address this, we devised a Go/No-go task in which head-fixed mice detect a brief whisker deflection to collect a reward. A visual cue signals the beginning of each trial. Upon whisker stimulation, mice need to withhold licking during a post-stimulus delay until an auditory go-cue is presented. This delay period separates sensory processing from action initiation, allowing us to investigate which brain areas contribute to sensory detection and action planning. We performed silicon probe recordings in multiple brain areas while Expert or Novice mice performed the task. Upon whisker stimulation, a short-latency response appeared with a sequential order in primary somatosensory cortex (wS1) and secondary somatosensory cortex (wS2) followed by primary (wM1) and secondary motor cortex (wM2) in both Expert and Novice mice. Anterior lateral motor cortex (ALM), tongue/jaw motor cortex (tjM1) and posterior parietal cortex (PPC) represented the whisker stimulus with longer latencies and their activity bridged the delay period in Expert but not in Novice mice. ALM and tjM1 transiently signaled initiation of licking earlier than other areas. Optogenetic inhibition during specific behavioral epochs revealed distinct involvement of cortical areas. During the whisker stimulus presentation, inactivation of all tested areas except primary visual cortex (V1) and wM1 disrupted performance with wS2 having the largest effect. Inactivation of ALM and tjM1 during the delay period, decreased hit rate while wM2 inhibition increased false alarms. During the response period, inactivation of tjM1, as well as ALM, wM2 and wM1, suppressed licking in both Go and No-go trials. Our data thus defines global and local entrainment of different cortical areas upon learning of a short-term memory goal-directed behavior.

I-69. Spatio-temporal dynamics of sensorimotor transformations in mouse dorsal cortex

Ivana Orsolic[1] ORSOLIC.IVA@GMAIL.COM
Petr Znamenskiy[2] P.ZNAMENSKIY@UCL.AC.UK
Thomas Mrsic-Flogel[2] T.MRSIC-FLOGEL@UCL.AC.UK

[1]Sainsbury Wellcome Centre, UCL
[2]University College London

Perceptual decisions are the product of concerted activity of a distributed network of brain areas. To characterise the activity within these pathways, we combined wide-field calcium imaging of dorsal cortex with a behavioural paradigm wherein animals had to detect a sustained increase in the temporal frequency of a visual stimulus. Temporal frequency was varied stochastically every 50 ms, allowing us to explore the animals' behavioural strategy and its neural underpinnings. Psychophysical reverse correlation showed that mice accumulated sensory evidence on the timescale of 100s of milliseconds, confirmed by model-based analyses. Consistent with recent reports, behavioural choices were associated with global modulation of neural activity. Transient fluctuations in the temporal frequency of the stimulus, which did not lead to overt motor responses, allowed us to examine how sensory evidence is processed and integrated across dorsal cortex leading up to choice. Such subthreshold sensory inputs triggered a localised cascade of activity. The cascade originated in visual cortex with low-latency and fast-decaying responses and spread to secondary motor and then primary motor cortices where responses

were sustained. On a moment-to-moment basis, stimulus fluctuations carried pro-licking or anti-licking information for the animal. While visual sensory areas were bidirectionally modulated by stimuli carrying either type of information, both in trained and naive animals, motor areas responded selectively to pro-licking stimuli and only in trained animals. These results provide a systematic survey of the engagement of the mouse dorsal cortex during perceptual decision-making. They reveal patterns of neural activity underlying the covert antecedents of behavioural choices and highlight secondary motor cortex as the site of integration of sensory evidence.

I-70. Memories in coupled winner-take-all networks

Roman Pogodin ROMAN.POGODIN.17@UCL.AC.UK
Peter Latham PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, University College London

A salient, and incredibly useful, feature of human memory is that it is content addressable: smell a chocolate cake, for example, and your 8th birthday comes flooding back. As shown by Hopfield over 30 years ago, implementing content addressable memory in a neural network is straightforward: a memory is nothing more than a stable fixed point of the network dynamics, and nearby activity is the “content”. However, as shown theoretically by Roudi and Latham (2007), Hopfield networks implemented with spiking neurons can contain at most several hundred memories; far too few to explain human performance. Recently, Chaudhuri and Fiete (2017) constructed a network, based on winner-take-all dynamics, for which the number of memories is exponential in the number of neurons. The exponential capacity was a huge breakthrough, and to reach human performance their network requires only a few thousand neurons. Here we analyze a memory network with similar architecture. It doesn't have exponential capacity, but it can also reach human level performance with only a few thousand neurons. And it has two other advantages: the architecture is simpler, and it is much more robust to errors in the input. Gripon and Berrou (2011) analyzed the capacity of this network; here we extend their work in several directions: we show how the number of memories depends on the size of the basin of attraction and the signal-to-noise ratio; we show that high capacity can be achieved with almost no spurious memories, a guarantee that is not, as far as we know, possible to make in other architectures; and finally, we show that implementation in networks of spiking neurons is reasonably straightforward. The proposed architecture could, therefore, serve as a substrate for human memory.

I-71. Arbitration between social and reward learning systems: Computational and neural signatures

Andreea Diaconescu[1] DIACONESCU@BIOMED.EE.ETHZ.CH
Christoph Mathys[2] CMATHYS@SISSA.IT
Lars Kasper[3,4] KASPER@BIOMED.EE.ETHZ.CH
Madeline Stecy[4] MSTECY@GMAIL.COM
Christopher J Burke[4] CHRISTOPHER.BURKE@ECON.UZH.CH
Zoltan Nagy[4] ZOLTAN.NAGY@UZH.CH
Klaas Enno Stephan[4,3] STEPHAN@BIOMED.EE.ETHZ.CH
Philippe N Tobler[4] PHIL.TOBLER@ECON.UZH.CH

[1]University of Basel
[2]International School for Advanced Studies (SISSA)
[3]ETH Zurich
[4]University of Zurich

As social primates navigating an uncertain world, humans use multiple information sources including expert advice to guide their decisions. Little is known about how the brain arbitrates between primary reward and social learning for optimal decision-making. In this study, we formalised arbitration as the precision afforded by each learning system using a series of hierarchical Bayesian models. We used a probabilistic learning task in which participants predicted the outcome of a lottery using recommendations from a more informed advisor and sampled card outcomes. While the card probabilities changed over time, the advisor was motivated to give helpful advice depending on the phase of the task, resulting in variations in his/her intentions (i.e., volatility). Bayesian model selection revealed one winning model with arbitration as the ratio of precisions dynamically weighting predictions and a social bias parameter reflecting that participants were unequally persuaded by the advisor (Fig. 1). Decisions of how many points to wager were determined by both arbitration and the volatility of the advisor's intentions: The higher the relative precision of one learning system over the other and the lower the intention volatility, the more points participants wagered. Using fMRI, we localised the arbitration signal to the dopaminergic midbrain such that activity increased with increasing relative precision of learning system relative to the other ($p<0.05$ FWE peak-level within an anatomical mask including all neuromodulatory nuclei) (Fig. 2). Precision-weighted predictions of fidelity activated the right amygdala, bilateral fusiform gyrus, dorsolateral PFC, and orbitofrontal cortex (whole-brain FWE cluster-level corrected, $p<0.05$). These results suggest that decisions to rely on one's own experience rather than advice depends on a general bias with which social information is weighted and the inferred reliability of each information source, arbitration, which was represented in the dopaminergic midbrain and regions previously associated with theory-of-mind functions.

I-72. Bidirectional coupling of grid cells and place cells as a mechanism for robust spatial representation across multiple maps

Haggai Agmon HAGGAI.AGMON@MAIL.HUJI.AC.IL
Yoram Burak YORAM.BURAK@ELSC.HUJI.AC.IL

Hebrew University of Jerusalem

Grid cells and place cells play an important role in spatial encoding. Grid cells in the medial entorhinal cortex have multiple periodically spaced firing fields that form a hexagonal lattice and are largely similar between environments. Hence, they are considered as reasonable candidates to be associated with path integration. Hippocampal place cells, on the other hand, typically display at most a single environment specific receptive field, and are associated with contextual memory since they exhibit 'global remapping', drastically changing their firing field between environments. In the past decade many experimental and theoretical studies examined the influence that grid cells and place cells might have on one another. Studies have demonstrated correlated phenomena between activities of grid cells and place cells under various manipulations, yet the exact mechanism and functional relationship between grid cells and place cells and its significance to encoding of spatial location is still unclear. Some studies have suggested that grid firing patterns are the main determinant of place cell firing while other studies have challenged this view arguing that grid fields are formed as a result of place cell input. Overall, studies so far were primarily concerned with uni-directional manipulations and their implications, namely, they dealt either only with the effects and emergence of grid cells from place cell inputs or vice versa. Instead of treating grid cells and place cells as two separate populations in successive stages of a processing hierarchy, we develop in this work a detailed computational attractor model with mutual interacting connections between grid cells and place cells and demonstrate the functional significance of this coupling. In this framework grid cell and place cell activities serve as complementary and interacting representations that work in combination to support the reliable coding of large-scale space over multiple environments.

I-73. Classification of time series using the covariance perceptron

Matthieu Gilson[1] MATTHIEU.GILSON@UPF.EDU
Ruben Moreno-Bote[1] RUBEN.MORENO@UPF.EDU
Andrea Insabato[1] ANDREA.INSABATO@UPF.EDU
Moritz Helias[2] M.HELIAS@FZ-JUELICH.DE

[1]Universitat Pompeu Fabra
[2]Research Centre Julich

Many efforts in the study of the brain have focused on representations of stimuli by neuronal networks. The perceptron [1] is a central concept that has brought many fruitful theories in the fields of neural coding and learning in networks. The present study considers the problem of categorizing fluctuating time series determined by their hidden dynamics. We propose a new version of the perceptron based on the multivariate autoregressive (MAR) dynamics [2] to learn covariance patterns that are computed using an observation window, thereby implementing batch learning. For each stimulus presentation, the covariances of the simulated output is compared to objective covariances for the corresponding category of input patterns. Here binary classification is implemented by mapping each of two input categories to larger variance in one of two output nodes. While the classical perceptron treats inputs as i.i.d. variables (i.e. vectors of mean activity), the covariance perceptron explicitly focuses on the temporal structure of input time series. We explore both analytically and numerically the robustness of this covariance-based classification scheme. In particular, its capacity (how many patterns can be successfully categorized) compares favorably with the classical perceptron. Conceptually, variability in the time series is the basis for information to be learned, via the second-order statistics of fluctuating inputs involving time lags. Our approach based on transitions for temporally correlated activity is thus a radical change of perspective compared to classical approaches that typically transform time series into a succession of static patterns where fluctuations are noise. This opens a promising new perspective for learning and coding in spiking neuronal networks, for example in relation with spike-timing-dependent plasticity that is also based on covariances.

[1] CM Bishop (2006) Pattern Recognition and Machine Learning. Springer. [2] H Lutkepohl (2005) New introduction to multiple time series analysis. Springer

I-74. Homeostatic plasticity and external input explain difference in neural spiking activity in vitro and in vivo

Johannes Zierenberg JOHANNES.ZIERENBERG@DS.MPG.DE
Jens Wilting JENS.WILTING@DS.MPG.DE
Viola Priesemann VIOLA.PRIESEMANN@DS.MPG.DE

Max Planck Institute for Dynamics and Self-Organization

In vitro and in vivo spiking activity clearly differ. Networks in vitro develop strong bursts separated by periods of very little spiking activity, whereas in vivo cortical networks show continuous activity. This is puzzling considering that both networks presumably share similar single-neuron dynamics and plasticity rules. We propose that the defining difference between in vitro and in vivo dynamics is the strength of external input. In vitro, networks are virtually isolated, whereas in vivo every brain area receives continuous input. We analyze a model of spiking neurons in which the input strength, mediated by spike rate homeostasis, determines the characteristics of the dynamical state. In more detail, our analytical and numerical results on various network topologies consistently show that under increasing input, homeostatic plasticity generates distinct dynamic states, from bursting, to close-to-critical, reverberating and irregular states. This implies that the dynamic state of a neural network is not fixed but can readily adapt to the input strength. Indeed, our results match experimental spike recordings in vitro and in vivo: the in vitro bursting behavior is consistent with a state generated by very low network input ($<0.1\%$), whereas in vivo activity suggests that on the order of 1% recorded spikes are input-driven, resulting in reverberating dynamics. Importantly, this predicts that one can abolish the ubiquitous bursts of in vitro preparations, and instead impose dynamics comparable to in vivo activity by exposing the system to weak long-term stimulation, thereby opening new paths to establish an in vivo-like assay in vitro for basic as well as neurological studies.

I-75. Multi-step planning of eye movements in visual search

David Hoppe HOPPE@PSYCHOLOGIE.TU-DARMSTADT.DE
Constantin A Rothkopf CONSTANTIN.ROTHKOPF@COGSCI.TU-DARMSTADT.DE

Technical University Darmstadt

The capability of directing gaze to relevant parts in the environment is crucial for our survival. Computational models have proposed quantitative accounts of human gaze selection in a range of visual search tasks. Initially, models suggested that gaze is directed to the locations in a visual scene at which some criterion such as the probability of target location, the reduction of uncertainty or the maximization of reward appear to be maximal. But subsequent studies established, that in some tasks humans instead direct their gaze to locations, such that after the single next look the criterion is expected to become maximal. However, in tasks going beyond a single action, the entire action sequence may determine future rewards thereby necessitating planning beyond a single next gaze shift. While previous empirical studies have suggested that human gaze sequences are planned, quantitative evidence for whether the human visual system is capable of finding optimal eye movement sequences according to probabilistic planning is missing. Here we employ a series of computational models to investigate whether humans are capable of looking ahead more than the next single eye movement. We found clear evidence that subjects' behavior was better explained by the model of a planning observer compared to a myopic, greedy observer, which selects only a single saccade at a time. In particular, the location of our subjects' first fixation differed depending on the stimulus and the time available for the search, which was well predicted quantitatively by a probabilistic planning model. Overall, our results are the first evidence that the human visual system's gaze selection agrees with optimal planning under uncertainty.

I-76. Is simple value-based decision-making suboptimal?

Nisheet Patel[1] NISHEET.PATEL@UNIGE.CH
Jan Drugowitsch[2] JAN_DRUGOWITSCH@HMS.HARVARD.EDU
Alexandre Pouget[1] ALEXANDRE.POUGET@UNIGE.CH

[1]University of Geneva
[2]Harvard University

Every-day decisions frequently require choosing among multiple alternatives. Whether humans and other animals decide optimally in situations involving more than two choices has been debated in the past. For instance, previous physiological and behavioural experiments show that the introduction of a low-valued option given two high-valued options impacts the choice between the two high-valued options even if the newly introduced option is rarely chosen. This is known as the violation of the independence of the irrelevant alternative (IIA). To address the question of optimality, we had previously derived the normative strategy for multi-alternative value-based decision-making. The optimal policy required evidence accumulation to nonlinear, time-dependent bounds that trigger choices, and this was near-perfectly approximated by a neural circuit implementing a normalized diffusion process with fixed decision bounds and an urgency signal.

Here, we explore whether our previously derived optimal neural circuit can account for several behaviourally observed “suboptimal” phenomena such as the violation of the IIA, the similarity effect and the violation of the regularity principle. We show that our circuit model can account for these effects in the presence of divisive normalization, a putative canonical computation in the brain, and internal variability due for instance to ongoing learning of the decision stopping bounds. We also compare our optimal circuit model to a race-model and other variants of both models, and show that our model copes with such internal variability much better than the race model variants. Thus, our results demonstrate that a variety of puzzling interactions between choices, that have

been deemed as suboptimal, may, in fact, be the consequence of the nonlinear implementation of the optimal policy in neural circuits combined with ongoing learning, or any other sources of internal variability.

I-77. Generalisation of structural knowledge in hippocampal – prefrontal circuits

Veronika Samborska[1] VERONIKA.SAMBORSKA@NDCN.OX.AC.UK
Thomas Akam[1] THOMAS.AKAM@NEURO.FCHAMPALIMAUD.ORG
James Butler[2] JAMES.BUTLER@UCL.AC.UK
Mark Walton[1] MARK.WALTON@PSY.OX.AC.UK
Tim EJ Behrens[1] BEHRENS@FMRIB.OX.AC.UK

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

A hallmark of intelligence is the ability to generalise previously learned knowledge to solve novel analogous problems. For instance, someone who knows how to drive a car in Europe can quickly adapt to driving in the U.K without having to relearn driving from scratch simply because some rules and motor actions are different. Such transfer of knowledge relies on formation of representations that are abstracted from sensory states and humans excel at this task. However, little is known about whether other animals can learn abstract representations that are detached from sensory experience. Even less is known about how the brain represents such abstract representations while maintaining the content of individual experiences.

Here we present a novel behavioural paradigm for investigating generalisation of structural knowledge in mice, and report preliminary electrophysiological findings from single neurons in hippocampus and prefrontal cortex. Mice serially performed a set of reversal learning tasks, which shared the same structure (e.g., one choice port is good at a time), but had different physical configurations and hence different sensory and motor representations. Subjects' performance on novel configurations improved with the number of configurations they had already learned, demonstrating generalisation of knowledge. As in spatial remapping experiments, many hippocampal neurons responded differently in different configurations – here tasks rather than spatial environments. In contrast, preliminary analyses suggest prefrontal representations were more general and reflected different stages of the trial irrespective of the current physical configuration.

I-78. Amygdala antagonistic inhibitory clusters control fear switch

Kenta Hagihara[1] KENTA.HAGIHARA@FMI.CH
Andreas Luthi[1] ANDREAS.LUTHI@FMI.CH
Olena Bukalo[2] OLENA.BUKALO@NIH.GOV
Martin Zeller[3] MARTIN.ZELLER@STUDENT.UNI-TUEBINGEN.DE
Ayla Aksoy-Aksel[3] AYLA.AKSOY-AKSEL@UNI-TUEBINGEN.DE
Ingrid Ehlich[3] INGRID.EHRLICH@UNI-TUEBINGEN.DE
Andrew Holmes[2] HOLMESAN@MAIL.NIH.GOV

[1]Friedrich Miescher Institute for Biomedical Research
[2]NIH/NIAAA
[3]University of Tubingen

Excessive fear expression leads to loss of opportunities in many aspects and disadvantages animal survival. Thus, properly suppressing fear memory is as important for survival as the initial formation of a fear memory. Intercalated cells (ITCs) of the amygdala, clusters of small-sized unique inhibitory neurons encapsulating the lateral and basal amygdala, are anatomically well situated to provide inhibition on amygdala circuits underlying fear expression, and thus, have been implicated in fear suppression upon extinction. However, largely due to their

small size and lack of genetic markers for targeting, it has been difficult to functionally dissect roles of individual ITC clusters. Here, to address this issue, we employed a multidisciplinary approach including in vivo calcium imaging from freely moving mice engaging a classical auditory fear conditioning and extinction paradigm, opto-/pharmaco-genetic activity manipulations, slice physiology, virus-based circuit tracings, and RNA-sequencing. We found that two transcriptionally distinguishable major clusters of ITCs, dmITC and vmITC oppositely represent fear-eliciting and extinguished auditory stimuli, and that manipulations of the two clusters changed the fear state of animals in opposite directions. We also found that dmITC and vmITC clusters directly and mutually inhibit each other while innervating to the basolateral amygdala (BLA), suggesting that they comprise antagonistic inhibitory control over BLA fear circuit. Furthermore, virus based projection mapping revealed that they target different brain regions. Overall, those results suggest that defined ITC clusters interact with each other to control the fear state of animals by influencing distributed brain-wide computation. The current study provides a new circuit motif where two mutually inhibiting units perform “push-pull” calculation.

I-79. A theory of inference amid fixational drift explains the contrast paradox in vernier hyperacuity

Nadav Ben Shushan NADAV.BEN-SHUSHAN@MAIL.HUJI.AC.IL
Yoram Burak YORAM.BURAK@ELSC.HUJI.AC.IL

Hebrew University of Jerusalem

Psychophysical experiments reveal a counterintuitive effect, the “contrast paradox”, in the Vernier hyperacuity task (SB Stevenson, LK Cormack, Vision Res, 2000). In this task a subject reports on the relative position of two misaligned vertical bars. Increasing the contrast of both bars improves performance. However, increasing the contrast of only one bar degrades performance - even though naively more information about the position of the higher-contrast bar is arriving at the retina. Here we argue that the contrast paradox may arise due to the continuous, random movement of the eye during fixation, known as fixational drift. Motion of the eyes might shift the apparent position of the bars, relative to their instantaneous location. If the upper and lower bars have the same contrast, their apparent shifts are identical. On the other hand, if the upper and lower bars differ in contrast, their apparent shifts may differ. A bias may arise in their perceived separation, depending on the detailed trajectory of the eye. This trajectory-dependent bias may contribute to an increased error rate in the task. To quantitatively examine this hypothesis, we consider a biologically plausible decoding strategy, a quadratic decoder, which exploits pairs of nearly synchronous spikes to perform the task. We train the quadratic decoder on simulated spike trains, generated by a model retina in response to the vernier stimulus, amid fixational drift. We find that a simple, feed-forward model retina which includes spatio-temporal filtering by ganglion cell receptive fields is insufficient to explain the contrast paradox. However, a previously reported component of retinal processing (DW Crevier, M Meister, J Neurophysiol, 1998), the ‘contrast-gain control’, generates a trajectory and contrast dependent shift in the apparent location of the bars. The contrast paradox is reproduced by the quadratic decoder when taking this nonlinear effect into account.

I-80. Distributed polarization of sensorimotor cortex by tDCS modulates functional population coding in macaques

Andrew Bogaard[1] ABOGAARD@GMAIL.COM
Guillaume Lajoie[2] GUILLAUME.LAJOIE@GMAIL.COM
Hayley Boyd[3] HMBOYD@UW.EDU
Andrew Morse[3] AGM22@UW.EDU
Stavros Zanos[4] STAVROSZANOS@GMAIL.COM
Eberhard Fetz[3] FETZ@UW.EDU

[1]Swiss Federal Institute of Technology in Lausanne
[2]Universite de Montreal
[3]University of Washington
[4]Feinstein Institute for Medical Research

Transcranial direct current stimulation (tDCS) is a form of neuromodulation that has been under continuous development for nearly two decades. Its appeal lies in its simplicity and non-invasiveness, requiring only two scalp electrodes to deliver a weak electric current that can produce physiological [1,2], behavioral [3], and therapeutic [4] effects. Despite this, tDCS has not been adopted outside of the lab because results are variable and the mechanism of action is not well understood, especially at the circuit level.

The most common theory posits that intracranial current from tDCS directly polarizes neurons in the brain. Depending on the pattern of depolarization and hyperpolarization in each neuron, various compartments may undergo a modest ($<$1mV) change in membrane potential [5], leading to small biases in spike timing or frequency. These small effects, simultaneously occurring across large areas of brain, are hypothesized to propagate and “amplify” through network interactions to produce meaningful changes in the spatiotemporal patterns of cortical ensembles, although this has not been demonstrated.

To quantify the nature of possible changes in cortical spiking, we recorded from populations of neurons in sensorimotor cortex before, during, and after tDCS. We found changes in neural coding at the population level without changes in the tuning of single neurons. About 15% of recorded neurons demonstrated reliable firing rate modulation during tDCS, and the responses of putative pyramidal cells were sensitive to polarity, whereas responses of non-pyramidal cells were not. We used dimensionality reduction techniques to uncover ensemble dynamics defined by manifolds present in joint spiking activity [6–8], and developed novel metrics to quantify changes in these spatiotemporal dynamics. These analyses revealed that tDCS changed population coding and expanded neural ensembles. Overall, these results are consistent with the notion that tDCS alters brain activity through a combination of direct cell modulation and network interactions.

I-81. A cortico-collicular circuit for decision maintenance in memory-guided behavior

Chunyu Duan[1] ANNDUAN2@GMAIL.COM
Yuxin Pan[1] YUXINPAN@ION.AC.CN
Signe Fruekilde[2] SIGNEFRUEKILDE2@HOTMAIL.COM
Taotao Zhou[1] CPUZTT@ION.AC.CN
Ning-long Xu[1] XUNL@ION.AC.CN

[1]Chinese Academy of Sciences
[2]Sino-Danish Center for Education and Research

Maintenance of decision-related information via short-term memory bridges past events with future actions. The neural circuits underlying this fundamental cognitive process include multiple cortical and subcortical structures, such as the frontal motor cortex, thalamus, basal ganglia, and the midbrain superior colliculus (SC). The frontal motor cortex and the SC both exhibit motor planning-related activity, and have been proposed to be important for short-term memory. But the precise causal contribution of the cortico-collicular pathway to decision maintenance remains unclear. Due to the complex interconnectivity between these regions, information may be represented in a distributed and redundant fashion; thus studying one region at a time cannot adequately constrain the circuit logic of short-term memory. It is therefore essential to identify the causal roles of well-defined pathways during memory-guided behaviors. Here, we conducted circuit-level investigation of the projections from the secondary motor cortex (M2) to SC in mice performing an auditory discrimination task with delayed responses. We combined parametric behavioral task, optogenetic and chemogenetic inactivations to reveal that not only are M2 and SC both key nodes in the motor planning circuit, the direct projections from M2 to SC are required for decision maintenance. Two-photon calcium imaging of SC-projecting M2 neurons shows choice-related information during the memory and response period, both on the single-neuron level and on the population level. Furthermore, direct monitoring of M2 axonal activity in the SC using fiber photometry reveals a progressive recruitment of the M2-SC pathway as memory demand increases. Together, our data suggest that the M2-SC pathway contains choice-related information that is essential for decision maintenance, providing critical constraints for the circuit architecture of short-term memory models.

I-82. Distinct prefrontal neural ensembles contribute to solving the explore/exploit tradeoff

Fanny Cazettes FANNY.CAZETTES@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen ZMAINEN@NEURO.FCHAMPALIMAUD.ORG
Vertechi Pietro PIETRO.VERTECHI@NEURO.FCHAMPALIMAUD.ORG
Masayoshi Murakami MASAYOSHI.MURAKAMI@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Centre for the Unknown

Prefrontal cortical regions are thought to orchestrate the selection and timing of goal-directed actions by representing internal estimates of environmental statistics. Here, we sought to investigate how neural dynamics in prefrontal circuits unfold and interact to support an explore/exploit decision. We first developed a foraging task in head-fixed mice wherein reward delivery was probabilistic and also depleted stochastically. Thus, in order to maximize reward intake, animals had to infer the value of the site (fresh or depleted) from the history of rewards and omissions. Consistent with the predictions of an inference model, mice could accumulate evidence of depletion during unrewarded attempts, at a rate set by the statistics of the environment, to successfully exploit a resource site and to leave when the given site was depleted. Next, we recorded cortex-wide neural activity with Neuropixels probes to dissect the role of two prefrontal regions thought to be involved in evidence accumulation. Semi-automatic sorting algorithms yielded well-isolated spiking activity from hundreds of neurons recorded simultaneously in the secondary motor cortex (M2) and in the orbitofrontal cortex (OFC). In both regions, we observed a reorganization of response correlations between the period of evidence accumulation and the running epoch, suggesting that the same neural populations acted as separate circuits during exploitation and exploration. Using dimensionality reduction methods we found that the time spent exploiting resources could be well decoded from the activity of M2. In contrast, OFC population did not contain as much information about exploitation time. Remarkably, the response dynamics in M2 also correlated with a graded value of evidence accumulation and captured the scalar invariance of exploitation time distributions predicted by the normative model. Altogether, these results revealed a rich representation of evidence in prefrontal circuitry and put M2 forward as a candidate accumulator for the inference of hidden states.

I-83. Inheritance and amplification of grid-cell activity in the medial entorhinal cortex

Tiziano D'Albis TIZIANO.DALBIS@HU-BERLIN.DE
Richard Kempter RICHARD.KEMPTER@BIOLOGIE.HU-BERLIN.DE
Humboldt-Universitat zu Berlin

Grid cells are neurons of the medial entorhinal cortex (MEC) that are tuned to the animal's location in the environment and whose firing fields form a regular triangular pattern (Hafting et al., 2005). Grid-cell activity is thought to support animals' navigation and spatial memory, but the fundamental principles that govern its dynamics are still debated.

Recent experiments suggest that grid cells belong to distinct principal-cell populations of the MEC (Sun et al., 2015; Rowland et al., 2018), which are embedded in excitatory feedforward and recurrent circuits (Fuchs et al., 2016; Winterer et al., 2017). Yet it remains unclear how grid-cell activity is affected by these connections. Do they disrupt or rather amplify grid patterns? Can grids be inherited via feedforward projections? And how are they affected by recurrent network dynamics?

We address these questions theoretically. First, we focus on feedforward networks. We hypothesize that grids originate in an upstream MEC region via a single-cell mechanism (e.g., Kropff and Treves, 2008; D'Albis and Kempter, 2017), and that projections to downstream neurons are plastic according to a Hebbian rule. In this scenario, we show that a structured synaptic connectivity spontaneously emerges in the network and, as a result, grid-cell activity propagates to downstream cells. We further show that, after learning, both feedforward and recurrent circuits can amplify grid patterns, i.e., noisy grid inputs become more spatially-regular at the output. The amplification strength is derived analytically as a function of the connection strength and the properties of the input tuning.

Considering the layout of the MEC microcircuitry, we propose that layer II stellate cells inherit grid tuning from layer II pyramidal cells and that grid-cell activity is amplified locally via structured recurrent connections. The present work sheds light on the role of local excitatory circuits in grid-cell activity.

I-84. Persistent activity encodes variables for flexible decision making

Bilal Bari BBARI1@JHMI.EDU
Cooper Grossman CGROSSM4@JHMI.EDU
Emily Lubin ELUBIN1@JHU.EDU
Adithya Rajagopalan ARAJAGO6@JHMI.EDU
Jianna Cressy JCRESSY1@JHU.EDU
Jeremiah Cohen JEREMIAH.COHEN@JHMI.EDU
Johns Hopkins University

Decisions take place in dynamic environments. The nervous system must continually learn the best actions to obtain rewards. In the theoretical framework of optimal control and reinforcement learning, policies (the probability of performing an action given a state of the environment) are updated by feedback arising from errors in the predicted reward. Whereas these reward prediction errors have been mapped to dopaminergic neurons in the midbrain, how the decision variables that generate policies themselves are represented is unclear. Here, we trained mice on a dynamic foraging task in which they freely chose between two alternatives that delivered reward with changing probabilities. Mice exhibited flexible behavior, using recency-weighed reward history to both select actions and influence response times, suggesting a common generative process. To model this process, we adapted an action-value-based reinforcement-learning model. This model generated two decision variables – relative value (the difference between action values) to bias choices and total value (the sum of action values) to bias response times. We found excellent agreement between the model and real behavior. To determine where these decision variables may be represented, we reversibly inactivated the medial prefrontal cortex, a region implicated in flexible decision making. This manipulation prevented mice from updating actions adaptively and slowed response times, consistent with an ablation of relative and total value representations. Experiments using other behavioral tasks revealed neither deficit was due to a motor impairment. We next recorded action potentials from 3,045 mPFC neurons in ten mice performing the foraging task. Most neurons (>80%) maintained persistent changes in firing rates that represented relative and total action values over long timescales. Furthermore, we found that corticostriatal neurons encoded this information, suggesting a basal ganglia circuit mechanism by which these decision variables may modulate actions. Thus, we define a stable neural mechanism used to drive flexible behavior.

I-85. Coordinated grid and place cell replay as offline spatial inference

Talfan Evans TALFAN.EVANS.13@UCL.AC.UK
Neil Burgess N.BURGESS@UCL.AC.UK
University College London

Effective navigation can rely on two sources of information. Firstly, unique sensory inputs may indicate absolute position in space. Secondly, path integration (PI) can be used to update a previous estimate by integrating perceived movement. An optimal system performing spatial navigation would treat these sources of information as complementary. In general, simultaneously learning and navigating under uncertainty requires one to maintain a probability distribution over possible states, given previous observations and movements (SLAM). Neurobiologically, entorhinal grid cells (GC) are thought to perform PI, whereas hippocampal place cells (PC) are thought to be driven predominantly by sensory cues.

Here, we present a novel algorithmic and mechanistic account of probabilistic spatial localization and learning, mediated by grid and place cells. Firstly, GCs represent arbitrary belief distributions over the current location in their population firing rates, but nonetheless show ‘attractor-like’ dynamics in their relative spatial phases. Secondly, weighing sensory input from PCs with PI information allows GCs to both drive PC remapping and embed local environments in a global structure. Thirdly, our simple model predicts characteristic distortions to the grid pattern, which are confirmed by a novel re-analysis of existing data. These distortions are interpreted as a salience-driven warping of the metric cognitive map.

Fourthly, we show that place and grid cells could perform structural inference via a message passing process, which generates structured reactivations of place and grid cells resembling those observed during coordinated hippocampal/entorhinal replay. Our model poses this hypothesis on both an algorithmic and neural level, the latter making several experimentally testable predictions. The overall framework proposes a dual-systems theory of spatial localization and learning in rodents. Under low computational demand, an online system may support cheap and efficient localization and simple learning. However, prediction errors trigger an offline inference process in more demanding or dynamic environments.

I-86. Dendritic NMDA receptors in parvalbumin neurons enable strong and stable neuronal assemblies

Jonathan Cornford JONATHAN.CORNFORD@GMAIL.COM
Marion Mercier M.MERCIER@UCL.AC.UK
Marco Leite MARCO.LEITE.11@UCL.AC.UK
Vincent Magloire V.MAGLOIRE@UCL.AC.UK
Michael Hausser M.HAUSSER@UCL.AC.UK
Dimitri Kullmann D.KULLMANN@UCL.AC.UK
University College London

Fast-spiking, parvalbumin-expressing (PV+) GABAergic interneurons mediate feedforward and feedback inhibition and have a key role in gamma oscillations and information processing. The importance of fast PV+ cell signalling for these circuit functions is well established. In contrast, the adaptive significance of NMDA receptors (NMDARs), which generate relatively slow postsynaptic currents, in PV+ cells is unclear. Underlining their potential importance, several studies implicate PV+ cell NMDAR disruption in pathology and impaired network function. Here, we show that NMDARs underlie supralinear integration of feedback excitation from local pyramidal neurons onto mouse CA1 PV+ cells. Furthermore, by incorporating NMDARs at feedback connections onto PV+ cells in spiking neural networks, we show that these receptors enable cooperative recruitment of PV+ interneurons, strengthening and stabilising principal cell assemblies. Failure of this phenomenon provides a parsimonious explanation for cognitive and sensory gating deficits in pathologies with impaired PV+ NMDAR signalling

I-87. Activity-dependent synaptic plasticity shapes dendritic synaptic inputs on developing cortical neurons

Jan Hendrik Kirchner JAN.KIRCHNER@BRAIN.MPG.DE
Julijana Gjorgjieva GJORGJIEVA@BRAIN.MPG.DE

Max Planck Institute for Brain Research

Most neural circuits experience minimal sensory stimulation during early postnatal development, yet they are capable of spontaneously generating highly structured activity even before the maturation of sensory organs. Spontaneous activity can shape circuit connectivity down to the fine-scale organization of synapses on dendritic trees. In fact, synapses in the developing mouse cortex are functionally clustered, where nearby synapses at most 15 microns apart tend to be more coactive than synapses farther away from each other [Kleindienst et al, Neuron, 2011].

We propose a biologically realistic, local synaptic plasticity mechanism of activity-dependent cooperation and competition that can induce the functional clustering of synapses based on patterns of spontaneous activity. The model is based on the interaction between two neurotrophic factors, BDNF and proBDNF, which determine the amounts of synaptic depression and potentiation, respectively. We mapped the biophysical model to a phenomenological rule known as burst-timing-dependent plasticity (BTDP), shown to operate in the developing visual system. Introducing a spatial component and a state-dependent mechanism of synaptic turnover in the model led to the depression of nearby synapses that received desynchronized input; these were eventually were replaced by new synapses. In agreement with experiments, distant or well-synchronized synapses did not depress. To further investigate the factors underlying synaptic competition vs. cooperation, we implemented orientation selectivity at the level of individual synapses by associating them with Gabor receptive fields of varying orientation. When stimulating the synapses with realistic retinal wave input, an initially high rate of synaptic turnover decreased exponentially until the system reached a steady state where nearby synapses shared a similar orientation preference. The resulting synaptic clusters became more homogeneous when we implemented backpropagating action potentials, thus providing a possible mechanistic explanation for a related experimental observation in the ferret visual cortex that distinguished between clustered and heterogeneous branches [Wilson et al, Nat Neur, 2016].

I-88. Inferring the dynamics of neural populations from single-trial spike trains using mechanistic models

Christian Donner[1] CHRISTIAN.DONNER@BCCN-BERLIN.DE
Josef Ladenbauer[2] JOSEF.LADENBAUER@GMAIL.COM
Manfred Opper[1] MANFRED.OPPER@TU-BERLIN.DE

[1]Technische Universitat Berlin
[2]Ecole Normale Superieure

Multi-neuronal spike-train data recorded in vivo often exhibit rich dynamics as well as considerable variability across cells and repetitions of identical experimental conditions (trials). Efforts to characterize and predict the population dynamics and the contributions of individual neurons require model-based tools. Abstract statistical models allow for principled parameter estimation and model selection, but possess only limited interpretive power because they typically do not incorporate prior biophysical constraints. Here we present a statistically principled approach based on a population of doubly-stochastic integrate-and-fire neurons, taking into account basic biophysics. This model class comprises an idealized description for the dynamics of the neuronal membrane voltage in response to fast independent and slower shared input fluctuations. To efficiently estimate the model parameters and compare different model variants we compute the likelihood of observed single-trail spike trains by leveraging analytical methods for spiking neuron models combined with inference techniques for hidden Markov models. This allows us to reconstruct the shared input variations, classify their dynamics, obtain precise spike rate estimates, and quantify how individual neurons couple to the low-dimensional overall population dynamics, all from a single trial. Extensive evaluations based on simulated data show that our method correctly identifies the dynamics of the shared input process and accurately estimates the model parameters. Validations on ground truth recordings of neurons in vitro demonstrate that our approach successfully reconstructs the dynamics of hidden inputs and yields improved fits compared to a typical phenomenological model. Finally, we apply the method to a neuronal population recorded in vivo, for which we assess the contributions of individual neurons to the overall spiking dynamics. Altogether, our work provides statistical inference tools for a class of reasonably constrained, mechanistic models and demonstrates the benefits of this approach to analyze measured spike train data.

I-89. One step back, two steps forward: interference and learning

Chen Beer CHEN.COHEN.BEER@GMAIL.COM
Omri Barak OMRI.BARAK@GMAIL.COM

Technion

Artificial neural networks, trained to perform cognitive tasks, have recently been used as models for neural recordings from animals performing these tasks. While some progress has been made in performing such comparisons, the evolution of network dynamics throughout learning remains unexplored. This is paralleled by an experimental focus on recording from trained animals, with few studies following neural activity throughout training. In this work, we address this gap in the realm of artificial networks by analyzing networks that are trained to perform memory and pattern generation tasks. The functional aspect of these tasks corresponds to dynamical objects in the fully trained network – a line attractor or a set of limit cycles for the two respective tasks. We use these dynamical objects as anchors to study the effect of learning on their emergence. We find that the sequential nature of learning – one trial at a time – has major consequences for the learning trajectory and its final outcome. Specifically, we show that Least Mean Squares (LMS), a simple gradient descent suggested as a biologically plausible version of the FORCE algorithm (Sussillo & Abbott, 2009), is constantly obstructed by forgetting, which is manifested as the destruction of dynamical objects from previous trials. We show the specific ingredients of FORCE that avoid this phenomenon. For LMS, we show that the correlation between different trials leads to destructive interference. Overall, this difference results in convergence that is orders of magnitude slower for LMS. Learning implies accumulating information across multiple trials to form the overall concept of the task. Our results show that interference between trials can greatly affect learning, in a learning rule dependent manner. These insights can help design experimental protocols that minimize such interference, and possibly infer underlying learning rules by observing behavior and neural activity throughout learning.

I-90. A biologically plausible neural network for semi-supervised learning on a manifold

Alexander Genkin[1] ALEXANDER.GENKIN@GMAIL.COM
Anirvan Sengupta[2] ANIRVANS@PHYSICS.RUTGERS.EDU
Dmitri Chklovskii[3] DCHKLOVSKII@FLATIRONINSTITUTE.ORG

[1]New York University
[2]Rutgers University
[3]Flatiron Institute and NYU Langone Medical Center

As large labeled datasets are rarely available in real-life biological learning, we expect the brain to utilize unlabeled data as much as possible. Whereas early stages of sensory processing can be modeled using fully unsupervised learning, later stages incorporate supervision signal whenever it is available. Here, we propose a two-layer biologically plausible network where the first layer learns low-dimensional manifolds in the data streamed by sensory organs [1] and the second classifies data using both occasional supervision and similarity of the manifold-tiling representation of the input data. Whereas online methods of semi-supervised learning on manifolds exist, they lack biological plausibility. Here we suggest a simple algorithm for a semi-supervised neuron, which, in addition to input data, also receives a channel carrying label information. This could be a neuron with a dendritic and a somatic compartments, where the label channel connects directly to the soma [2]. The label channel is assumed to be ”silent” most of the time, making it a truly semi-supervised setup. Addition of unlabeled data to supervised learning allows using fewer labeled examples and accelerates learning. We are taking a normative approach and formulate an objective function combining the loss on the label prediction with the similarity preserving term [3] between the manifold tiling representation and the output. The online algorithm derived from this objective maps naturally onto a neural network with a local Hebbian learning rule. We think of manifold tiling channels as vertices and correlations between them as weighted edges of the graph representing the manifold. The similarity term allows the label information to propagate along the graph edges without the graph being represented explicitly, so the correct classifier can be learned with very few labeled data points. Our algorithm learns to classify data points from two interleaving but non-overlapping arcs from only two labeled samples.

I-91. A neurally inspired circuit model for human sensorimotor timing

Seth Egger SWEGGER@MIT.EDU
Mehrdad Jazayeri MJAZ@MIT.EDU
Nhat Le NMLE@MIT.EDU

Massachusetts Institute of Technology

The richness of our behavioral repertoire is critically dependent on our ability to flexibly coordinate movement timing with internal and external cues. Recent physiological recordings indicate that this flexibility is mediated by controlling the speed at which neural activity evolves over time. Mechanistically, this can be achieved with a simple circuit module consisting of two pools of neurons that reciprocally inhibit each other and receive a shared input. This input adjusts the latent dynamics of the system and, therefore, acts as a dial to control the speed at which output evolves over time. When output is the difference between the two pools of neurons, a ramp to threshold with adjustable rate is created, allowing input to flexibly control movement initiation time. Here we show that this circuit motif may serve as a core module of a sensorimotor controller by acting not only to govern motor timing, but also to implement forward prediction (i.e. simulation) and feedback control. Combining these elements allows a single circuit to implement a variety of timing and temporal coordination tasks including time interval production, rhythmic production, synchronization-continuation, and Bayesian time interval estimation. Further, the model emulates nontrivial aspects of human behavior in these tasks, suggesting the circuit model might capture the essence of the mechanisms that support interval timing behavior.

I-92. Thalamic drive of cortical interneurons during down states

Pedro Lagomarsino de Leon Roig PEDRO.LAGOMARSINO@IIT.IT
Stefano Zucca S.ZUCCA@UCL.AC.UK
Valentina Pasquale VALENTINA.PASQUALE@IIT.IT
Stefano Panzeri STEFANO.PANZERI@IIT.IT
Tommaso Fellin TOMMASO.FELLIN@IIT.IT

Istituto Italiano di Tecnologia

Slow oscillations are among the most prominent brain dynamics during NREM sleep and many forms of anesthesia. Recent evidence demonstrates that cortical inhibitory cells play a major role in the modulation of up and down state transitions during such oscillations. More specifically, parvalbumin-positive fast spiking interneurons discharge action potentials also during cortical down states and this gabaergic signaling is associated with prolonged down state durations. However, what drives these inhibitory cells to fire during down states, when excitatory neurons are largely silent, remains unclear. Here, we performed two-photon guided juxtasomal recordings from PV interneurons in the barrel field of the somatosensory cortex (S1bf) of anesthetized mice, while simultaneously collecting the local field potential (LFP) in S1bf and the multi-unit activity (MUA) in the ventral posteromedial (VPM) thalamic nucleus. We found that activity in the VPM was associated with longer down state duration and that PV firing during down states was temporally correlated with higher VPM activity. Moreover, pharmacological inactivation of the thalamus significantly reduced the percentage of PV spikes observed during down states while significantly increasing the locking of PV spikes to cortical up states. Finally, we optogenetically inactivated PV cells during down states and we found increased latency of the light-evoked down-to-up transitions upon thalamic pharmacological blockade compared to control conditions. These results contribute to dissect out the circuit mechanisms underlying the regulation of up and down state transitions during slow oscillations and demonstrate that the thalamus is a potent drive of cortical interneurons during down states.

I-93. Task-irrelevant plasticity constrains optimal neural network size

Dhruva Raman DHRUVARAMAN@GMAIL.COM
Adriana Perez-Rotondo AP2013@CAM.AC.UK
Timothy O'Leary TSO24@CAM.AC.UK

University of Cambridge

During learning, information on task performance is somehow converted into local plasticity signals at synapses. However, in a biological network, synapses are often involved in multiple tasks and subject to biological processes besides learning. We therefore expect that for a given task, a significant component of synaptic change during learning is not relevant to that task. In this work we show that overall network size can directly overcome some kinds of task-irrelevant synaptic changes, but not others. For arbitrary learning rules, we show that plasticity at a synapse can be decomposed into three components reflecting different underlying sources: 1. task-relevant plasticity 2. task-irrelevant plasticity due to imperfect learning rules, or ongoing plasticity unrelated to the task 3. intrinsic noise in synaptic strength due to the stochasticity of molecular mechanisms We show mathematically that the performance of any learning rule can be quantified by the ratio of source 1 to source 2. For any given quality of learning rule (in the absence of source 3), we show that increasing network size generically improves learning rate and steady state task performance in a quantifiable way. In- formally, increasing network size in an appropriate way 'flattens out' the error surface, facilitating learning and reducing impact of noise. Therefore networks can buffer the effect of imperfect, biologically plausible learning rules by increasing the number of neurons involved in learning the task. However, when source 3 is present, as we expect biologically, excessively large networks are penalised. Thus two competing factors (sources 2 and 3) determine an optimal network size. We calculate this optimum for nonlinear, feedfor- ward networks, providing a template for calculating optima of generic networks. Our analysis suggests that overall network size is constrained by the need to learn efficiently with unreliable synapses, and provides a rigorous theory of how network size impacts learning performance.

I-94. Emergence of spatiotemporal sequences in spiking neuronal networks with spatial connectivity

Sebastian Spreizer[1] SEBASTIAN.SPREIZER@BCF.UNI-FREIBURG.DE
Ad Aertsen[2] AERTSEN@BIOLOGIE.UNI-FREIBURG.DE
Arvind Kumar[3] ARVKUMAR@KTH.SE

[1]Bernstein Center Freiburg, Faculty of Biology, University of Freiburg, Germany
[2]Bernstein Center Freiburg, University of Freiburg
[3]KTH Royal Institute of Technology, Stockholm

Spatio-temporal sequences (STS) of neuronal activity are observed in many brain regions in a variety of tasks and are thought to form the basis of meaningful behavior. However, mechanisms by which a neuronal network can generate STS have remained obscure. Existing models are biologically untenable because they either require manual embedding of a feedforward network within a random network or supervised learning to train the connectivity of a network to generate sequences. Here, we propose a biologically plausible, generative rule to create STS in a network model of spiking neurons with distance-dependent connectivity. We show that the emergence of STS requires: (1) individual neurons preferentially project a small fraction of their axons in a specific direction, and (2) the preferential projection direction of neighboring neurons is similar. Thus, an anisotropic but correlated connectivity of neurons suffices to generate STS in an otherwise random neuronal network model.

I-95. Decoding spatial information from local field potentials in the human MTL

Nora Herweg[1] NHERWEG@SAS.UPENN.EDU
Michael Kahana[1] KAHANA@PSYCH.UPENN.EDU
Paul A Wanda[1] PWANDA@SAS.UPENN.EDU
Lukas Kunz[2] LUKAS.KUNZ@UNIKLINIK-FREIBURG.DE
Armin Brandt[2] ARMIN.BRANDT@UNIKLINIK-FREIBURG.DE
Michael R Sperling[3] MICHAEL.SPERLING@JEFFERSON.EDU
Ashwini D Sharan[3] ASHWINI.SHARAN@JEFFERSON.EDU
Andreas Schulze-Bonhage[2] ANDREAS.SCHULZE-BONHAGE@UNIKLINIK-FREIBURG.DE

[1]University of Pennsylvania
[2]University Medical Center Freiburg
[3]Thomas Jefferson University

The activity of cell-populations in the MTL represents an animal's trajectory through space. These representations are thought to support spatial navigation as well as memory for the spatial contexts associated with specific events. To characterize spatial representations at different stages of a spatial-episodic memory task, we recorded local field potentials (LFPs) from micro-wires in the MTL of 19 patients with medication-resistant epilepsy. Participants navigated through a virtual town consisting of target stores and non-target buildings to deliver objects to a sequence of target stores. We used multinomial logistic regression to decode current location and goal information from LFPs (power in 10 frequencies between 3 and 200 Hz). We obtained significant classification of both location and goal in a subset of participants (68 and 42%, respectively) and across the group ($p<0.01$). Accuracy of both, location and goal decoding increased over time, suggesting that classification is based on a representation of the environment that is acquired through experience. Equipped with a tool to uncover latent location and goal representations during navigation, we asked when and where specific representation are activated during task performance. Specifically, we explored classifier output probability for all stores as a function of spatial, temporal and semantic distance between them. We find that location and goal representations are mostly distinct in physical and semantic space. In temporal space, location representations linger after visiting a certain location and goal representations are activated in a sustained way during navigation. Taken together, these results provide compelling evidence that neuronal populations in the human MTL represent local and distant spatial locations at

different temporal scales in a reference frame that is acquired through experience to support goal-oriented navigation. From a general perspective, they further highlight the utility of decoding techniques as a tool to visualize latent cognitive representations during complex tasks.

I-96. Coordination of inhibitory Golgi cell population activity in the cerebellar cortex

Harsha Gurnani HARSHAGURNANI@GMAIL.COM
R Angus Silver A.SILVER@UCL.AC.UK

University College London

Sensorimotor encoding by cerebellar granule cells (GCs) is important for downstream associative learning and motor control. Golgi cells (GoCs) provide both feedforward and feedback inhibition onto GCs and can regulate their excitability and spike timing, shaping their population response. Theoretical work predicts that parallel fibre (PF) and mossy fibre (MF) inputs on GoCs, and electrical synapses between GoCs, have a differential effect on GoC synchrony [Maex 1998, Vervaeke 2010]. However, it has not been possible to study GoC networks experimentally, due to their sparse distribution in the input layer. We present the first GoC population recording in awake animals. We performed calcium imaging of sparsely distributed, GCaMP6f-expressing GoCs in awake, head-fixed mice using high-speed acousto-optic lens 3D two-photon microscopy [Nadella 2016]. GoCs in Crus I/II region revealed strong, coherent activation across the recorded population (20-70 cells/region) during spontaneous whisking epochs. This is consistent with a net increase in MF and PF activity observed under similar conditions (Ros and Lanore, unpublished observations). This broad increase in GoC population activity is consistent with network activity-dependent thresholding of GCs that has previously been shown to improve decorrelation of population responses [Cayco-Gajic 2017]. Pairwise correlations of GoC activity were high across distances of at least $200\mu m$. Despite this slow-timescale (hundreds of milliseconds) GoC network coherence, nearby GoCs did not necessarily show synchronous activity. Lagged and anti-phasic modulation superimposed on an envelop of activation within whisking epochs, was observed. Thus, GoC populations show local modulation superimposed on network-wide signals, suggesting a mixed role of shared and tuned inhibition on potentially different aspects of sensorimotor processing. Global inhibition performing adaptive gain control can support pattern separation in the GC layer by increasing decorrelation and sparsening of GC activity [Billings 2014, Cayco-Gajic 2017], while functional inhibitory subnetworks may shape temporal and plasticity profile of GCs.

I-97. Tactile sensory specialization in the secondary somatosensory cortex

Matias A Goldin[1] GOLDIN@UNIC.CNRS-GIF.FR
Evan Harrell[2] EVAN.HARRELL@UNIC.CNRS-GIF.FR
Luc Estebanez[1] ESTEBANEZ@UNIC.CNRS-GIF.FR
Daniel E Shulz[3] SHULZ@UNIC.CNRS-GIF.FR

[1]Unite de Neurosciences Information et Complexite
[2]CNRS
[3]CNRS – UNIC – UPR 3293

Tactile perception in rodents depends on simultaneous, multi-whisker contacts with objects. Although it is known that neurons in secondary somatosensory cortex (wS2) respond to individual deflections of many whiskers, wS2‘s precise function remains unknown. The convergence of information from multiple whiskers into wS2 neurons suggests that they are good candidates for integrating multi-whisker information. Here, we apply stimulation patterns with rich dynamics simultaneously to 24 macro-vibrissae of rats while recording large populations of single neurons in wS2 and in primary somatosensory cortex (wS1). We first analyze the responses of the neurons to a stimulus in which all whiskers receive deflections sampled from the same Gaussian white noise (correlated

GWN). We apply a spike-triggered covariance (STC) approach to obtain the whisker movements, or filters, responsible for eliciting the most spikes. These filters are well represented by low dimensional filter subspaces specific for wS2 and wS1. We then developed the Temporal Projection Method (TPM), which uncovers extended stimulus dependencies across time. TPM is based on a change of coordinates that allows a dimensionality reduction and it proved to be necessary to capture the functional responses of wS2 neurons for a stimulus where each whisker receives a unique uncorrelated GWN. With our new method, we observe pronounced supra-linear multi-whisker integration and we show that continuous multi-whisker movements contribute to the firing of wS2 neurons over long temporal windows, facilitating spatio-temporal integration. In contrast, wS1 neurons encode fine features of whisker movements on precise temporal scales. These results provide the first description of wS2's representation during continuous multi-whisker stimulation and outline its specialized role in parallel to wS1 tactile processing.

I-98. Rapid spatiotemporal coding in trained multi-layer and recurrent spiking neural networks

Friedemann Zenke FRIEDEMANN.ZENKE@CNCB.OX.AC.UK
Tim P Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Our senses acquire large amounts of information which needs to be processed rapidly. To achieve this feat, our nervous system relies on spikes for rapid encoding, filtering and transmission of information. Experimental work has shown that, while the number of spikes involved in this computation can be surprisingly low, spike timing is often precise. Thus neural processing may proceed in large part by dint of rapid, sparse temporal codes. Yet, how such a coding scheme can be employed to solve specific computational problems remains largely elusive. A major impediment to understanding the neural code has been the lack of suitable network models that both perform specific computations and exhibit spiking activity comparable to circuits in neurobiology.

Here we address this shortcoming by training spiking neural networks on a range of spatiotemporal tasks. Specifically, we use supervised loss functions to train spiking neural networks and rely on recently developed surrogate gradient techniques to overcome the difficulties typically associated with their training. Surrogate gradients are smoothly differentiable continuous relaxations of the true gradient which permit training spiking neural networks using standard machine learning tools. We illustrate the effectiveness of this approach by training multi-layer and recurrent spiking classifiers on a variety of real-world and synthetic temporal data problems. Specifically, we demonstrate that spiking neural networks can solve standard machine learning benchmarks and accurately classify spikes generated from spoken words fed through a cochlea model. Finally, using synthetic data, instructed by the latent structure of smooth random manifolds, we characterize the robustness of spiking neural networks to noise and their ability to generalize at different levels of sparsity and temporal precision.

Our work demonstrates not only how to build functional spiking neural networks, but also gives a glimpse of the extraordinary computing power of spiking neural networks operating in the sparse spiking regime.

I-99. A basis set of operations captures cell assembly dynamics under basal conditions and during learning

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

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

[3]Johannes Gutenberg University Mainz

In the mouse auditory cortex population response patterns to short sounds cluster into a near discrete set of response modes that predict behavioral discrimination [Bathellier et al, 2012]. These response modes are formed by the co-activity of sparse cell assemblies displaying winner-takes-all-like dynamics. Such cell assemblies are believed to form basic functional units within cortical circuits. Therefore, stable brain function could simply rely on the stability of such cell assemblies. However, with the advent of technologies allowing chronic measurements of synaptic connections [Loewenstein et al, 2015] and functional tuning of neurons [Driscoll et al, 2017], the emerging picture is that cortical networks undergo significant remodeling over multiple timescales, even in the absence of experimenter-induced learning or sensory perturbation. Here, we studied the long-term dynamics of cell assemblies in the mouse auditory cortex under basal conditions and the impact of auditory cued fear conditioning, a widely-used associative learning paradigm. Using chronic 2-photon microscopy in layer 2/3, we imaged 38388 neurons expressing GCaMP6m from 20 mice. We find that cell assemblies show significant remodeling across the time span of days under basal conditions. While some response modes are stable, others emerge or disappear. We identified a basis set of ten operations capturing all transitions between imaging sessions. These operations cover remapping of stimuli between persistent cell assemblies as well as cell assembly dynamics. Deploying the framework to dissect the effects of learning in mice undergoing auditory cued fear conditioning, we observe an increase in transitions mapping stimuli together onto one cell assembly as well as a decrease in transitions separating stimuli that were previously mapped together. Our findings show that cortical representations of stimuli are dynamically organized under basal conditions and that learning weaves into these dynamics by biasing them towards the formation of new associations and stabilization of existing associations.

I-100. Short-term plasticity as a spatial distribution filter

Luiz Tauffer TAUFFER@KTH.SE
Arvind Kumar ARVKUMAR@KTH.SE

KTH Royal Institute of Technology, Stockholm

How could a single neuron, reading out the noisy activity of a large presynaptic population, discriminate between inputs of the same intensity (number of spikes in a period of time) but with different spatial arrangements? How could this postsynaptic neuron distinguish, for example, a burst of spikes coming from the same presynaptic neuron encoding for a particular stimulus from an equal number of spikes generated by different presynaptic neurons due to fluctuations of the background activity? We describe how a population of synapses with short-term plasticity (STP) could provide a solution to this, acting as a filter for the spatial distribution of the population activity. We show how STP modulates the output of presynaptic populations as a function of the spatial distribution of a given extra spiking activity and find a strong relationship between STP features and sparseness of the population code. Furthermore, we show that feedforward excitation followed by inhibition (FF-EI), combined with target-dependent STP, promote substantial increase in the filter gain even for considerable deviations from the optimal conditions, granting robustness to the filter system. A simulated neuron driven by a spiking FF-EI network is reliably modulated as predicted by the rate analysis and inherits the ability to differentiate spatially sparse signals from spatially smooth background activity changes of the same magnitude, even at very low signal-to-noise conditions. We propose that the STP-based spatial distribution filter is likely a latent functionality serving several brain networks such as the granular layer of the cerebellum and the Schaffer collaterals of the hippocampus.

I-101. Amortised inference for mechanistic models of neural dynamics

Jan-Matthis Lueckmann[1,2] JAN-MATTHIS.LUECKMANN@CAESAR.DE
Pedro J Goncalves[1] PPJGONCALVES@GMAIL.COM
Chaitanya Chintaluri[3] CHAITANYA.CHINTALURI@DPAG.OX.AC.UK
William F Podlaski[3] WILLIAM.PODLASKI@CNCB.OX.AC.UK
Giacomo Bassetto[1,2] GIACOMO.BASSETTO@CAESAR.DE
Tim P Vogels[3] TIM.VOGELS@CNCB.OX.AC.UK
Jakob H Macke[2] MACKE@TUM.DE

[1]Research Center Caesar
[2]Technical University of Munich
[3]University of Oxford

Bayesian statistical inference provides a principled framework for linking mechanistic models of neural dynamics with empirical measurements. However, for many models of interest, in particular those relying on numerical simulations, statistical inference is difficult and requires bespoke and expensive inference algorithms. Furthermore, even within the same model class, each new measurement requires a full new inference – one can not leverage knowledge from past inferences to facilitate new ones. This limits the use of Bayesian inference in time-critical, large-scale, or fully-automated applications.

We overcome these limitations by presenting a method for statistical inference on simulation-based models which can be applied in a 'black box' manner to a wide range of models in neuroscience. The key idea is to generate a large number of simulations from the model of interest and use them to train a neural network to perform statistical inference. Once the network is trained, performing inference given any observed data is very fast, requiring only a single-forward pass through the network, i.e. inference is amortised.

We explain how our approach can be used to perform parameter-estimation, and illustrate it in the context of ion channel models. We train a network on a large diversity of simulated current responses to voltage-clamp protocols. After training, the network is able to instantaneously provide the posterior distribution over the channel model parameters given current responses from a publicly available database of ion channel models. The approach will enable neuroscientists to perform scalable Bayesian inference on large-scale data sets and complex models without having to design model-specific algorithms, closing the gap between mechanistic and statistical approaches to neural dynamics.

I-102. Recent reward history impacts ventral pallidum estimates of present and future reward value

David Ottenheimer[1] DAVID.OTTENHEIMER@JHU.EDU
Elissa Sutlief[1] ESUTLIE1@JHMI.EDU
Jocelyn Richard[2] RICHARDJ@UMN.EDU
Patricia Janak[1] PATRICIA.JANAK@JHU.EDU

[1]Johns Hopkins University
[2]University of Minnesota

A critical avenue of inquiry in systems neuroscience is understanding how the brain implements value computations. Valuation of a given outcome relative to all available options is necessary to direct adaptive reward-seeking behavior. In two experiments, we sought to characterize the computations underlying evaluative and predictive neural activity in ventral pallidum (VP)—a basal ganglia output nucleus implicated in the orchestration of reward-seeking behavior—during a task with multiple reward outcomes. In the first experiment, rats responded to a non-predictive white noise cue, triggering the delivery of 10% solutions of either sucrose (preferred) or maltodextrin. We found a prominent value signal following reward delivery related to preference; phasic excitation for sucrose and a reduction in firing for maltodextrin. We then analyzed this signal according to previously received

outcomes using multiple linear regression and found strong modulation according to the most recently received outcome. Moreover, when analyzing the anticipatory neural activity surrounding entry into the reward port, we found a significant impact of each of the previous two trials on the firing rate. In the second task, to probe whether these reward history-sensitive signals reflected expected value, we introduced two additional cues that predicted sucrose and maltodextrin, respectively, with 100% probability (predictive trials) in addition to the original 50/50 cue (non-predictive trials). On non-predictive trials, as before, the anticipatory signal reflected the outcome on the previous two trials. On predictive trials, there was a reduction in coefficient weights for the previous trials and, instead, a strong influence of the cued upcoming reward. Our results illustrate in VP a scheme for representing expected value that integrates previous experience and predictive sensory stimuli. These data highlight VP as an important site to test models of how the brain calculates the subjective value of outcomes using both internally and externally generated information.

I-103. Synaptic plasticity dynamics for deep continuous local learning

Jacques Kaiser[1] JKAISER@FZI.DE
Hesham Mostafa[2] HESHAM@UCSD.EDU
Emre Neftci[3] ENEFTCI@UCI.EDU

[1]FZI Research Center For Information Technology, Karlsruhe
[2]University of California, San Diego
[3]University of California, Irvine

A growing body of work underlines striking similarities between spiking neural networks modeling biological networks and recurrent, binary neural networks. A relatively smaller body of work, however, demonstrate similarities between learning dynamics employed in deep artificial neural networks and synaptic plasticity in spiking neural networks. The challenge preventing this is largely due to the discrepancy between dynamical properties of synaptic plasticity and the requirements for gradient backpropagation. Here, we demonstrate that deep learning algorithms that locally approximate the gradient backpropagation updates using locally synthesized surrogate gradients overcome this challenge. Locally synthesized gradients were initially proposed to decouple one or more layers from the rest of the network so as to improve parallelism. Here, we exploit these properties to derive gradient-based synaptic plasticity rules in spiking neural networks. Our approach results in highly efficient spiking neural networks and synaptic plasticity capable of training deep networks, while being compatible with biological neural networks. Furthermore, our method utilizes existing autodifferentation methods in machine learning frameworks to systematically derive synaptic plasticity rules from task-relevant cost functions and neural dynamics. We benchmark our approach on a event-based DVS Gestures dataset, and report state-of-the-art results. Our results provide continuously learning machines that are not only relevant to biology, but suggestive of a brain-inspired computer architecture that matches the performances of conventional neural network accelerators on target tasks.

I-104. Network dimensionality reflects input at the edge of chaos

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

Technion

The interplay between internal spontaneous activity of neural networks and their response to external inputs is a central problem in Neuroscience. While spontaneous activity presumably arises from network dynamics, it is often overlooked that external inputs are also typically generated by dynamical systems – be it other neural networks, or generative dynamics of the external world. It is thus natural to inquire how these external dynamics reflect in the network. We address this question by studying an ensemble of random networks, defined by connection strength

and degree of symmetry (Sommers et al. 1988, Marti et al. 2018). These properties affect network dimensionality, timescales and stability, thus offering a rich testbed for input interaction. For spontaneous dynamics, a curve on the connectivity-symmetry plane – the 'edge of chaos' (EOC) – sharply separates chaotic from regular behavior. At the same time, network dimensionality varies smoothly over the plane, peaking at intermediate symmetry. We couple the network with scalar signals, arising from dynamical systems with a specified number of degrees of freedom (order). Driving the networks with signals of varying properties generally shifts the EOC (Rajan et al. 2010); however, appropriate normalization reveals collapse of all EOCs onto a universal curve. In contrast, network dimensionality is signal dependent, conforming to the input order along the EOC for all signals. Training the networks to autonomously produce specified signals, we find a limited region of learnability delineated between the spontaneous and driven EOC curves. Within this region, learning is not uniform: it is optimal where network dimensionality most closely matches input order, while performance itself depends on task parameters. Our results point to input order as a fundamental attribute in the network's coupling to its environment, and indicate that networks learn best by entraining their dynamics to those of external stimuli near the EOC.

I-105. Emerging functional connectivity in a mesoscale brain network during sensorimotor task learning

Yaroslav Sych[1] SYCH@HIFO.UZH.CH
Leonardo Novelli[2] LNOV6504@UNI.SYDNEY.EDU.AU
Fritjof Helmchen[1] HELMCHEN@HIFO.UZH.CH

[1]University of Zurich
[2]University of Sydney

Adaptive behaviors require coordinated activation of neural populations distributed across multiple brain regions. To study learning-induced transformation in such mesoscale network, we applied high-density-fiber photometry to simultaneously record bulk calcium signals from 12 brain regions. Multi-fiber arrays were chronically implanted in the mouse brain and photometry measurements were repeatedly performed over 2-4 weeks while mice were trained in a whisker-based sensorimotor texture discrimination task. Our goal was to analyze changes in functional connectivity of the targeted mesoscale subnetwork based on multi-region calcium signals. During learning of the task salient calcium signal peaks shifted from action-reward period to the time of stimulus presentation in essentially all recorded sub-regions, including basal ganglia (CPu, GP), thalamus (VM, LD and RT), S1 barrel cortex, primary motor cortex (M1), and hippocampus (CA1). Signal correlations to whisking envelope and texture touch increased in multiple brain regions along with the improvement in task performance, while correlation to licking rate increased only in the primary somatosensory cortex. We estimated functional connectivity using the transfer entropy (TE), a proxy for the directed information flow between the brain regions. The resulting functional networks were distinct for Hit vs. Correct Rejection trials and thus predictive for the trial outcome. The number of functional links in the network increased with a learning onset. Clustering coefficient for the basal ganglia was the highest as compared to other brain regions. Measures of network robustness indicated a stabilization of the network during learning, i.e. the number of shared links in expert sessions was higher as compared to sessions during naive or learning phases. Taken together, our results highlight the cooperative action of distributed brain regions during learning to establish the functional connectivity and mesoscale brain dynamics required to solve the task.

I-106. Detecting information-limiting correlations in early visual cortex

Jorrit Montijn[1] JSMONTIJN@GMAIL.COM
Rex Liu[1] REX.LIU@UNIGE.CH
Amir Aschner[2] AASCHNER@MAIL.EINSTEIN.YU.EDU
Adam Kohn[2] ADAM.KOHN@EINSTEIN.YU.EDU
Peter Latham[3] PEL@GATSBY.UCL.AC.UK
Alexandre Pouget[1] ALEX.POUGET@GMAIL.COM

[1]University of Geneva
[2]Albert Einstein College of Medicine
[3]Gatsby Computational Neuroscience Unit, University College London

If the brain processes incoming data efficiently, information in the relevant neural circuits should match behavioral performance. For instance, if there is enough information in a visual cortical area to determine the orientation of a grating to within 1 degree, and the code is simple enough to be read out by downstream circuits, then animals should be able to achieve that performance behaviorally. Despite over 30 years of research, it is still not known how efficient the brain is. For tasks involving a large number of neurons, the amount of encoded information is limited by differential correlations, so determining the amount of encoded information requires quantifying the strength of differential correlations. Detecting them, however, is difficult; current methods require recording spikes from thousands of neurons simultaneously, which is not yet feasible. We report here a new method, which requires on the order of 100s of neurons and trials. This method relies on computing the alignment of the neural stimulus encoding direction, f', with the principal components of the noise covariance matrix, Σ. In the presence of strong differential correlations, f' must be spanned by a small number of the largest principal components of Σ. Using simulations with a leaky-integrate-and-fire neuron model of the LGN-V1 circuit, we confirmed that this method can indeed detect even small levels of differential correlations, consistent with those that would limit orientation discrimination thresholds to 0.1 degrees. We applied this technique to V1 recordings in awake monkeys. This revealed differential correlations consistent with a discrimination threshold of 0.5 degrees, which is not far from typical discrimination thresholds (1-2 deg). These results suggest that V1 contains about as much information as is seen in behaviour, which would imply that downstream circuits are efficient at extracting the information available in V1.

I-107. Hierarchical inference interactions in dynamic environments

Tahra Eissa[1] TAHRA.EISSA@COLORADO.EDU
Nicholas W Barendregt[1] NICHOLAS.BARENDREGT@COLORADO.EDU
Joshua Gold[2] JIGOLD@PENNMEDICINE.UPENN.EDU
Krešimir Josić[3] JOSIC@MATH.UH.EDU
Zachary Kilpatrick[1] ZPKILPAT@COLORADO.EDU

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

In a constantly changing world, accurate decisions require flexible evidence accumulation. As old information becomes less relevant, it should be discounted at a rate adapted to the frequency of environmental changes.However, sometimes humans and other animals must simultaneously infer the state of the environment and its volatility (hazard rate). How do such inference processes interact when performed hierarchically? To address this question, we developed and analyzed a model of an ideal observer who must report either the state or the hazard rate. We find that the speed of both state and hazard rate inference is mostly determined by information integration across change points. Our observer infers the state and hazard rate by integrating noisy observations and discounting them according to an evolving hazard rate estimate. To analyze this model and its variants, we developed a new method for computing the observer's state and hazard rate beliefs. Instead of sampling, we solve a set of nonlinear

partial differential equations (PDEs), leading to faster and more accurate estimates. We characterize how optimal and suboptimal (those with mistuned evidence discounting rates or other discounting functions) observers infer the state and hazard rate and compare their performance in tasks with varying difficulty. Evidence near change points strongly perturbs the state belief and supports higher hazard rates. Thus, state and hazard rate inference are linked, and the speed of hazard rate learning is primarily determined by how well the observer accounts for change points. Early in a trial, changes may not be well tracked, as the observer's hazard rate estimate is poor, but this estimate improves as the trial evolves. We measure how biases in hazard rate learning influence an observer's state inference process. Our setup can therefore be used to improve dynamic decision task design by identifying parameterizations that reveal hierarchical inference strategies.

I-108. Arousal, behavioral and neural determinants of exploration behavior

Anjali Raja Beharelle[1] ANJALI.RAJABEHARELLE@ECON.UZH.CH
Marcus Grueschow[1] MARCUS.GRUESCHOW@ECON.UZH.CH
Rafael Polania[2,3] RAFAEL.POLANIA@HEST.ETHZ.CH
Marius Moisa[1] MARIUS.MOISA@ECON.UZH.CH
Todd Hare[1] TODD.HARE@ECON.UZH.CH
Christian Ruff[1] CHRISTIAN.RUFF@ECON.UZH.CH

[1]University of Zurich
[2]ETH Zurich
[3]Swiss Federal Institute of Technology

Adaptive gain theory (AGT) posits that switches between exploration and exploitation behavior are arbitrated by a locus-coeruleus – noradrenergic arousal system (LC-NE) that responds to changes in the expected utility of selecting an action. Here, we test key predictions of AGT using fMRI and fluctuations in pupil diameter, which have been shown to index LC-NE activity. Participants (N = 66) played a restless bandit task while undergoing simultaneous measurements of fMRI-BOLD activity and pupil diameter. Participants' exploratory decisions were classified using an established reinforcement learning model. Changes in the pupil baseline (PBL) diameter and neural activity were compared between explore vs. exploit choices, for both the current trial and the trials preceding the present choice. We found that PBL diameter was increased prior to explore relative to exploit choices (Figure 1a). This PBL increase was linked to a higher probability of exploration, even when controlling for relevant decision variables and was accompanied by elevated activity in anterior cingulate cortex (ACC) for explore vs. exploit choices. Participants with increased PBL and ACC activity just prior to explore relative to exploit choices showed more optimal decision-making (increased earnings on the task). Interestingly, PBL was already elevated up to two trials prior to exploration, even though participants mostly exploited on these trials (Figure 1b). These increases in PBL were accompanied by concomitant increases in frontopolar cortex activity (Figure 2). Our results provide novel evidence for the central prediction of AGT that exploratory behavior is triggered by context-adaptive responses of the arousal system, and they highlight the importance of LC-noradrenergic interactions with ACC in initiating exploration. Increased sensitivity of neural arousal states to exploration vs. exploitation may facilitate behavioral flexibility and improve performance.

I-109. Whole-brain calcium imaging during physiological vestibular stimulation in larval zebrafish

Geoffrey Migault[1] MIGAULT@HOTMAIL.FR
Bernhard Englitz[2] B.ENGLITZ@DONDERS.RU.NL
Thijs L van der Plas[2] VDPLASTHIJS@GMAIL.COM
Hugo Trentesaux[1] HUGO.TRENTESAUX@UPMC.FR
Thomas Thomas[1] THOMAS.PANIER@UPMC.FR
Natalia Belen Beiza Canelo[1] NBBEIZA@GMAIL.COM
Raphael Candelier[1] RAPHAEL.CANDELIER@UPMC.FR
Remi Proville[3] REMIPROVILLE@GMAIL.COM
Georges Debregeas[1] GEORGES.DEBREGEAS@UPMC.FR
Volker Bormuth[1] VOLKER.BORMUTH@UPMC.FR

[1]Sorbonne Universite
[2]Radboud University
[3]Neurocentre Magendie, Physiopathologie de la Plasticite Neuronale, INSERM

The vestibular apparatus provides animals with postural and movement-related information that are essential to adequately execute numerous sensorimotor tasks. In order to activate this sensory system in a physiological manner, one needs to macroscopically rotate or translate the animal's head, which in turn renders simultaneous neural recordings highly challenging. Here we report on a novel miniaturized, light-sheet microscope that can be dynamically co-rotated with a head-restrained zebrafish larva, enabling controlled vestibular stimulation. The mechanical rigidity of the microscope allows one to perform whole-brain functional imaging with state-of-the art resolution and signal-to-noise ratio while imposing up to $25°$ in angular position and $6000°/s^2$ in rotational acceleration. We illustrate the potential of this novel setup by producing the first whole-brain response maps to sinusoidal and stepwise vestibular stimulation. The responsive population spans across multiple brain areas, displays bilateral symmetry and its organization is highly stereotypic across individuals. Using Fourier and regression analysis, we identified three major functional clusters that exhibit well-defined phasic and tonic response patterns to vestibular stimulation. Our rotatable light-sheet microscope provides a unique tool to systematically study vestibular processing in the vertebrate brain, and extends the potential of virtual-reality systems to explore complex multisensory and motor integration during simulated 3D navigation.

I-110. Investigating the active generation of behavior variability

Marwen Belkaid BELKAID@ISIR.UPMC.FR
Elise Bousseyrol ELISE.BOUSSEYROL@GMAIL.COM
Romain Durand-de Cuttoli ROM.DURAND@GMAIL.COM
Malou Dongelmans MALOU.DONGELMANS@GMAIL.COM
Etienne K Durante ETIENNE.DURANTE@GMAIL.COM
Tarek Ahmed Yahia TAREKAHMEDYAHIA@GMAIL.COM
Jeremie Naude JEREMIE.NAUDE@GMAIL.COM
Olivier Sigaud OLIVIER.SIGAUD@UPMC.FR
Philippe Faure PHFAURE@GMAIL.COM

Sorbonne Universite

Numerous studies have delved into the neurobiological mechanisms by which individuals use past experience to repeat successful actions. Typically, reinforcement learning is conceived as a process that increases the propensity to perform actions that have proved successful in the past [Sutton and Barto, 1998; Schultz, 2002], thus leading to a repetitive behavior. Nevertheless, even when repeatedly facing the same situation, behavioral variability can constitute a strategic advantage. Ecological examples of such contexts range from preys escaping predators to humans playing competitive games. Strikingly, an abundant literature suggests that voluntary gen-

eration of random behavior is an arduous process that involves (costly) executive functions and in which human performance is rather weak (or suboptimal) [Wagenaar, 1972; Oomens et al., 2015]. Therefore, it is unclear whether the brain is able to actively generate variability; and if so, how.

To investigate this question, we designed a mice experiment reinforcing non-repetitive choice sequences. Mice were trained to perform a sequence of binary choices in an open-field where three locations were explicitly associated with intra-cranial self-stimulation rewards. Animals were rewarded when their decisions increased the grammatical complexity of their last sequence of 9 choices. Unlike other approaches [Tervo et al. 2014], the task is fully deterministic, thus allowing a comprehensive characterization of the animal's choices. Overall, mice could use either random selection or a more cognitive, memory-based process.

In line with related works [Grunow and Neuringer, 2002; Tervo et al. 2014], we found that mice progressively increased the variability of their choice sequences until they reached the highest possible performance using random selection. More importantly, our experimental and computational results suggest that they do so using a memory-free, pseudo-random selection. They also point toward an active process allowing the tuning of the decision-making parameters in order to adapt to the need of variability.

I-111. Reliability and precision of single spike propagation in the reptilian cortical circuit

Juan Luis Riquelme JUAN.RIQUELME@BRAIN.MPG.DE
Mike Hemberger MIKE.HEMBERGER@BRAIN.MPG.DE
Gilles Laurent GILLES.LAURENT@BRAIN.MPG.DE
Julijana Gjorgjieva GJORGJIEVA@BRAIN.MPG.DE

Max Planck Institute for Brain Research

How do the properties of a network of neurons affect a signal traveling through it? Theoretical studies have classically addressed this question from the perspective of firing rates in populations of neurons. However, recent experimental studies have shown that single cells can impact global network activity and behavior, while spiking patterns in behaving animals may display fine temporal structure. Here we explore the effects of single spikes on the properties of cortical network activity using numerical simulations. We address this question in the three-layered cortex of turtles, which provides a simple architecture and enables the simultaneous recording of hundreds of cells. Recent experiments reveal that inducing single spikes in pyramidal cells evokes spatio-temporal patterns of activity in the surrounding population that unfold over tens to hundreds of milliseconds. A subset of neurons preserves the temporal precision of their first spike over multiple trials, resulting in a repeatable order of activation that spreads radially from the source cell. We developed a network model of spiking neurons, constrained by experimental measurements, to study the properties of the network response to induced excitatory spikes. We showed that single-spike transmission can be mediated through a heavy-tailed distribution of synaptic strengths in densely connected networks. Spatially-dependent connectivity supported signal propagation between groups of strongly coupled excitatory neurons that amplified network activity. Although initial groups of cells reliably activate, the later path taken by the propagating activity varied over trials, reducing reliability of the entire excitatory population. In contrast, the response of the inhibitory population was less temporally precise over trials but more reliable. Our work demonstrates that the response of a cortical circuit to single spikes can be strong and lasting, while displaying characteristic levels of temporal precision and reliability that are linked to the connectivity properties of the network.

I-112. Prediction of future input explains lateral connectivity in primary visual cortex

Emil Fristed EMIL.FRISTED@DPAG.OX.AC.UK
Yosef Singer YOSEF.SINGER@STCATZ.OX.AC.UK
Andrew King ANDREW.KING@DPAG.OX.AC.UK
Maria Iacaruso FLORENCIA.IACARUSO@GMAIL.COM
Nicol Harper NICOL.HARPER@ENG.OX.AC.UK

University of Oxford

Neurons in primary visual cortex (V1) receive abundant synaptic input via recurrent connections of other cortical neurons. The local synaptic connectivity in layer 2/3 of V1 is functionally specific, whereby neurons with similar visually-driven activity patterns, or similar orientation tuning, connect preferentially. This specificity likely plays an important role in information processing, but it is unknown what determines these connectivity patterns. We hypothesized that the connectivity of neurons in V1 is optimized to predict the immediate future of visual inputs from recent past inputs. To test this hypothesis, we implemented a recurrent form of a temporal prediction model, in which a one-hidden-layer recurrent neural network (RNN) was trained to predict the next frame in movies of natural visual scenes from the previous frames. Tuning properties of the model's hidden units were obtained by measuring their responses to drifting grating stimuli and natural visual scenes, and network connectivity was obtained from the recurrent weight matrix. The network developed oriented receptive fields (RFs) and tuning properties resembling those of pyramidal cells in layer 2/3 in mouse V1, including orientation- and direction-selectivity, and both simple- and complex-like cells emerged in the network. Crucially, the network developed functionally-specific connectivity resembling local connectivity between pyramidal cells of layer 2/3 in V1 of mice: Model units with similar responses to natural movies or with similar orientation tuning connected preferentially. Furthermore, reducing the connectivity in the network changed the responses of complex-cell-like units to simple-like. Our findings provide a normative model for lateral connections in V1, and supports the hypothesis that primary visual cortex is optimized to predict the immediate future from the recent past.

I-113. Conjunctive map for odors and space in posterior piriform cortex

Cindy Poo[1] CINDY.POO@NEURO.FCHAMPALIMAUD.ORG
Gautam Agarwal[2] GAUTAM.AGARWAL@NEURO.FCHAMPALIMAUD.ORG
Niccolo Bonacchi[2] NBONACCHI@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[3] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Research
[2]Champalimaud Neuroscience Programme
[3]Champalimaud Centre for the Unknown

Odors and space are two of the most important types of information used by rodents to guide behavior. It is hypothesized that spatial and olfactory brain systems co-evolved to guide a wide range of natural behaviors, and olfaction and spatial memory are linked to overlapping brain areas. Despite this, little is known about how olfactory and spatial information are combined in the brain. Posterior piriform cortex (pPir) receives input from olfactory bulb and anterior piriform cortex, as well as hippocampus and entorhinal cortex, making it a strong candidate to combine these information. However, there are few reported recordings from pPir. To investigate, we trained rats on a novel four-alternative odor-guided spatial choice task. On a plus maze with four water/odor ports, rats learned an odor-guided allocentric navigation rule (e.g. go to the north for odor A). We recorded from a large population of pPir neurons for novel and trained odors. We found that pPir neurons were selective to trained odors and the spatial location in which odors were sampled. Individual pPir neurons fired both independently and jointly to odors and locations. Odors and spatial locations can be decoded from pPir neurons in a single session. In a subset of interleaved recordings, rats were presented with novel vs trained odors. We found that pPir neurons were more selective for novel odors, while location selectivity remained constant. In addition, a smaller fraction

(32%) of neurons conjunctively encoded novel odors and space, while for odors which rats have learned spatial navigational rules the majority (56%) of odor selective neurons conjunctively encoded for space. This indicates that individual pPir neurons flexibly associates odors to a spatial map based on behavioral relevance. Together, we provide the first evidence for pPir as a site for integration across olfactory and spatial memory systems.

I-114. Representations and causal contributions of frontal regions during flexible decision-making

Marino Pagan MPAGAN@PRINCETON.EDU
Vincent Tang VDTANG@PRINCETON.EDU
Carlos Brody BRODY@PRINCETON.EDU

Princeton University

Our ability to flexibly select, based on context, the relevant information to guide our decisions is a fundamental cognitive process, yet its neural underpinnings are still largely unknown. To address this issue, we trained rats to perform a task requiring context-dependent selection and integration of sensory information (adapted from Mante et al., Nature, 2013). In our task, rats are presented with a train of randomly-timed auditory pulses, where each pulse varies in its location (right or left) and its tone pitch (high or low). In separate blocks of trials, rats are cued to report either the prevalent location of the pulses, or their prevalent pitch. Using an automated training procedure, we trained 40 rats to perform this task with high accuracy. To greatly increase the speed of data collection, we developed an automated, high-throughput pipeline to wirelessly record and perturb the activity of large ensembles of single neurons during behavior. Using this system we recorded neural responses in the Frontal Orienting Fields (FOF; 5 rats) and in medial Prefrontal Cortex (mPFC; 2 rats). Population decoding analyses revealed that context information (i.e. Location vs Frequency) peaks earlier in mPFC, whereas choice information (i.e. Right vs Left) is first represented in FOF. Furthermore, the encoding of choice in FOF followed the rats' decisions even on incorrect trials, consistent with a role of FOF in choice selection. Preliminary optogenetics results were also consistent with this hypothesis, as unilateral inactivation of FOF produced an ipsilateral choice bias. Our results suggest complementary roles for FOF and mPFC during context-dependent evidence accumulation. They also demonstrate that automated behavior, electrophysiology and optogenetic experiments in rats can vastly increase the amount of collected data compared to traditional techniques and provide a powerful tool to dissect the neural mechanisms underlying flexible behavior.

II-1. Studying the neuronal substrates of internal models via a novel closed-loop olfactory task for mice

Priyanka Gupta PGUPTA@CSHL.EDU
Marie Dussauze MARIE.DUSSAUZE@GMAIL.COM
Uri Livneh ULIVNEH@CSHL.EDU
Dinu Albeanu ALBEANU@CSHL.EDU

Cold Spring Harbor Laboratory

During behavior, sensation and action operate in closed-loop. Movements shape sensory input, and sensory inputs guide motor commands: where one looks determines what one sees. Through experience, the brain learns the reciprocal relationship between sensory inputs and movements to build internal models that predict the sensory consequences of upcoming actions (sensorimotor predictions). Comparing internal sensory predictions to actual sensory observations generates prediction errors that can be minimized by learning increasingly accurate models of the world. This exchange of sensory inputs and egocentric expectations is at the core of active perception. Experimental investigation of this idea has been sparse and split between behavioral interrogation of sensory-guided, precise motor control in primates (visuomotor adaption tasks) and the search for neuronal

substrates of sensory predictions in rodents via simpler running-based closed-loop behaviors. To study internal models both at behavioral and circuit-level, we developed a novel behavioral task where head-fixed mice are trained to steer the left-right location of an odor source by controlling a light-weight lever with their forepaws. In this manner, 1) we link a precise motor action to well-defined sensory expectations (odor location) and 2) subsequently violate the learnt expectations via online feedback perturbations in trained animals. Expert mice (6 out of 6 trained mice, training period $<$ 2 weeks, $>$90% accuracy, 400-800 trials/session) showed precise movements that were locked to the instantaneous odor feedback during normal closed-loop coupling. However, when sensory feedback was transiently interrupted (halting of odor source) or distorted (displacement of odor source or change in movement gain), movements were initially guided by each animal's learnt internal model and further, quickly adapted (within few sniffs in single-trials) in accordance with the instantaneous sensory error. We are currently probing activity in olfactory and motor cortex and the olfactory striatum to understand the sensorimotor transformations that enable this behavior.

II-2. Neuronal mechanisms underlying multisensory stimulus detection in the primary and secondary visual cortex of the mouse

Guido Meijer[1,2] GUIDO.MEIJER@RESEARCH.FCHAMPALIMAUD.ORG
Cyriel MA Pennartz[3] C.M.A.PENNARTZ@UVA.NL
Pietro Marchesi[3] P.MARCHESI@UVA.NL
Jorge Mejias[3] J.F.MEJIAS@UVA.NL
Jorrit Montijn[3] JORRITMONTIJN@GMAIL.COM
Carien Lansink[3] C.S.LANSINK@UVA.NL

[1]Champalimaud Centre for the Unknown
[2]International Brain Laboratory
[3]University of Amsterdam

How multisensory integration of sensory information supports behavioral actions and where in the brain this process takes place are still open questions. Area anterio-lateral (AL) in the secondary visual cortex is a prime candidate to serve as a hub for sensorimotor transformation of audiovisual information, because it receives input from primary visual (V1) and auditory (A1) cortex and projects strongly to motor areas. Using two-photon calcium imaging we recorded the activity of large groups of single units in V1 and AL in mice performing on a detection task using visual, auditory and audiovisual stimuli. We found AL neurons to respond more strongly than V1 neurons to stimuli with a contrast near the detection threshold (1% contrast) but not to full contrast stimuli. A computational model of V1-AL connectivity showed that the elevated response of AL neurons to low contrast visual stimuli can be explained by convergent input from V1 to AL neurons. Neuronal activity in AL, but not V1, was closely aligned with the psychophysical performance of the mouse because decoding of stimulus presence based on the population activity in AL resulted in a psychometric function which was statistically indistinguishable from the behavioral psychometric function. During multisensory detection behavior, mice showed enhanced detection performance for audiovisual compared to visual stimuli. In both V1 and AL we found a fraction neurons (25%) which showed a significant modulation of their response during visual versus audiovisual stimulation regardless of the behavioral choice of the animal. AL neurons, however, showed stronger differentiation of behaviorally reported versus unreported stimuli compared to V1 during audiovisual trials but not visual-only trials. These results suggest that area AL more closely correlates with multisensory detection behavior than V1 indicating that this area is a stage of the sensorimotor transformation pathway of multisensory information.

II-3. Intrinsic timescales define a cortical hierarchy and suggests network-tuning to task requirements

Jonas Dehning[1] JONAS.DEHNING@DS.MPG.DE
Viola Priesemann[1] VIOLA.PRIESEMANN@DS.MPG.DE
Nicholas Dotson[2] NICHOLAS.DOTSON@LIVE.COM
Steven Hoffman[3] STEVEHOFFMAN89@HOTMAIL.COM
Charles Gray[3] CHARLES.GRAY@MONTANA.EDU

[1]Max Planck Institute for Dynamics and Self-Organization
[2]University of California, Berkeley
[3]Montana State University

In cortex, information processing is likely to be organized hierarchically. Recently, a functional hierarchy was identified across five visual areas and two somatosensory areas in different macaques, based on estimates of an intrinsic timescales for each area (Murray et al., Nat. Commun., 2014). This intrinsic timescale τ implies that information can be maintained in ongoing activity without the need of synaptic changes, and can thereby be used for processing across an extended time window of the order of the intrinsic timescale. We here present results that encompass 39 areas recorded simultaneously from the same monkey and 16 areas from a second one. We find in both monkeys the same hierarchical organization: Visual cortex exhibits the shortest timescales, prefrontal and posterial parietal intermediate ones and motor and somatosensory cortex the highest. More importantly, we found a dependence of the intrinsic timescales on the type of task performed: timescales are higher during visual memory tasks than during fixation periods. Furthermore, neurons that are presumably inside the visual receptive field of the relevant stimuli tend to increase their timescale when confronted with a task, whereas the ones that are outside decrease theirs. Such a decrease (increase) in the timescale is related to a direct decrease (increase) in computational properties like the sensitivity, dynamic range and amplification strength of the network. Therefore, we interpret the increase of timescale of the neurons in the visual field as a tuning-in to task requirements (higher sensitivity) whereas the others tune-out, decreasing their sensitivity. Thereby we did not only expand on earlier insights about a hierarchy of information processing, but also showed that changes of the intrinsic timescales may be explained by changing task requirements. This advances our understanding of the relation between temporal dynamics and the function of cortical networks.

II-4. Neocortical inhibitory engrams protect against memory interference

Renee S Koolschijn[1] RENEE.KOOLSCHIJN@KEBLE.OX.AC.UK
William F Podlaski[1] WILLIAM.PODLASKI@CNCB.OX.AC.UK
Uzay E Emir[2] UEMIR@PURDUE.EDU
Tim P Vogels[1] TIM.VOGELS@CNCB.OX.AC.UK
Tim EJ Behrens[1] BEHRENS@FMRIB.OX.AC.UK
Helen C Barron[1] HELEN.C.BARRON@GMAIL.COM

[1]University of Oxford
[2]Purdue University

Memories often overlap with each other, sharing features, stimuli or higher-order information. Despite this overlap, we are able to selectively recall individual memories to guide our decisions and future actions. Previous theoretical work and empirical evidence suggests that inhibition plays a role in quenching memory expression, to retain excitatory inhibitory (EI) balance and prevent run-away excitation (Vogels et al., 2011; Barron et al., 2016). Here, we hypothesise that neocortical inhibition also prevents inference between overlapping context-dependent memories. To investigate this possibility, we built a model network involving two partially overlapping memory structures. Using this model, we predicted that the neural signatures of the overlapping elements in the memories would co-activate if the inhibitory tone of the network is transiently reduced, causing memory interference. We then tested these predictions in humans using ultra-high field 7T MRI. We designed a memory task involving the model's overlapping memory structures, and developed an index for neural memory interference in the human neocortex, using a representational fMRI paradigm to measure the relative overlap between context-dependent associated stimuli. We then reduced the concentration of neocortical GABA using trans-cranial direct current stimulation (tDCS) to transiently disturb EI balance, and quantified the change in GABA using MR spectroscopy. We observed an increase in neocortical memory interference in proportion to the reduction in GABA, which in turn predicted behavioural memory performance. These findings support a role for neocortical inhibition in mediating memory interference, by preventing unwanted co-activation between overlapping memories.

II-5. Execution of logical negation in pre-verbal infants

Milad Ekramnia[1,2] MLD.EKRAMNIA@GMAIL.COM

[1]NeuroSpin, CEA Paris-Saclay
[2]unicog

Assessing the minimal set of logical primitives in preverbal humans can shed light on the framework under which a putative language of thought may function. To this end we proposed if pre-verbal infants possess an abstract negation operator a brief exposure to any arbitrary label that functions like a negation word should suffice to trigger this operator in novel instances well before they acquire the negation word in natural language. To address this hypothesis we home-trained the infants to four toys paired with their corresponding pseudo-words for a period of one week. Further during the experiment, in an EEG design, we trained subjects that when the labels are accompanied by a pseudo-word 'Kou' any of the toys can subsequently be shown on the screen except the toy that the label belongs to. 40 infants at the age of 5 months were introduced with exemplars of application of 'kou' as a negation word on three of the trained pairs in a total of 10 trials and then were tested on the 4th pair as well as on a novel object to assess the extents of generalization. The results of cluster and permutation analyses suggest that infants can successfully discriminate between the incongruent and congruent application of the negation word and subsequently apply this word to a novel object. We observed significant effects both early on at p400 from image onset as well as a late effect around 250 milliseconds from image offset. The results here suggest for the first time that pre-verbal infants can either learn the negation concept as a rule and link it to an arbitrary word in a very few trials or trigger a primitive negation operator based on the role exhibited by this word.

II-6. Dorsocentral striatum integrates sensory and motor signals during visually guided behavior

Andrew Peters PETERS.ANDREW.J@GMAIL.COM
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini MATTEO@CORTEXLAB.NET
Nicholas Steinmetz NICK.STEINMETZ@GMAIL.COM

University College London

The striatum has been implicated in learning and executing stimulus-guided movements, and is a large and heterogeneous structure with diverse anatomically organized inputs. To determine the role of the striatum in goal-directed actions, we sought to identify functionally distinct domains within the striatum and associate them with the cortical regions that serve as their major input. We achieved this in mice by inserting a Neuropixels probe in the striatum while imaging widefield calcium signals on the dorsal surface of the cortex. By regressing cortical fluorescence to striatal activity, we revealed a stereotypical functional topography between the cortex and striatum. This relationship was highly similar to anatomical projections, indicating that corticostriatal domains can be identified through population activity and that corticostriatal anatomy largely predicts activity correlations. We then investigated the activity within cortically-defined striatal domains during a visually guided task. Mice rested their forepaws on a steering wheel and were trained to turn the wheel left or right when a stimulus appeared on the right or left respectively. As the mice engaged in each trial, cortical fluorescence progressed from visual to frontal to retrosplenial and somatomotor areas. A similar sensory-to-motor gradient was present from dorsomedial to dorsolateral striatum. Along this gradient, a domain in dorsocentral striatum stood out as the only domain to show strong movement direction-specific signals. This domain was correlated with activity in parietal and frontal cortex. This domain also combined responses to both the stimulus and movement. Passive conditions revealed larger stimulus responses in the dorsocentral striatum for task-relevant stimuli compared to other stimuli despite similar visual cortical responses, suggesting an influence of learning in this domain of the striatum. We conclude that the striatum can be divided into functional domains, one of which may act as a site of sensorimotor integration during the task.

II-7. Dimensional reduction in networks of non-Markovian spiking neurons

Andrea Galluzzi[1] ANDREA.GALLUZZI@ISS.IT
Matteo Biggio[2] MATTEO.BIGGIO@UNIGE.IT
Marco Storace[2] MARCO.STORACE@UNIGE.IT
Maurizio Mattia[1] MAURIZIO.MATTIA@ISS.IT

[1]Istituto Superiore di Sanita
[2]University of Genoa

Message passing between components of a distributed physical system is non-instantaneous and contributes to determine the time scales of the emerging collective dynamics. In biological neuron networks this is due in part to local synaptic filtering of exchanged spikes, and in part to the distribution of the axonal transmission delays. How differently these two kinds of communication protocols affect the network dynamics is still an open issue due to the difficulties in dealing with the non-Markovian nature of synaptic transmission (Brunel et al., Phys. Rev. Lett. 2001; Moreno-Bote & Parga, Phys. Rev. Lett. 2004; Schuecker et al., Phys. Rev. E 2015; Schwalger et al., J. Comput. Neurosci. 2015). Here, we develop a mean-field dimensional reduction yielding to an effective Markovian dynamics of the population density of the neuronal membrane potential, valid under the hypothesis of small fluctuations of the synaptic current. The resulting theory allows us to prove the formal equivalence between the two transmission mechanisms, holding for any synaptic time scale, integrate-and-fire neuron model, spike emission regimes and for different network states even when the neuron number is finite. The theoretical approach we developed can be straightforwardly extended to the general case in which both axonal delays and non-instantaneous synaptic transmission coexist in the model network leading to a generalization of the previous dynamic mean-field equation. Overall, these networks have firing rate dynamics equivalent to those obtained with instantaneous synaptic transmission and a suited unimodal distribution of transmission delays, showing that our perturbative approach is capable to describe also Markovian neurons with dimension larger than two

II-8. Flexible cortical dynamics in adaptive control of sensorimotor behavior

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

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

Humans can efficiently adapt to changes in their environment. The neural basis of such adaptability, however, is not understood. A normative description of changes in the nervous system that are necessary to accommodate specific behavioral adjustments is especially lacking. In a highly interconnected system such as the brain, this problem becomes rapidly intractable unless we can understand how individual computational modules are altered to support adaptation. In the framework of dynamical systems, adaptive control can be accomplished in two ways associated with distinct timescales and neural signatures. First, controlling external inputs enables the system to visit unexplored regions of the state space. Alternatively, changes in synaptic coupling within the system modifies its latent dynamics, creating new activity patterns in the same region of state space. Because synaptic modifications occur on relatively long timescales, we hypothesize that input-control strategies support fast adaptation, while adjustments of latent dynamics serve slower adaptation. To test these hypotheses, we build on a sensorimotor task in which monkeys reproduce measured time intervals. In this task, responses are biased toward the mean of previously encountered intervals indicating that animals learn the underlying interval distribution. We leveraged animals' reliance on the distribution to elicit adaptation at two different timescales: fast adaptation (within a behavioral session) by covertly changing the mean of the distribution, and slow adaptation (across multiple sessions) by changing the variance. Recordings of population activity in frontal cortex during fast adaptation revealed a concomitant change of cortical responses. Consistent with our prediction, this adjustment appeared to be mediated by an input displacing neural trajectories to a different part of state space. These results provide evidence that fast adaptation does not involve changes of synaptic coupling within cortical areas; instead, inputs from elsewhere in the brain are adjusted giving access to suitable regions of cortical latent dynamics.

II-9. Thirst regulates motivated behavior through modulation of brain-wide neural population dynamics

William Allen[1] WALLEN@STANFORD.EDU
Michael Chen[2] MICHAELZHUCHEN@ALUMNI.STANFORD.EDU
Nandini Pichamoorthy[1] NANDINI.PICHAMOORTHY@GMAIL.COM
Rebecca Tien[1] RHYTIEN@STANFORD.EDU
Marius Pachitariu[3] MARIUS10P@GMAIL.COM
Liqun Luo[1] LLUO@STANFORD.EDU
Karl Deisseroth[1] DEISSERO@STANFORD.EDU

[1]Stanford University
[2]Oxford University
[3]HHMI Janelia Research Campus

Physiological needs produce motivational drives, such as thirst and hunger, that regulate behaviors essential to survival. Specific populations of hypothalamic neurons continually receive information on the status of these needs and must coordinate neuronal activity in many diverse circuits underlying sensation and action to produce the appropriate motivated behavior. How this vital coordination is achieved remains unclear. Here, we record dynamics from 24,000 neurons in 34 brain areas during thirst-motivated choice behavior, as mice gradually consumed water and became satiated. We found that water-predicting sensory cues produced a remarkably pervasive wave of activity that rapidly spread throughout the brain of thirsty animals, modulating the activity of over half of all recorded neurons. Surprisingly, these activity dynamics were gated by a specific, persistent, and brain-wide mode of population activity that encoded motivational state. Optogenetic activation of a small population of hypothalamic thirst-sensing neurons after satiation restored activity along this mode to its pre-satiation global state, as well as brain-wide task-related dynamics. These results reveal a simple mechanism by which motivation globally coordinates brain activity: motivational states specify initial conditions that determine how a brain-wide dynamical system transforms sensory input into behavioral output.

II-10. Adaptive erasure of spurious sequences in cortical circuits

Alberto Bernacchia[1] AB2349@CAM.AC.UK
Jozsef Fiser[2] JOOZSEFFISER@GMAIL.COM
Guillaume Hennequin[1] G.HENNEQUIN@ENG.CAM.AC.UK
Mate Lengyel[1] M.LENGYEL@ENG.CAM.AC.UK

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

The sequential activation of neurons, reflecting a previously experienced temporal sequence of stimuli, is believed to be a hallmark of learning across cortical areas, including the primary visual cortex (V1). While circuit mechanisms of sequence learning have been studied extensively, the converse problem, that is equally important for robust performance, has so far received much less attention: how to avoid producing spurious sequential activity that does not reflect actual sequences in the input? Here, we developed a new measure of sequentiality for multivariate time series, and a theory that allowed us to predict the sequentiality of the output of a recurrent neural circuit. Our theory suggested that avoiding spurious sequential activity is non-trivial for neural circuits: e.g. even with a completely non-sequential input and perfectly symmetric synaptic weights, the output of a neural circuit will still be sequential in general. We then show that the most celebrated principles of synaptic organization, those of Hebb and Dale, jointly act to effectively prevent spurious sequences. We tested the prediction that cortical circuits actively diminish sequential activity, in an experience-dependent way, in multielectrode recordings from awake ferret V1. We found that activity in response to natural stimuli, to which animals were continually adapted, was largely non-sequential. In contrast, when animals were shown entirely non-sequential artificial stimuli, to which they had not been adapted yet, neural activity was sequential at first, and then gradually became non-sequential within a few minutes of extended exposure. Furthermore, this difference between responses to natural and artificial stimuli was not present at eye opening but developed over several days. Our work identifies fundamental requirements for the reliable learning of temporal information, and reveals new functional roles for Dale's principle and Hebbian experience-dependent plasticity in neural circuit self-organization.

II-11. Navigating in neural and perceptual manifolds with closed-loop multi-site electrical microstimulation system

Sina Tafazoli[1] TAFAZOLI@PRINCETON.EDU
Camden MacDowell[1,2] CAMDENM@PRINCETON.EDU
Kate Letai[3] KCLETAI@GMAIL.COM
Daniel Che[1] ZCHE@PRINCETON.EDU
Tim Buschman[1] TBUSCHMA@PRINCETON.EDU

[1]Princeton University
[2]Rutgers Medical School
[3]Peking University

Recent work has suggested neural activity is restricted to a low-dimensional manifold: a few independent latent variables capture a significant fraction of neural variability. However, it remains unclear how the low-dimensional manifold relates to behavior. To address this, one must develop a causal method for navigating the manifold. To this end, we have developed a multi-dimensional stimulation system that allows neural activity to be perturbed in several independent directions. Using many-contact silicon probes, we electrically stimulated and recorded neural activity across 64 channels simultaneously. We found that response manifolds generated by visual and electrical stimulation were largely overlapping, suggesting large-scale electrical microstimulation can be used to navigate perceptual manifolds. To directly test this, we developed a novel closed-loop stimulation system which learns to generate specific patterns of neural activity. In this way, our system is able to navigate through the manifold of neural activity.

II-12. Localized random projections challenge benchmarks for bio-plausible deep learning

Bernd Illing[1,2] BERND.ILLING@EPFL.CH
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH
Johanni Brea[1] JOHANNI.BREA@EPFL.CH

[1]Ecole Polytechnique Federale de Lausanne
[2]Brain Mind Institute

Similar to models of brain-like computation, artificial deep neural networks rely on distributed coding, parallel processing and plastic synaptic weights. Training deep neural networks with the error-backpropagation algorithm, however, is considered bio-implausible. Numerous recent publications suggest various models for bio-plausible deep learning, typically defining success as reaching around 98% test accuracy on the MNIST data set. Here, we investigate a simplistic benchmark model: a network with one hidden layer and a single readout layer. The hidden layer weights are either fixed and random or trained with an unsupervised, local learning rule implementing either Principal Component Analysis or Sparse Coding. The readout layer is trained with a supervised, local learning rule. We first implement these models with rate neurons and systematically compare the MNIST classification performance between them and to networks trained with the backpropagation algorithm. This comparison revealed, first, that unsupervised learning does not lead to better performance than fixed random projections for large hidden layers on digit classification (MNIST) and object recognition (CIFAR10). Second, networks with random projections and localized receptive fields perform significantly better than networks with all-to-all connectivity and almost reach the performance of networks trained with the backpropagation algorithm on MNIST. Based on these results we implement our model with leaky-integrate-and-fire neurons with fixed random, localized receptive fields in the hidden layer and spike timing dependent plasticity to train the readout layer. %after a transient phase of the pattern presentation. This model achieves 98.1% test accuracy on MNIST, which is close to the optimal result achievable with error-backpropagation in non-convolutional rate networks with one hidden layer. The performance of these simple random projection networks is comparable to most current models of bio-plausible deep learning. Our results thus provide an important benchmark and instruct using datasets other than MNIST for testing deep learning capabilities of future models of bio-plausible deep learning.

II-13. Cortical basis of audiovisual spatial localization in mice

Philip Coen PIPCOEN@GMAIL.COM
Miles Wells MILES.WELLS@UCL.AC.UK
Matteo Carandini MATTEO@CORTEXLAB.NET
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK

University College London

The ability to combine visual and auditory cues to localize objects in space is critical to many organisms, whether prey, predator, or pedestrian crossing the street. Early inactivation experiments in cat suggested that audiovisual spatial localization relies on cortex, but the techniques available at the time lacked temporal precision and reversibility. These limitations could be readily overcome by using mice, where one can use optogenetics for comprehensive and reversible inactivation across dorsal cortex. However, recent work suggests that mice do not perform optimal audiovisual integration, displaying auditory dominance when auditory and visual cues are in conflict. Do mice combine audiovisual cues appropriately when performing spatial localization, and do they use cortex to do so? To answer these questions, we developed a two-alternative forced choice task where head-fixed mice turn a wheel to indicate whether a stimulus appeared on the left or right. The stimuli can be auditory, visual, or a combination of the two, presented in coherent or conflicting locations. In this task, mice showed no unisensory dominance, and integrated auditory and visual cues appropriately in all conditions. Indeed, a simple additive model, with independent weights for auditory and visual stimuli, completely predicted the response of mice to all cue combinations. To test the role of cortex we used optogenetics to randomly inactivate different spots across

dorsal cortex while mice performed this task. We identified distinct roles for different cortical regions. For example, V1 was required on visual, but not auditory, trials. Conversely, inactivating secondary motor cortex impaired performance independent of sensory modality. These results demonstrate that mice combine audiovisual cues without any unisensory dominance during spatial localization, suggesting that the underlying neural framework is homologous to other mammals. Further, we establish that different regions of cortex have distinct roles in this multimodal decision process.

II-14. Hidden neural states underlie history-dependent canary song sequences

Yarden Cohen[1] YARDENC@BU.EDU
Jun Shen[1] JUNSHEN@BU.EDU
Dawit Semu[1] DAWITS@BU.EDU
Daniel P Leman[1] DPLEMAN@BU.EDU
William A, III Liberti[2] BLIBERTI@BU.EDU
Nathan L Perkins[1] LNP@BU.EDU
Timothy J Gardner[1] TIMOTHYG@BU.EDU

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

In skills like speech or dance, motor sequences follow long-range syntactic rules – behavioral transitions depend on actions multiple time steps in the past. The neural substrates that support these high-order Markov chains, or 'complex' transitions, are unknown. Songbirds have a learned and naturally recurring behavior, song, whose temporal structure is largely governed by an avian 'cortico-thalamic loop' that includes the premotor nucleus HVC (Hahnloser 2002, Long 2008, Wang 2008, Nottebohm 1976). Canary songs are defined by syllables repeats, called phrases, that shape during song maturation from the syllables in the birds' repertoire (Gardner 2005) and exhibit long-range syntactic rules (Markowitz 2013). Specifically, certain phrase transitions have history dependence, that extends 2-3 phrases or dozens of syllables. To study the neural basis of long-range syntax rules in song, we use head mounted miniature microscopes to record [Ca2+] dynamics from canary HVC neurons. In repeating sequences, spanning up to four phrases, we discover hidden phrase-locked network states that reflect behaviorally-relevant information about song history. We find that higher order Markov dependence in neural activity correlates more often with the song's past than its future, and the cells that manifest more complex dependencies also occur during syllables that demonstrate long-range history dependence behaviorally. That behavior with long-range syntax rules is reflected by hidden network states suggest that such features should be included in simplified models of birdsong sequence generation. Investigating the underlying mechanisms may reveal neural processing common to a wider class of complex motor and cognitive functions such as inference and decision making.

II-15. Deep learning with segregated dendrites and multiplexing

Jordan Guerguiev[1] JORDAN.GUERGUIEV@UTORONTO.CA
Thomas Mesnard[2] THOMAS.MESNARD@GMAIL.COM
Richard Naud[3] RNAUD@UOTTAWA.CA
Blake Richards[4] BLAKE.RICHARDS@UTORONTO.CA

[1]University of Toronto Scarborough
[2]Mila
[3]University of Ottawa
[4]University of Toronto

In the machine learning field, deep learning has led to human-level performance in computer vision, speech

recognition, natural language processing and other complicated learning tasks. Interestingly, the most common algorithm that is used to train deep (many-layer) neural networks, backpropagation of error (backprop), leads to emergent feature representations that resemble those seen in the neocortex. However, backprop is not biologically realistic for several reasons, including the need for symmetric feedforward and feedback connections and separate networks states for feedforward and feedback processing. It remains unclear how the brain might implement backprop or some approximation to it in order to achieve powerful learning in the cortex. Here, we demonstrate a biologically plausible form of deep learning at the level of neuronal ensembles, which combines the unique morphology of cortical pyramidal neurons and the theory of burst ensemble multiplexing developed by Richard Naud and Henning Sprekeler. Our model uses computational units that represent ensembles of cortical pyramidal neurons. Each unit has two compartments, a somatic compartment receiving feedforward input and a dendritic compartment receiving feedback input. The feedforward weights in the network are trained in a way that approximates backprop using only information that is locally accessible to each ensemble in the network. Importantly, the network does not need to switch between feedforward and feedback states in order to learn. Using this model, we can match the performance of backprop on standard visual recognition tasks, as well as generate experimental predictions about circuit structure and how neuronal ensembles encode feedforward and feedback information. If evidence for this learning framework is found in the brain, it would constitute an important step in advancing our understanding of learning in the cortex.

II-16. Targeted comodulation supports accurate decoding in V1

Caroline Haimerl CH2880@NYU.EDU
Cristina Savin CS5360@NYU.EDU
Eero Simoncelli EERO.SIMONCELLI@NYU.EDU

New York University

Accurate decision making requires reliable encoding and flexible decoding of sensory information. Given a specific task, only a small subset of neurons encode relevant information and need to be identified by a decoder depending on that task. No current theory explains how such decoding flexibility can be achieved, while respecting biological constraints on available knowledge of the encoding population. Haimerl & Simoncelli (Cosyne, 2018) introduced a framework for neural decoding based on functionally-targeted stochastic modulators, motivated by experimental observations in monkey V4 (Rabinowitz et al, 2015). That work revealed that despite its detrimental effect on encoding, the introduction of shared modulatory noise in task-informative neurons could serve as a labeling mechanism for decoding. Specifically, it introduced a modulator-guided decoding scheme which does not require biologically unrealistic knowledge assumed by an ideal observer. It only needs access to the low-dimensional modulator. Here we test predictions of this theory using population recordings from macaque V1 during a discrimination task (Ruff &Cohen, 2016). We first fit a Poisson Linear Dynamical System (Macke et al, 2015) to the data and extract two latent sources of co-variability, a slow drift and a fast modulator. We find that the fast modulator preferentially targets task-informative neurons as predicted by the theory, whereas the slow drift is task irrelevant. Second, we compare the performance of decoders that differ in their assumed knowledge about the encoding population and hence in their biological plausibility. We find that the modulator-guided decoder reaches the upper bound set by an ideal observer. Moreover, its parameters can be learned from a small number of training trials. These results confirm that targeted modulation could provide a mechanism enabling flexible and accurate decoding of neural responses in V1.

II-17. From adult dentate gyrus neurogenesis to unsupervised clustering

Olivia Gozel OLIVIA.GOZEL@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

Ecole Polytechnique Federale de Lausanne

Adult neurogenesis of dentate granule cells (DGC) only concerns a small percentage of the dentate gyrus (DG) principal cells population, and yet it has been shown to strikingly promote behavioral pattern separation of similar stimuli in a variety of tasks. The properties of newborn DGC evolve as a function of maturation. Ultimately, only the ones that are well integrated into the DG survive, a process which has been observed to critically depend on GABAergic input. In the early maturation phase of newborn DGC, the GABAergic input they receive has an excitatory effect, and it becomes inhibitory in the late phase of their maturation. It is still unknown, however, why the switch from excitation to inhibition in adult DG neurogenesis is crucial for proper integration. The main input to DG comes from the entorhinal cortex (EC). We model the EC-to-DGC circuit as an input layer (EC) fully feedforwardly connected to a winner-take-all (WTA) network that represents DG. We consider stimuli organized in several clusters. Stimuli coming from a novel cluster are presented intermingled with stimuli from previously stored clusters. We suggest that adult DG neurogenesis provides a biological solution to the problem of dead units, which is a common issue in competitive unsupervised learning. Indeed, in the early phase of maturation, GABAergic input is excitatory. Hence, cooperativity makes the EC connections towards newborn DGC grow in the direction of the subspace of previously presented stimuli. When GABAergic input switches to inhibitory in the late phase of maturation, the DG network becomes competitive. Thus the feedforward weights towards newborn DGC are driven away from previously stored clusters and towards the novel cluster. To our knowledge, we present the first model that can explain both how adult newborn DGC integrate into the preexisting network and why they promote pattern separation of similar stimuli.

II-18. Stable attractor dynamics in prefrontal cortex during oculomotor decision making

Aniruddh Galgali[1,2] GALGALIA@INI.UZH.CH
Maneesh Sahani[3] MANEESH@GATSBY.UCL.AC.UK
Valerio Mante[1] VALERIO@INI.UZH.CH

[1]University of Zurich
[2]ETH Zurich
[3]Gatsby Computational Neuroscience Unit, University College London

The emerging idea that neural population responses in frontal and motor cortices are the manifestation of a latent dynamical system has been used to obtain critical insights about the neural underpinnings of evidence accumulation and movement generation. However, evidence for the specific type of latent dynamics being implemented (for e.g. line attractors/point attractors/rotational dynamics) has often been rather indirect, as it relied on generating low dimensional visualizations of the recorded condition-average trajectories and comparing them to ones from simulated recurrent neural networks. We propose a novel approach of quantifying the temporal dynamics of trial-by-trial variability in population responses as a way to better distinguish between different classes of latent neural dynamics. Our approach involves fitting a linear time-varying dynamical system to population spike-count residuals using concepts from subspace identification and instrumental variable regression. Applying this approach to simultaneous recordings from monkey pre-frontal cortex during a perceptual decision-making task revealed that latent dynamics during evidence accumulation and saccadic eye movements is consistent with a single fixed point under the influence of a time-varying, external input. This finding is at odds with several previously proposed models for evidence integration (e.g. a line attractor) and movement generation (e.g. rotational dynamics) that could not have been excluded based on the properties of condition-averaged trajectories alone. An analysis of trial-by-trial variability can thus provide insights into the latent dynamics underlying neural population responses, even in settings when the properties of external inputs are not precisely known.

II-19. Flexible categorization in perceptual decision making

Genis Prat Ortega[1] GENISPRAT@GMAIL.COM
Jaime de la Rocha[1] JROCHAV@CLINIC.CAT
Klaus Wimmer[2] KWIMMER@CRM.CAT
Alex Roxin[2] AROXIN@CRM.CAT

[1]IDIBAPS
[2]Centre de Recerca Matematica

Our brains interpret ambiguous streams of information to make decisions and guide our behavior. Canonical approaches to model this cognitive function are based on diffusion processes that assume bounded or unbounded perfect stimulus integration. However the relationship of such models with the underlying neural circuitry is unclear. Here we study the integration process in neurobiological models with winner-take-all dynamics. Such models can actually be reduced to a nonlinear diffusion process which in the case of binary categorizations, can be described by a double well potential (DW). To show the key mechanisms that differentiate the DW model from the canonical ones, we calculated Psychophysical Kernels (PK) using stimuli with zero net evidence in a fixed-duration simulated task. We found that the DW showed different integration regimes, from transient (primacy) to leaky (recency) integration, as the magnitude of the stimulus fluctuations (σs) or the stimulus duration (T) increased yielding an increased transition rate between the two wells. Between these two extreme regimes lied the flexible categorization regime for which fluctuations during the whole trial robustly generate decision reversals by overcoming the internal attractor dynamics only when the initial choice was incorrect. These correcting transition caused a non-monotonically relation of accuracy and σs. We tested the existence of the flexible categorization regime using previous data from human subjects (n= 22) performing a luminance categorization task. As predicted by the DW model, but not by the canonical models, the PKs changed from primacy to recency as stimulus duration increased. Our findings show that winner-take-all attractor models are more versatile than previously thought when operating in the flexible categorization regime, a regime that seems at play during perceptual categorization tasks.

II-20. A mechanistic model of cortical contributions towards mouse visual discrimination

Peter Zatka-Haas P.ZATKA@UCL.AC.UK
Nicholas A Steinmetz N.STEINMETZ@UCL.AC.UK
Matteo Carandini MATTEO@CORTEXLAB.NET
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK

University College London

The neocortex is thought to play a crucial role in decision behaviour. However it remains unclear what the precise contribution of cortex is, and how multiple areas coordinate to drive decision behaviour. Here, we establish quantitative laws for how cortex drives decisions in mice performing a visual task, and use these laws to show that two cortical regions, visual cortex and secondary motor cortex, play precise and distinct roles in the decision process. We trained mice in a two-alternative unforced visual discrimination task. Mice were rewarded with water for turning a wheel to indicate which of two Gabor stimuli had higher contrast, or for holding the wheel still if no stimulus was present. Widefield calcium imaging showed robust sequential activation of nearly the whole dorsal cortex, whereas scanning optogenetic inactivation revealed that only visual (VIS) and secondary motor (MOs) cortices play a causal role. To ascertain what type of computation VIS and MOs might be performing, we developed a mechanistic model of how VIS and MOs activity contribute to a decision. In the model, a linear weighted sum of each hemisphere's VIS and MOs activity is fed to a 3-alternative softmax decision function. We fit this model to widefield fluorescence data, and found that it could quantitatively account for average behavior and trial-by-trial choice variability. This model proposes distinct roles for VIS and MOs, suggesting that downstream decision circuits weight VIS cortical activity subtractively, but MOs activity additively. We further show that the

mechanistic model could predict the behavioural effect of optogenetic inactivation, despite the inactivation trials not participating in the model fit. We conclude that cortex's role in the decision task can be accurately modeled as providing a representation of sensory data to a subcortical decision circuit, that makes a stochastic choice based on linear weighted sum of this input.

II-21. Fluctuations in neural activity are reflected in the structure of associative memory networks

Chi Zhang C.ZHANG@NEU.EDU
Danke Zhang D.ZHANG@NORTHEASTERN.EDU
Armen Stepanyants A.STEPANYANTS@NORTHEASTERN.EDU

Northeastern University

Neural networks in the brain can function reliably despite various sources of noise present at every step of signal transmission. These sources include errors in presynaptic firing activity, noise in synaptic transmission, and fluctuations in postsynaptic potentials. Collectively they lead to errors in neurons' outputs which are, in turn, injected into the network. Here, we examine the effects of different sources of noise on structure of associative McCulloch and Pitts networks of inhibitory and excitatory neurons. Each neuron in the network is required to associate noisy input patterns with desired outputs. In contrast to previous studies, we assume that associations of individual neurons may contain errors, so long as they can be corrected at subsequent steps of activity propagation through the network. In addition, learning in the model is accompanied with sign-constraints on the weights of postsynaptic connections of neurons, and homeostatic constraints on the weights of the neurons' presynaptic inputs. We use the replica theory to determine the critical memory storage capacity of the network and examine its structural properties as a function of the presynaptic errors-rate and postsynaptic noise strength. Our results are consistent with the available experimental data on probabilities of inhibitory and excitatory connections in local cortical networks. Noisy connections in our model are depressed or completely eliminated during learning, and the resulting distributions of connection weights can have long, super-exponential tails, which is consistent with published electrophysiological data.

II-22. Learning divisive normalization from natural images to predict responses in primary visual cortex

Max Gunthner[1] MAX.GUENTHNER@BETHGELAB.ORG
Santiago Cadena[1] SANTIAGO.CADENA@BETHGELAB.ORG
George Denfield[2] GDENFIELD@CNS.BCM.EDU
Edgar Walker[3] EWALKER@CNS.BCM.EDU
Leon Gatys[1] LEON.GATYS@BETHGELAB.ORG
Andreas Tolias[2] ASTOLIAS@BCM.EDU
Matthias Bethge[1] MATTHIAS.BETHGE@UNI-TUEBINGEN.DE
Alexander Ecker[1] ALEXANDER.ECKER@UNI-TUEBINGEN.DE

[1]University of Tubingen
[2]Baylor College of Medicine
[3]Baylor College of Medicine, USA

Deep neural networks (DNNs) are state-of-the-art for predicting single-unit monkey V1 responses to natural stimuli and outperform classical linear-nonlinear and wavelet-based feature representations, but we currently do not know what kind of nonlinear computations DNNs approximate, limiting our understanding of V1 function. A good candidate for such nonlinearities is divisive normalization (DN), which explains a wide range of neurophysiological phenomena observed with simple stimuli. However, it is currently unknown whether DN plays an important role

in processing natural stimuli in the brain and, if so, which surrounding neurons contribute to any given neuron's normalization pool and with what weight. Here, we propose an end-to-end trainable DN module consisting of a collection of normalization pools factorized as a spatial integration window and feature weightings that are directly learned from the data. We applied this model to natural image responses in monkey V1 and found that it achieves an accuracy comparable to the state-of-the-art convolutional neural network (CNN) while using less than half the number of parameters for its feature space and being much more interpretable. We found that the strongest contribution to normalization came from units with matching receptive fields while pooling from the surround played a much weaker role. Moreover, unlike expected from earlier work using simple stimuli, we found that oriented features were normalized preferentially by channels with similar orientation rather than pooling non-specifically. Thus, our work reveals novel mechanisms of divisive normalization under stimulation with natural stimuli.

II-23. Bayesian inference for an exploration-exploitation model of human gaze control

Noa Malem-Shinitski NOA.MALEM@UNI-POTSDAM.DE
Sebastian Reich SREICH@MATH.UNI-POTSDAM.DE
Ralf Engbert RALF.ENGBERT@UNI-POTSDAM.DE

University of Potsdam

Understanding human gaze control and the underlying saccadic selection process is an important question in cognitive-neuroscience with many interesting applications in different fields from psychology to computer vision. In recent years there has been a considerable progress in the development of spatial models for fixation distribution prediction, known as saliency maps, with the goal of predicting where observers fixate on average, given a specific image. However, typical saliency models do not attempt to explain the spatial interaction between fixations and the temporal structure of a sequence of fixations, known as scanpaths.

In comparison to the body of research on saliency maps, the field of dynamic fixation modeling is at its first stages. In this work we present a new parametric, data-driven model for scanpaths generation. This model attempts to capture the wide range of saccade lengths found in experimental studies. We develop a discrete-time probabilistic generative model, with a Markovian structure, where at each step the next fixation selection, given the current fixation and a given saliency map, is generated from one of two strategies – exploitation or exploration. An exploitation step results in a random walk with a small step size around the current fixation location. In an exploration step the location of the next fixation is drawn from the saliency map, under some constrains. The simplicity of the model allows us to implement an efficient Bayesian inference for the hyperparameter using an Hamiltonian Monte Carlo (HMC) within Gibbs approach.

We demonstrate the validity our model on a dataset containing fixations from 35 participants viewing 30 different natural scene images. We fit a separate models for each participant and compare the statistics of the data generated from the model to the experimental data and evaluate our results against the existing state of the art models.

II-24. Invariant, stochastic, and developmentally regulated synapses constitute the C. elegans connectome from isogenic individuals

Daniel Witvliet[1] TEILVTIW@GMAIL.COM
Ben Mulcahy[2] BENMULCAHY406@GMAIL.COM
James K Mitchell[3] JMITCHELL@G.HARVARD.EDU
Yaron Meirovitch[4] YARON.MR@GMAIL.COM
Daniel R Berger[3] DANIELBERGER@FAS.HARVARD.EDU
Douglas Holmyard[2] DHOLMYARD@MTSINAI.ON.CA
Richard L Schalek[3] RSCHALEK@MCB.HARVARD.EDU
Steven J Cook[5] STEVEN.COOK@GMAIL.COM
Wan Xian Koh[1] WANXIAN.KOH@MAIL.UTORONTO.CA
Marianna Neubauer[3] M.NEUBAUER330@YAHOO.COM
Christine Rehaluk[2,1] CHRISTINEREHALUK@YAHOO.CA
ZiTong Wang[1] JERRY.WANG95@YAHOO.CA
David Kersen[3] DAVIDKERSEN712@G.HARVARD.EDU
Andrew D Chisholm[6] CHISHOLM@UCSD.EDU
Nir Shavit[4] SHANIR@CSAIL.MIT.EDU
Jeffrey W Lichtman[3] JEFF@MCB.HARVARD.EDU
Aravinthan Samuel[3] SAMUEL@G.HARVARD.EDU
Mei Zhen[7] ZHEN@LUNENFELD.CA

[1]University of Toronto
[2]Lunenfeld-Tanenbaum Research Institute
[3]Harvard University
[4]Massachusetts Institute of Technology
[5]Albert Einstein College of Medicine
[6]University of California, San Diego
[7]Lunenfeld-Tanenbaum Research Institute and University of Toronto

The developmental assembly of neural circuits is often argued to be highly optimized to wire precise and efficient networks. However, for a species to adapt and evolve, a dynamic network is essential. This is especially evident during postnatal development, when newly born neurons are integrated into functioning neurocircuits. Rules for circuit variability and remodeling are poorly characterized, because circuit-level analyses across multiple animals and developmental time points are missing. We used an isogenic C. elegans population to examine the dynamics of neural circuits during development. At birth, C. elegans has 218 neurons; 82 new neurons are incorporated into the nervous system before the end of larval development. Using serial-section electron microscopy, we mapped the synaptic connectivity in the central nervous system for seven animals spanning from birth to adulthood. These datasets reveal a five-fold increase in synapse number during larval development. Neurons with low centrality increase synaptic output, while neurons with high centrality acquire more synaptic input, suggesting that central neurons have mature output at birth, and their regulation is modulated by development or experience. The connectivity can be separated into core connections, synapses present throughout the lifetime of the animal, developmentally regulated connections, synapses present only during early or late development, and variable connections, synapses that are not conserved between individuals. The majority of new synapses are added to strengthen core connections. We propose that core connections drive hard-wired behaviors, and developmentally regulated connections exert stage-specific roles, while the surprisingly high prevalence of variable connections may confer inter-individual variability. Developmentally regulated connections appear in all circuit layers, but connections between interneurons are remarkably stable. This implies that core decision-making circuits are already established at birth, but their modulation and multi-sensory integration are refined or shaped during development.

II-25. Learning enhances sensory representations via selective temporal integration of feedforward inputs

Angus Chadwick[1] ANGUS.CHADWICK@UCL.AC.UK
Thomas Mrsic-Flogel[2] T.MRSIC-FLOGEL@UCL.AC.UK
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

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

Cortical circuits process sensory information through both feedforward and recurrent synaptic interactions. While feedforward synapses can be tuned to filter and propagate relevant information, the functional role of recurrent synaptic interactions and their changes over learning remain unclear. We recently found that as mice learn to behaviorally discriminate two full-field oriented grating stimuli, recurrent interactions between excitatory and inhibitory neurons in visual cortex reorganize, leading to improved cortical representations of those stimuli (Khan et al., 2018). How changes in recurrent interactions can achieve such an improvement, particularly for a simple orientation discrimination task involving perceptually distinct stimuli, is yet to be resolved. While the instantaneous information provided by feedforward inputs cannot be increased through recurrent processing, such information may be more faithfully preserved if integrated through time. We show analytically that, for recurrently connected networks of nonlinear units, information transmission is optimized when slow modes of dynamics align with information-carrying dimensions of feedforward inputs. Thus, the integration and transmission of feedforward information can be improved through either a selective slowing or realignment of recurrent network dynamics. We re-analyzed experimental data and found that learning involves a selective alignment of slow modes of recurrent network dynamics with task-relevant inputs. To understand how such changes in recurrent dynamics could emerge from synaptic reorganization within cortical circuits of excitatory and inhibitory neurons, we numerically computed the response dynamics and information content of networks while varying their synaptic connectivity profiles. We found that the experimental results are consistent with a model in which both excitatory-excitatory and excitatory-inhibitory synapses become more tuned to stimulus preference over learning. In contrast to other changes in synaptic connectivity which increase information, these changes produce a more selective temporal integration of feedforward inputs without introducing an overall slowing of network dynamics, thereby enhancing information transmission without compromising network stability.

II-26. A dendritic mechanism for dynamic routing and control in the thalamus

Carmen Varela[1] CARMENV.WORK@GMAIL.COM
Subutai Ahmad[2,3] SAHMAD@NUMENTA.COM

[1]Massachusetts Institute of Technology
[2]Numenta
[3]Research

A prevalent hypothesis about the thalamus is that thalamocortical cells (TC) function as ‘gated relays’, switching between two different modes of response (burst and tonic) controlled by modulators such as the corticothalamic input from layer 6 (L6). Predictive theories of perception and cognition model neocortical feedback as carrying expectations generated by learned models of the world. However, few predictive coding hypotheses consider the implications of corticothalamic (CT) feedback to thalamus, despite the fact that this is an integral component of the neocortical circuit, and despite the fact that the number of -feedback CT axons is an order of magnitude greater than feedforward TC axons. Here we propose a dendritic mechanism by which L6 corticothalamic cells could exert precise contextual modulatory control over thalamic activity. A calcium mediated plasticity rule enables TC cells to detect precise sparse codes from L6 on dendrites in a branch specific manner. We describe how the dendrite specific mechanism enables TC cells to act as a flexible multiplexer, a shifter, or a contextual filter. If L6 feedback encodes a location signal, the thalamic circuit can implement spatial and contextual selective attention mechanisms. The proposed mechanism closely matches the detailed anatomy and physiology of thalamocortical

circuits, and leads to specific experimentally testable hypotheses. We suggest that L6 corticothalamic feedback enables the thalamus to support complex routing operations that are far more powerful than traditionally assumed.

II-27. Double-dissociation between catecholaminergic and cholinergic effects on cortex-wide intrinsic correlations of neural activity

Thomas Pfeffer[1] T.PFEFFER@UKE.DE
Adrian Ponce-Alvarez[2] ADRIAN.PONCE@UPF.EDU
Ruud L Van den Brink[1] R.VAN-DEN-BRINK@UKE.DE
Andreas K Engel[1] AK.ENGEL@UKE.DE
Guido Nolte[1] G.NOLTE@UKE.DE
Gustavo Deco[2] GUSTAVO.DECO@UPF.EDU
Tobias Donner[1] T.DONNER@UKE.DE

[1]University Medical Center Hamburg Eppendorf
[2]Universitat Pompeu Fabra

The catecholaminergic (norepinephrine and dopamine) and cholinergic neuromodulatory systems widely project to the cerebral cortex and shape the operating mode of cortical circuits. Previous work in rodents and non-human primates has revealed similar effects of these neuromodulators on cortical neurons, such as the depolarization and decorrelation of pyramidal cells and an increase in the gain of their input-output functions. Here, we report a double-dissociation between the catecholaminergic and cholinergic effects on large-scale cortical circuit dynamics in humans during different behavioral contexts. We quantified changes of intrinsic correlations in band-limited magnetoencephalography (MEG) signals under placebo-controlled pharmacological interventions during an active task entailing constant visual input (but not motor output transients) as well as during “rest” (passive fixation). Elevation of catecholamine levels increased cortical correlations, but only during task, whereas increased acetylcholine levels decreased correlations, but only during rest. Multi-scale computational modelling explained the context-dependent catecholamine effects by a homogenous increase in the effective gain of all neural populations, in turn resulting from an increase in the synaptic excitation/inhibition ratio (e.g. through disinhibition). By contrast, the acetylcholine effects were due to a balanced decrease in external excitatory and inhibitory drive, accompanied by no change (or even a decrease) in gain. In sum, the context-dependent distinction between catecholamine and acetylcholine effects on large-scale cortical dynamics uncovered here provides new insight into the specific circuit effects of these neuromodulators. Those circuit effects may be the substrate of the distinct computational roles postulated by influential theories (e.g., Yu & Dayan, Neuron, 2005). Our model further provides a foundation for relating our mechanistic insights to new cellular-resolution data obtained in animals.

II-28. Rapid perceptual modulation of PPC activity predicts concomitant changes in urgent-decision accuracy

Joshua Seideman JSEIDEMA@WAKEHEALTH.EDU
Emilio Salinas ESALINAS@WAKEHEALTH.EDU
Terrence Stanford STANFORD@WAKEHEALTH.EDU

Wake Forest University

The lateral intraparietal area (LIP) is posited to play an essential role in guiding choice behavior based on sensory information. This long-held, yet currently contentious hypothesis is central to many working models of perceptual decision making, and to our present understanding of the functional role of the posterior parietal cortex (PPC) in general (Katz et al., 2016, Nature; Hanks & Summerfield, 2017, Neuron). However, it is currently unknown whether relevant sensory information is first processed by LIP neurons before it is available for incorporation into/guidance of an overt perceptual choice. We investigated this using a recently developed, urgent-decision task in which

perceptual performance varies from chance to near 100% correct as a function of sensory cue viewing time prior to choice execution (i.e., processing time; Stanford et al., 2010, Nat. Neurosci.). Using this urgent-decision task, we characterized and directly compared (1) the time course by which sensory evidence informs perceptual choice behavior, to (2) the time course by which sensory evidence informs LIP activity. Single-unit recordings revealed extremely rapid perceptual modulation of LIP neuronal activity that preceded and predicted concomitant changes in urgent-decision accuracy. In the urgent-decision task, sensory evidence modulated LIP neuronal activity at the same rate that it informed perceptual choice behavior. Even though confidence estimates were in no way explicitly required by the task, LIP activity correlated with urgent-decision accuracy and exhibited additional features characteristic of the statistical definition of confidence – adhering to the predictions of normative, as well as signal detection theory-based models of decision making (Kepecs et al., 2008, Nature; Hangya et al., 2016, Neural Comput.). Together, these empirical findings suggest that sensory evidence rapidly informs urgent perceptually-guided choice behavior by way of parietal cortex, and that confidence computations are inherent to this dynamic process.

II-29. DeepLabCut: a tool for fast, robust, and efficient 3D pose estimation

Alexander Mathis[1] AMATHIS@FAS.HARVARD.EDU
Tanmay Nath[1] NATH@ROWLAND.HARVARD.EDU
Richard A Warren[2] RAW2163@CUMC.COLUMBIA.EDU
An Chi Chen[3] ANCHI.CHEN.127@GMAIL.COM
Amir Patel[3] A.PATEL@UCT.AC.ZA
Venkatesh N Murthy[1] VNMURTHY@FAS.HARVARD.EDU
Mackenzie Mathis[1] MATHIS@ROWLAND.HARVARD.EDU
Matthias Bethge[4] MATTHIAS.BETHGE@UNI-TUEBINGEN.DE

[1]Harvard University
[2]Columbia University
[3]University of Cape Town
[4]University of Tubingen

Quantifying behavior is crucial for many applications in neuroscience. While cameras provides easy, noninvasive methods for the observation of animal behavior in diverse settings, the extraction of particular aspects of a behavior for further analysis can be highly time consuming. Recently, we presented an efficient method for markerless pose estimation based on transfer learning with deep neural networks that achieves excellent results with minimal training data. We demonstrated the versatility of this python package, called DeepLabCut, by tracking various body parts in multiple species across a broad collection of behaviors [1]. Current experiments produce vast amounts of video data, which pose challenges for both storage and analysis. Here we improve the inference speed of DeepLabCut, and show that poses can be inferred at up to 1200 frames per second (FPS). For instance, 278 x 278 images can be processed at 225 FPS on a GTX 1080 Ti graphics card. Furthermore, we show that DeepLabCut is highly robust to standard video compression (ffmpeg). Compression rates of greater than 1,000 only decrease accuracy by about half a pixel (for 640 x 480 frame size). DeepLabCut's speed and robustness to compression can save both time and hardware expenses [2]. We will illustrate DeepLabCut's workflow and capabilities based on various datasets and further demonstrate the scope of DeepLabCut by tracking 3D poses of hunting cheetahs.

[1] Mathis et al. (2018) DeepLabCut: markerless pose estimation of user-defined body parts with deep learning – doi: https://doi.org/10.1038/s41593-018-0209-y

[2] Mathis, Warren (2018) On the inference speed and video-compression robustness of DeepLabCut – doi: https://doi.org/10.1101/457242

II-30. Critical correlation between memory engrams in associative memory

Chiara Gastaldi[1] CHIARA.GASTALDI@EPFL.CH
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH
Tilo Schwalger[2] TILO.SCHWALGER@BCCN-BERLIN.DE

[1]Ecole Polytechnique Federale de Lausanne
[2]Bernstein Center for Computational Neuroscience Berlin

Human memory operates with associations. If you visited a famous place with your best friend, any postcard of that famous place will remind you of him or her. Experimental evidence (Ison et al. 2015) suggests that associative memory is stored in the area CA3 of the hippocampus, by cell assemblies that respond selectively to single concepts (Quiroga et al. 2005). Associations between different concepts are encoded in the overlap between the respective cell assemblies: assemblies relative to non-associated concepts share less than 1% of neurons, while those relative to highly associated concepts share up to 4-5% of neurons (Waydo et al. 2006, De Falco et al. 2016, Rey et al. 2018). Associative memory is traditionally modeled through attractor neural networks (ANN). Memory engrams are represented by binary patterns with low mean activity. While there is extensive literature on independent patterns, only a few studies can be found on correlated patterns (Tsodysk 1989). In particular, how the stability of a single memory pattern is affected by the correlation with other patterns remains an open question. Using mean-field approximation, we derive analytic equations for the network dynamics in the case of correlated patterns. Moreover, in the simple case of two correlated patterns, we prove the existence of a critical correlation $\hat{C}^*$ (which corresponds to a percentage of shared neurons) above which patterns are not distinguishable. For $\hat{C}>\hat{C}^*$, the stability of a single pattern is lost and the network can be either in the resting state or retrieve both patterns together, which act as a single memory. Importantly, for $\hat{C}<\hat{C}^*$ patterns can be retrieved together or as isolated concepts. Our results provide a theoretical framework that can explain the experimentally observed value of shared neurons in memory assemblies.

II-31. Bayesian Regression best explains how participants handle parameter uncertainty

Jannes Jegminat[1,2] JANNES@INI.UZH.CH
Maya A Jastrzębowska[3,4] MAYA.JASTRZEBOWSKA@EPFL.CH
Matthew V Pachai[3,5] MATTPACHAI@GMAIL.COM
Michael H Herzog[3] MICHAEL.HERZOG@EPFL.CH
Jean-Pascal Pfister[6] JPFISTER@INI.UZH.CH

[1]University of Zurich
[2]ETH Zurich
[3]Ecole Polytechnique Federale de Lausanne
[4]University of Lausanne
[5]York University
[6]University of Bern

Psychophysical experiments have shown that the human brain treats sensory uncertainty in accordance with Bayes' theorem. However, most of this psychophysical evidence comes from inference tasks about hidden variables in the presence of sensory uncertainty. Here, we consider a regression task instead and test if the Bayesian framework applies as well. We propose a new psychophysical experiment where participants have to extrapolate a curve (parabola) from a limited number of points where the curve parameter is unknown but drawn from some prior distribution. The (Bayesian) optimal solution to this problem is to perform Bayesian regression and thereby improve the generalisation performance of the prediction. We therefore compared the responses from 7 different participants to the Bayesian regression solution, as well as to other sub-optimal solutions such as maximum likelihood regression and maximum a posterior regression. We found that the participants' data are best explained with the Bayesian regression model. Overall, this study further extends the scope of the so-called Bayesian brain hypothesis by showing that parameter uncertainty is relevant for brain function.

II-32. Quantifying information conveyed by large neuronal populations

John Berkowitz[1] JACKBERKOWITZ88@GMAIL.COM
Tatyana Sharpee[2] SHARPEE@SNL.SALK.EDU

[1]University of California, San Diego
[2]Salk Institute for Biological Sciences

Quantifying mutual information between inputs and outputs of a large neural circuit is an important open problem in both machine learning and neuroscience. However, evaluation of the mutual informa- tion is known to be generally intractable for large systems due to the exponential growth in the number of terms that need to be evaluated. Here we show how information contained in the responses of large neural populations can be effectively computed provided the input-output functions of individual neu- rons can be described by a logistic function, applied to a potentially nonlinear function of the stimulus. Neural responses in this model can remain sensitive to multiple stimulus components. We describe both an approach for tractably estimating information numerically, as well as several approximations that may be more amenable to analytic treatment. These approximations are obtained by deriving a vector quantity that is guaranteed to capture all information contained in the neural responses and by expressing the mutual information as a sum of lower-dimensional conditional mutual information terms. These terms can be evaluated with a number of information estimators. The cost of this computation grows linearly in the dimension of the input, and compares favourably with other approximations.

II-33. Heterosynaptic plasticity determines the set-point for cortical excitatory-inhibitory balance

Christoph Miehl[1] CHRISTOPH.MIEHL@BRAIN.MPG.DE
Rachel Field[2] RACHEL.FIELD2@NYUMC.ORG
James D'amour[3] JIM.D'AMOUR@NIH.GOV
Robin Tremblay[2] ROBIN.TREMBLAY@MED.NYU.EDU
Illya Kruglikov[2] ILLYA.KRUGLIKOV@NYULANGONE.ORG
Bernardo Rudy[2] BERNARDO.RUDY@NYUMC.ORG
Robert Froemke[2] ROBERT.FROEMKE@MED.NYU.EDU
Julijana Gjorgjieva[1] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]New York University
[3]National Institutes of Health

Excitation in neural circuits must be carefully controlled by inhibition to regulate information processing and network excitability. During development, cortical inhibitory and excitatory inputs are initially mismatched but become co-tuned or 'balanced' with experience. However, little is known about how excitatory-inhibitory balance is defined at most synapses, or the mechanisms for establishing or maintaining this balance at specific set-points. Here we show how coordinated long-term plasticity calibrates populations of excitatory and inhibitory inputs onto mouse auditory cortical pyramidal neurons. Pairing pre- and postsynaptic activity induced plasticity at paired inputs and different forms of heterosynaptic plasticity at the strongest unpaired synapses, which required minutes of activity and dendritic Ca2+ signaling to be computed. Theoretical analysis of two different models demonstrated how the relative amount of heterosynaptic plasticity could normalize and stabilize synaptic strengths to achieve any possible excitatory-inhibitory correlation. Thus, excitatory-inhibitory balance is dynamic and cell-specific, determined by distinct plasticity rules across multiple excitatory and inhibitory synapses.

II-34. Recruitment and disruption of value encoding in models of alcohol seeking

Jocelyn Richard[1] RICHARDJ@UMN.EDU
David Ottenheimer[2] DAVID.OTTENHEIMER@JHU.EDU
Karen Wang[2] KARENWANG952@GMAIL.COM
Alexandria Haimbaugh[2] ALEXHAIMBAUGH@GMAIL.COM
Patricia Janak[2] PATRICIA.JANAK@JHU.EDU

[1]University of Minnesota
[2]Johns Hopkins University

Reward-seeking behaviors can be elicited by a variety of reward-related mental representations. Addiction involves the biased recruitment of some of these representations (e.g. cue or action values) at the expense of others (e.g. outcome value representations). Understanding how these neural representations are altered in models of addiction will provide insight into both addiction itself, and how these representations are computed and instantiated as adaptive reward-seeking behavior. We previously showed that cue-evoked activity in rat ventral pallidum (VP) robustly encodes the value of cues trained under both Pavlovian and instrumental contingencies (Richard et al., 2016, 2018). Here, we assessed VP representations of cue value in rats trained with a Pavlovian conditioned stimulus (CS+) that predicted alcohol delivery, or an instrumental discriminative stimulus (DS) that predicted alcohol availability if the rat entered the reward port during the cue. We then fit linear discriminant analysis models trained on the activity of VP single units or ensembles, to assess the strength of VP encoding of value for DS and CS cues. Decoding of cue value based on VP firing was blunted for an alcohol CS+ versus an alcohol DS, as well as in comparison to a sucrose DS or CS+. Non-associative alcohol exposure had opposing effects on VP encoding of cue value for a sucrose DS versus a sucrose CS+ at all ensemble sizes, enhancing decoding accuracy for the DS and reducing decoding accuracy for the CS+. Overall, we find that VP encoding of cue value is bidirectionally altered by exposure to and conditioning with alcohol, depending on the associative structure of the task, which may reflect biased engagement of specific mental representations in alcohol abuse.

II-35. Active navigation enhances visual responses in mouse cortex

Efthymia (Mika) Diamanti EFTHYMIA.DIAMANTI.11@UCL.AC.UK
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK
Aman Saleem AMAN.SALEEM@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

Vision is crucial not only for viewing the world but also for navigating in it. Do passive viewing and navigation result in different responses across the visual cortex? We used 2-photon calcium imaging in head-restrained mice to record neural activity across seven areas in visual cortex in three conditions: (1) passive viewing of drifting gratings at varying orientations; (2) active navigation in a virtual reality (VR) corridor comprising two visually matching segments; (3) passive viewing of previous VR sessions. We first compared passive viewing of drifting gratings to active navigation in VR. Both conditions elicited robust responses across visual areas. Nevertheless, cells that responded to drifting gratings tended to be inactive in VR, and vice versa. In addition, different visual areas favored the two conditions differently: for instance, primary visual cortex (V1) favored the drifting gratings, whereas parietal areas A and AM favored the VR. We next asked whether responses in VR differed during active navigation versus passive viewing of the same VR scenes. During active navigation, spatial context modulated responses in all visual areas, consistent with previous findings in V1 (Saleem et al. Nature 2018). Instead, during passive viewing the influence by spatial context dramatically decreased. In addition, responses during passive viewing were less reliable than during active navigation. These results indicate that active navigation drives a distinct population of neurons in visual cortex and enhances the reliability of visual responses. We conclude that visual cortex is capable of distinguishing between visual inputs controlled by self-motion and passive viewing.

II-36. Serial cells track discrete episodic events and their temporal relationships

Chen Sun[1] CSUN@MIT.EDU
Wannan Yang[2] S1449361@SMS.ED.AC.UK
Jared Martin[1] JMARTIN1@MIT.EDU
Susumu Tonegawa[1] TONEGAWA@MIT.EDU

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

An emerging view is that episodic experience is tracked by the hippocampus as a sequence of moment to moment continuous variations but whether this same episode is separately tracked as a sequence of discrete events and the pure sequential relationships between them, is unknown. We designed a task in which mice experience a series of materially indistinguishable events. We report hippocampal 'serial cells' that organize an episodic experience around discrete subdivisions ("events"). The serial code is unaffected by arbitrary changeability within the events, reflecting their flexible and abstract nature. The discrete serial code is conjunctively represented with the continuous spatial code but can be independently perturbed. The serial code may be one of the fundamental building blocks for representing an episode.

II-37. Synaptic plasticity in correlated balanced networks

Alan Akil[1] AAKIL@MATH.UH.EDU
Robert Rosenbaum[2] RROSENB1@ND.EDU
Krešimir Josić[1] JOSIC@MATH.UH.EDU

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

Neurons in the cortex exhibit temporally irregular, but correlated spiking during ongoing sensory experiences. It is still unclear what mechanisms drive the emergence of global activity patterns, while supporting local fluctuations. Models that exhibit an emergent balance between excitation and inhibition produce individual neuron activity and responses remarkably similar to those in the cortex. However, in their original version, they only lead to asynchronous states. Recently, a number of mechanisms that support correlated, balanced activity across the population have been proposed. Such correlated states can be generated intrinsically or extrinsically. Yet no theories about how this activity is maintained and shaped by ongoing changes in synaptic weights exists. How do emergent patterns in correlated activity drive changes in synaptic architecture, and how do these changes, in turn, shape the activity of the network? Under what conditions are correlations amplified or dampened by changes in synaptic strengths? Could changes in synaptic architecture self-amplify, and drive the network out of balance? These questions are nontrivial, as the details of the recurrent architecture and plasticity rule determine changes in weights and correlations. To answer these questions, we develop a general theory of plasticity in correlated balanced networks. We show that, in general, balance is attained and maintained both in asynchronous and correlated states. We find that the presence of externally generated correlations in the network does not drive significant changes in the synaptic connectivity and firing rates during spontaneous activity. However, for stimulus-evoked activity, changes in inhibitory to excitatory synapses can considerably dampen both firing rates and correlations across the population. Our general computational and theoretical framework allows us to describe how stimuli interact with correlated activity to drive changes in synaptic connectivity, which in turn, shape activity patterns in balanced neuronal networks.

II-38. Identification of unique pre-epileptic states via non-negative tensor decomposition of single unit recordings

Jordan Sorokin JORSOROKIN@GMAIL.COM
Surya Ganguli SURYA.GANGULI@GMAIL.COM
John Huguenard HUGUENARD@GMAIL.COM

Stanford University

The brain has a remarkable ability to dynamically switch between cognitive states, although this process is poorly understood. Absence epilepsy, characterized by sudden high voltage oscillations within the thalamocortical circuit, provides a unique window to study such state switching. We recently discovered that in rodents with absence epilepsy, seizures are preceded (e.g. the pre-ictal period) by elevated beta (20 – 40 Hz) power in the electrocorticogram (ECoG). Here we investigated this transition in more detail using single unit recordings within the thalamocortical circuit (Fig. 1A). PCA on trial-averaged neural firing revealed stereotyped activity trajectories during the transition into seizures, characterized by a migration away from a baseline point attractor (Fig. 1B). This pre-ictal migration was localized to thalamic neurons (Fig. 1C), despite previous literature suggesting that cortical activity initiates absence seizures. Notably, although trial-averaged trajectories were consistent across animals, individual units showed heterogeneous activity across trials (Fig. 1D). In order to identify whether this variability related to different well-defined pre-ictal states, we decomposed our data via non-negative canonical polyadic decomposition (nnCPD) (Fig. 2A). nnCPD uncovered latent firing patterns (Fig. 2B-C) that explained different subsets of trials, which were associated with diverse pre-ictal thalamic firing patterns and seizure lengths (Fig. 2D-E). Interestingly, the projection of trial-averaged neural activity for longer seizures resided in the same space at baseline as that of the pre-ictal transition for the equivalent projection of shorter seizures (Fig. 2F), suggesting that brief seizures arise from a distinct pre-ictal state. Our results are the first to show that pre-ictal activity in generalized epilepsy can localize to the thalamus, vary between seizures, and predict seizure severity. We hope to use this method for improving closed-loop seizure interruption by adapting stimulus parameters according to anticipated seizure type/severity. Moreover, we believe our methodology extends more broadly to the analysis of state bifurcations.

II-39. Inferring the parameters of neural simulations from high-dimensional observations

Marcel Nonnenmacher[1] MARCEL.NONNENMACHER@TUM.DE
Jan-Matthis Lueckmann[2,1] JAN-MATTHIS.LUECKMANN@CAESAR.DE
Giacomo Bassetto[2,1] GIACOMO.BASSETTO@CAESAR.DE
Pedro J Goncalves[2] PPJGONCALVES@GMAIL.COM
Jakob H Macke[1] MACKE@TUM.DE

[1]Technical University of Munich
[2]Research Center Caesar

Many models in neuroscience, such as networks of spiking neurons or complex biophysical models, are defined as numerical simulators. This means one can simulate data from the model, but calculating the likelihoods associated with specific observations is hard or intractable, which in turn makes statistical inference challenging. So-called Approximate Bayesian Computation (ABC) aims to make Bayesian inference possible for likelihood-free models. However, standard ABC algorithms do not scale to high-dimensional observations, e.g. inference of receptive fields from high-dimensional stimuli.

Here, we develop an approach to likelihood-free inference for high-dimensional data, where we train a neural network to perform statistical inference given adaptively simulated data sets. The network is composed of layers performing non-linear feature extraction, and fully connected layers for non-linear density estimation. Feature extraction layers are either convolutional or recurrent in structure, depending on whether the data is high-dimensional in space or time, respectively. This approach makes it possible to scale ABC to problems with high-dimensional inputs.

We illustrate this method in two canonical examples in neuroscience. First, we infer receptive field parameters of a V1 simple cell model from neural activity resulting from white-noise stimulation, a high-dimensional stimulus in the space domain. Second, we perform Bayesian inference on a Hodgkin-Huxley model of a single neuron, given full voltage traces resulting from intracellular current stimulation. On both applications, we retrieve the posterior distribution over the parameters, i.e. the manifold of parameters for which the model exhibits the same behaviour as the observations. Our approach will allow neuroscientists to leverage the power of deep neural networks to link high-dimensional data to complex simulations of neural dynamics.

II-40. Inferring the microstructure of pitch reinforcement learning in songbirds

Anja Tamara Zai[1,2] ZAIA@ETHZ.CH
Sophie Cave-Lopez[1,2] SOPHIE@INI.ETHZ.CH
Anna E Stepien[1,2] STEPIEN@INI.ETHZ.CH
Nicolas Giret[3] NICOLAS.GIRET@U-PSUD.FR
Richard Hahnloser[2] RICH@INI.PHYS.ETHZ.CH

[1]University of Zurich
[2]ETH Zurich
[3]University Paris Saclay

Reinforcement is one of the most effective means to selectively change motor behaviors in humans and animals. Yet, how does reinforcement act upon plastic motor gestures on a trial-by-trial basis? In songbirds subjected to operant conditioning of vocal pitch, we fit trial-by-trial song data to a simple stochastic model that can account for a broad range of experimental findings. Incorporating diverse sources of behavioral variability, our Kalman filter model produces excellent fits to pitch conditioning data even when learning trends are non-monotonic due to diurnal rhythms of pitch. Furthermore, our model can account for the firing behavior of dopaminergic neurons projecting to a basal ganglia homologue, Area X, as recently reported (Gadagkar et al. 2016). To test whether the role of Area X is to process performance errors or merely reward prediction errors, we conduct experiments in deaf zebra finches that we expose to reinforcement by switching off the light for low (or high) pitched renditions of a particular song syllable. We find that deaf birds reliably adapt their pitch in response to light-off and that such visual reinforcement of pitch critically depends on Area X. Hence, birds do not require auditory feedback for trial-and-error learning of pitch and reinforcement learning can occur without evaluation of vocal performance. The main benefits of our model are to provide a rich characterization of motor learning and to provide estimates of basal ganglia function via realistic modeling of dopaminergic neurons.

II-41. Abnormal locus coeruleus sleep activity alters sleep signatures and impairs spatial memory

Kevin Swift[1] SWIFTKE@UMICH.EDU
Brooks Gross[2] BROOKSGROSS@UCLA.EDU
Michelle Frazer[2] FRAZERMI@UCLA.EDU
David Bauer[1] DSBAUER@UMICH.EDU
Kyle Clark[1] CLARKLY@UMICH.EDU
Elena Vazey[3] EVAZEY@UMASS.EDU
Gary Aston-Jones[4] GSA35@RUTGERS.EDU
Yong Li[5] YONG.LI@BRISTOL.AC.UK
Anthony Pickering[5] TONY.PICKERING@BRISTOL.AC.UK
Susan Sara[6] SUSAN.SARA@COLLEGE-DE-FRANCE.FR
Gina Poe[2] GINAPOE@UCLA.EDU

[1]University of Michigan
[2]University of California, Los Angeles
[3]University of Massachusetts Amherst
[4]Rutgers University
[5]University of Bristol
[6]College de France

Sleep is critical for proper memory consolidation. The locus coeruleus (LC) releases norepinephrine throughout the brain except when the LC falls silent throughout rapid eye movement (REM) sleep and prior to each non-REM (NREM) sleep spindle. We hypothesize that these transient LC silences allow the synaptic plasticity that is necessary to incorporate new information into pre-existing memory circuits. We found that spontaneous endogenous LC activity within sleep spindles triggers a decrease in spindle power. By optogenetically stimulating norepinephrine-containing LC neurons at 2 Hz during sleep, we reduced sleep spindle occurrence, as well as NREM delta power and REM theta power, without causing arousals or changing sleep. Stimulating the LC during sleep following a hippocampus-dependent food location learning task interfered with consolidation of newly learned locations and reconsolidation of previous locations, disrupting next-day place cell activity. Further, the fields for place cells encoding information for the reconsolidated memory selectively expanded in size, whereas the fields for the newly consolidated memory did not. The LC stimulation-induced reduction in NREM sleep spindles, delta, and REM theta as well as reduced ripple-spindle coupling all correlated with decreased hippocampus-dependent performance on the task. Thus, we find periods of LC silence during sleep following learning are essential for normal spindle generation, delta and theta power, and consolidation of spatial memories. We believe this is the first work demonstrating that manipulation of the noradrenergic system during sleep can interfere with memory consolidation.

II-42. Population-level signatures of probabilistic computation

Edoardo Balzani EB162@NYU.EDU
Cristina Savin CS5360@NYU.EDU

New York University

Cortical responses exhibit coordinated fluctuations of population activity, in the form of network oscillations, low-dimensional latent dynamics and structured noise correlations. Experiments suggest that these fluctuations are actively shaped by stimuli, context, brain state and often correlate with behavioral performance. Still, their computational significance remains unclear. From the perspective of an experimentalist studying a single-cell encoding a stimulus, such global fluctuations seem to drown neural responses in noise, strongly limiting brain computations. Instead, here we argue that a range of observations on network fluctuations can be understood as signatures of efficient probabilistic computation. We propose a novel spiking network model for Bayesian inference which com-

bines a distributed sampling scheme (Savin & Deneve, 2014) with a new class of stochastic dynamics, designed to increase computational speed and efficiency (Hennequin et al, 2014). When probed using experimental-like input manipulations, the model reproduces a range of biological observations concerning the modulation of single-cell and population responses by uncertainty and stimulus statistics. First, population responses show tuning- and stimulus-dependent noise correlations. Second, the model reproduces the previously unexplained experimental observation that the peak frequency and power of cortical oscillations is modulated by uncertainty. Furthermore, the model predicts systematic changes in sparseness and oscillatory structure when manipulating behavioral uncertainty or when contrasting natural vs. artificial stimulus statistics.

II-43. Learning sequence disambiguation with synaptic traces in associative neural networks

Ramon Martinez RHMM@KTH.SE
Anders Lansner ALA@CSC.KTH.SE
Pawel Herman PAHERMAN@KTH.SE

KTH Royal Institute of Technology, Stockholm

In this work we study the problem of sequence disambiguation in an associative neural network with meta-stable attractor states. The network is capable of encoding sequences with the associative Bayesian Confidence Propagation Neural Network (BCPNN) learning rule through the use of synaptic z-traces. The problem of sequence disambiguation is parameterized by the the overlap in activation between two competing sequences. The sensitivity of the system to noise under various regimes is systematically tested. The network robustly differentiates between considerably overlapping sequences. The performance is non-monotonically modulated by the time constant of synaptic traces and a range for optimal disambiguation is shown as a function of noise level.

II-44. Lingering generalization in humans learning a complex skill

Gautam Agarwal[1] AGARWALLED@GMAIL.COM
Tiago Quendera[1] TIAGO.QUENDERA@NEURO.FCHAMPALIMAUD.ORG
Mattia Bergomi[1] MATTIA.BERGOMI@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[2] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Research
[2]Champalimaud Centre for the Unknown

Learning a complex skill requires traversing a potentially enormous search space. While reinforcement learning (RL) algorithms approach human levels of performance in complex tasks, they require much more training to do so. This may be because humans draw from extensive prior knowledge that constrains search problems, allowing them to more rapidly discover solutions. Importantly, the statistics that underlie this learning process are both poorly understood and hard to investigate in the large state spaces found in most complex tasks.To address this, we have designed a game-based variation of the traveling salesman problem, a standard NP-hard search problem. By changing the number and arrangement of targets, we densely sample a high-dimensional stimulus space. For each n-target stimulus, the reward associated with each possible action sequence can be visualized as an n-dimensional hypercube with roughly n! local maxima. We compared the behavior of human subjects with that of RL agents known as deep-Q networks (DQN's) in this task. We find that human subjects rapidly learn to solve the task by arriving at a strategy that exploits task invariances and generalizes to previously unencountered stimuli. However, as task dimensionality increases, this strategy leads to subjects getting "stuck" in an increasingly suboptimal region of the value landscape, a mistake that is corrected only after repeated failure. In contrast, the RL agents initially learn the task slowly, in part due to their inability to generalize their knowledge, but subsequently avoid getting stuck in local maxima. Thus, humans and agents improve in complementary ways, respectively showing flexibility in either the application or modification of their policy. Introducing a structural prior that allows the RL agent to exploit the translational invariance of the problem (as do our subjects) partially infuses it with both forms of flexibility.

II-45. Organization and modulation of cortical representations for decision-making in a dynamic environment

Selmaan Chettih SELMAAN.CHETTIH@GMAIL.COM
Shih-Yi Tseng SHIHYI_TSENG@G.HARVARD.EDU
Christopher Harvey HARVEY@HMS.HARVARD.EDU

Harvard University

Animals dynamically acquire knowledge of their environment, and combine this with new sensory information to make adaptive decisions in changing environments. However, the neural mechanisms underlying this cognitive flexibility, particularly at the population level, are unclear. We developed a new task for head-restrained mice navigating a virtual Y-maze, which is compatible with simultaneous two-photon imaging and optogenetics approaches. Abstract visual cues indicated which maze arm would contain reward, but the rewarded cue-action mapping (‘rule’) was changed unexpectedly within single sessions. By task design, mice had to remember both cue and choice until delayed feedback of the trial’s outcome was provided, in order to update or maintain an internal estimate of the underlying rule (‘belief’). Mice adapted to rule switches by gradually transitioning through a period of exploratory behavior before reliably following the new rule. Our task thus dissociated ‘belief’ from the identity of cues and choices themselves.

We have studied where task-related information is encoded, and how ‘belief’ modulated representations of cue and choice. We imaged 1,000 neurons simultaneously for fields-of-view collectively tiling primary and secondary visual areas, parietal cortex, and retrosplenial cortex. Neurons with different projection targets were identified by retrograde labeling of orbitofrontal cortex, dorsal striatum, and anterior cingulate cortex. We analyzed mesoscopic-scale activity to identify distinct modes of task-locked activity, which indicated functional specialization of cortical regions (e.g. for feedback-related activity). We then developed behavioral estimates of the mouse’s ‘belief’, which exhibited substantial correlations with single-trial neural encoding of cue and choice. For example, decoding of past cue and choice from neural activity when awaiting outcome feedback was significantly decreased when the animal was uncertain (low magnitude of ‘belief’). Further analysis of these data will provide novel characterization of the changes in cortical representations accompanying dynamically accumulated knowledge and flexible cognition.

II-46. Signatures of low-dimensional neural predictive manifolds

Stefano Recanatesi[1] STEFANO.RECANATESI@GMAIL.COM
Matthew Farrell[1] MSF9@UW.EDU
Guillaume Lajoie[2] LAJOIE@DMS.UMONTREAL.CA
Sophie Deneve[3] SOPHIE.DENEVE@ENS.FR
Mattia Rigotti[4] MRIGOTT@US.IBM.COM
Eric Shea-Brown[1] ETSB@UW.EDU

[1]University of Washington
[2]Universite de Montreal
[3]Ecole Normale Superieure
[4]IBM

Many of the recent advances of neural networks in time-dependent applications such as natural language processing hinge on the use of representations obtained by predictive models. This success seems to correlate

with the emergence of low-dimensional representations of latent structure, when networks are trained to perform semantic relational tasks [1,2]. Motivated by the recent theoretical proposal that the hippocampus performs its role in sequential planning by organizing semantically related episodes in a relational network [3], we investigate the hypothesis that this organization results from learning a predictive representation of the world. Using an artificial recurrent neural network model trained with predictive learning on a simulated spatial navigation task, we show that network dynamics exhibit low dimensional but non-linearly transformed representations of sensory input statistics. These neural activations that are strongly reminiscent of the place-related neural activity that is experimentally observed in the hippocampus and in the entorhinal cortex [4,5]. We establish a link between place-related activity and the low-dimensional latent structure formed by predictive learning by using novel measures of representation dimensionality. More precisely, we show that this relationship can be explained by computing a *dimensionality gain* between linear algebraic and intrinsic dimensionalities of network activity. Furthermore, we provide theoretical arguments as to why predictive learning objectives necessarily imply the emergence of low-dimensional representations. We thus suggest a unifying mechanism to aid the interpretation and analysis of several experimental observations in a commmon framework.

II-47. Combined visual and patterned optogenetic stimulation of ferret visual cortex reveals that cortical circuits respond to independent inputs in a sublinear manner

Shen Wang[1] SWBUMEM@BRANDEIS.EDU
Kenneth Miller[2] KEN@NEUROTHEORY.COLUMBIA.EDU
Stephen Van Hooser[1] VANHOOSR@BRANDEIS.EDU

[1]Brandeis University
[2]Columbia University

A recently proposed theory, the supralinear stabilized network (SSN), has shed light on how canonical computations may be implemented in neural circuits. The SSN shows that, given supralinear neuronal input/output functions, network stabilization requires that recurrent input largely cancel or "loosely balance" feedforward input. This leads to sublinear summation of responses to multiple stimuli. However, for weak stimuli, little or no stabilization is needed and response summation becomes linear or supralinear. Across multiple visual cortical areas, responses to presentation of multiple stimuli are typically closer to the average than the sum of the responses to the individual stimuli, that is, response summation is sublinear (except for very weak stimuli). In most cases, it is not clear whether the sublinear summation is created by the local cortical circuit or is already present in the inputs to that circuit. Combining visual and nonspecific optogenetic stimuli, previous studies found sublinear summation in monkey V1 and linear summation in mouse V1. Here we directly test a cortical component to the sublinearity by using patterned optogenetic stimuli to drive different sets of cortical neurons, or combined visual and optogenetic stimuli. Both methods should provide independently driven inputs to cortex and allow isolation of the cortical circuit's summation properties. Preliminary results showed that simultaneous optogenetic stimulation targeting independent orientation columns produced sublinear response summation. Combined visual and patterned optogenetic stimulation also elicited sublinear responses. We also found that even though optogenetic inputs were local to specific orientation columns, the resulting neural responses did not respect orientation column boundaries. These results show that the cortex integrated inputs sublinearly and support the hypothesis that sensory cortex operates in the SSN regime, though further experiments are needed to test this hypothesis. Results of ongoing such experiments will be presented.

II-48. The role of the corpus callosum and behavioral state in synaptic inter-hemispheric correlations

Ilan Lampl ILAN.LAMPL@WEIZMANN.AC.IL
Yonatan Katz YON.KATZ@WEIZMANN.AC.IL
Yael Oran YAELORAN@GMAIL.COM
Michael Sokoletsky MICHAEL.SOKOLETSKY@WEIZMANN.AC.IL

Weizmann Institute of Science

Synchronized neuronal activity is a major hallmark of cortical physiology, and is known to exist both locally and globally. Whereas local synchrony is expected due to common inputs to nearby cells, the origin of global synchrony is much less clear. An interesting example is that of interhemispheric correlations, which are suggested to stem from projections connecting homologous areas across the two hemispheres. Yet, it is unknown how these correlations are reflected in membrane potential, how deeply involved is the corpus callosum in mediating them as opposed to other potential mechanisms (such as common thalamic and neuromodulatory inputs) and how they might be affected by different behavioral states. To address these questions, we performed paired intracellular recordings from the two barrel cortices as well as calcium imaging of transcallosal axons in awake head-fixed mice. We found that spontaneous fluctuations in membrane potential were significantly correlated across the hemispheres. Although on average synaptic activity in each hemisphere neither preceded nor lagged the other, it stochastically shifted between preceding and lagging during each recording. To find how interhemispheric correlations depend on brain state, we tracked the mouse's pupil size, whisking and locomotion. During quiescent periods interhemispheric correlations were significantly higher compared to active states, suggesting that these correlations are not driven by ascending thalamic inputs. To manipulate corpus callosum activity, we unilaterally silenced callosal terminals with hM4D, which reduced the correlations in magnitude and, importantly, in symmetry. Finally, to find how the activity of transcallosal axons is modulated by state we imaged axonal calcium signals. Interestingly, axonal activity was negatively correlated with behavioral activity, which indicates transcallosal axons could be mediating the activity-dependence of correlations. In summary, we show that interhemispheric correlations depend on the corpus callosum and that both are reduced during active behavioral states.

II-49. Learning temporal structure of the input with a network of integrate-and-fire neurons

Lyudmila Kushnir[1,2] KUSHNIRLV@GMAIL.COM

[1]ENS
[2]LNC2, DEC

One important task of the brain is to look for structure in external input. While the detection of input patterns have been extensively studied, the time dimension remains least explored. We study a network of integrate-and-fire neurons with several types of recurrent connections that learns the structure of its time-varying feedforward input by attempting to efficiently represent this input with spikes. The efficiency of the representation arises from incorporating different time scales in a decoder, which is implicit in the learned synaptic connectivity of the network. While in the original work of (Boerlin, 2013) and (Brendel, 2017) the structure learned by the network was the low-dimensionality of the feedforward input, in the present work it is its temporal dynamics. The network learns to achieve this efficiency by adjusting its synaptic weights depending on presynaptic and postsynaptic activity, so that for any neuron in the network, the recurrent input cancels the feedforward for most of the time. We show that to learn that the input follows a linear differential equation of the first order, the network needs to have at least two types of slow synaptic currents with different time scales. When endowed with these two time scales, the network automatically learns to predict the derivative of its inputs, while more synaptic time scales capture higher derivatives. Spikes occur only in case of significant deviation of the input signals from this prediction. The better the predictions, the less spikes are needed to represent the signals, leading to massive improvement in efficiency. We test the learning rule on synthetic and naturalistic time-varying multi-dimensional input and demonstrate that

the structure of the input is reflected in the shape of postsynaptic potentials. Learnt decoding filters extract the spatio-temporal features best describing the statistics of the input, with potential application to any sensory signal with temporal dynamics.

II-50. Recurrent neural networks as a model organism to study multi-task decision making

Laura Driscoll[1] LNDRISCOLL@GMAIL.COM
Guangyu Robert Yang[2] GYYANG.NEURO@GMAIL.COM
Krishna Shenoy[1] SHENOY@STANFORD.EDU
David Sussillo[3,1] SUSSILLO@GOOGLE.COM

[1]Stanford University
[2]Columbia University
[3]Google AI

Human thought is inherently flexible. Preliminary work has identified how both artificial and biological networks perform context dependent operations within the framework of a single task. However, it is unknown how computations for many tasks interact within a single network. To advance our understanding of the flexibility of biological networks, we are interested in how computations for multiple tasks are implemented within a single network of neurons. How does the similarity or dissimilarity of the task structure for two trained tasks affect the structure of computation within a fixed set of network connections? If both tasks have a working memory period, does that imply the neural activity will share particular features? To address these questions and others like them, we employ a set of 20 tasks for artificial recurrent neural networks (RNNs) to perform. We study network activity during the performance of each task as a dynamical system, governed by the initial state and the inputs to the system. We found each task rule input produces a stable fixed point. Collectively, these fixed points are organized in neural activity space according to task similarity. Following stimulus onset, hidden unit activity evolves in mostly separate subspaces for each task. Training on multiple tasks compared to training on a single task changes the dimensionality of working memory tasks, suggesting multi-task training could change the algorithms that perform these computations. We will further characterize these effects in order to develop hypotheses for testing in biological cortical networks during the performance of multiple tasks.

II-51. Representation of latent states in a non-spatial cognitive map

Cheng Chen CHENG.CHEN@NDCN.OX.AC.UK
Xudong Wang XUDONG.WANG@NDCN.OX.AC.UK
Mark Walton MARK.WALTON@PSY.OX.AC.UK
Tim EJ Behrens BEHRENS@FMRIB.OX.AC.UK

University of Oxford

Understanding the implicit relationships between states in a non-spatial cognitive map of task space allows for quick and flexible inference on the basis of sparse observation (Gershman SJ & Niv Y 2010, Wilson RC & Niv Y 2011, Behrens TEJ, et al. 2018). Such relational reasoning, which depends on the unobservable latent states, plays a key role in enabling subjects to quickly predict changes of observable states before experiencing them. Recent research has proposed that ventral prefrontal cortex (vPFC) plays a key role in enabling fast inference and flexible behavior by doing model-based credit assignment (Walton ME, et al. 2010) and representing latent states (Wilson RC, et al. 2014). However, little is known about how the neural codes of latent states are represented in the vPFC and how this representation can be used for credit assignment. To investigate these questions, we developed a stimulus-guided probabilistic reversal learning task with anti-correlated world structures, in which rats could learn 4 time-varying stimulus-response mappings separately or could alternatively learn the 1 latent variable

that controlled all the 4 mappings. We demonstrate that rats learn this global structure using both model-based analyses and targeted tests in which we break the world structure. Their choices demonstrate rapid learning and flexible inferences. Our initial neural data recorded by wireless silicon probes in PrL and lOFC demonstrate that the latent state impacts on the coding of other task variables in both regions. Intriguingly, some lOFC neurons suggest credit assignment to the latent state and use of the latent state to control credit assignment to stimulus-response associations.

II-52. Arousal unlocks interneuron heterogeneity in olfactory codes

Hannah Bos HANNAH.BOS@POSTEO.DE
Anne-Marie Oswald AMOSWALD@PITT.EDU
Brent Doiron BDOIRON@PITT.EDU

University of Pittsburgh

How are heterogeneities in cortical circuits reflected in neural activity? Here we explore how a spatial heterogeneity in interneuron circuits interacts with top-down inputs to shape cortical responses and facilitate state dependent distributed computing. Recently we have reported a density gradient of Somatostatin (SST) interneurons across Layer 3 (L3) of the olfactory (piriform) cortex, with more SST-cells in the rostral versus the caudal end. We use established excitatory and inhibitory circuitry in piriform cortex to build a network model of spiking neurons and derive mean-field theory to investigate the dynamical consequences of this SST gradient. The initial model prediction is that the SST gradient is reflected in an opposing gradient of pyramidal cell activity. However, experimental findings show spatially homogeneous activity as assayed by Fos expression in piriform cortex when mice are at rest in the home cage. Extending the model to include homeostatic plasticity of inhibition from Parvalbumin (PV) interneurons onto pyramidal cells compensates for the SST-cell gradient so that network activity is homogeneous across the piriform axis, in agreement with the data. This balance between SST-cell activity and PV-cell homeostatic plasticity allowed us to make a targeted prediction. Top-down modulation resulting from arousal or movement has been associated with activation of Vasoactive Intestinal Peptide (VIP) interneurons in other cortices. Since VIP-cells inhibit SST neurons, we predict that arousal-mediated activation of VIP-cells suppresses SST-cells and unlocks the PV inhibitory gradient across piriform cortex. Pyramidal cell activity is subsequently enhanced on the rostral and suppressed on the caudal end. This prediction is confirmed by a gradient of Fos expression, following active exploration of a novel environment. The rostral end of piriform cortex preferentially projects to the olfactory bulb. We hypothesize that the amplification of rostral activity with arousal modulates responsiveness of the olfactory bulb during active exploration in novel environments.

II-53. A simplified model of a pyramidal neuron network derived from canonical correlation analysis

Xinyuan Zhao[1] 1500017808@PKU.EDU.CN
Cengiz Pehlevan[2] CPEHLEVAN@SEAS.HARVARD.EDU
Anirvan Sengupta[3] ANIRVANS.PHYSICS@GMAIL.COM
Dmitri Chklovskii[4,5] DCHKLOVSKII@FLATIRONINSTITUTE.ORG

[1]Peking University
[2]Harvard University
[3]Flatiron Institute and Rutgers University
[4]Flatiron Institute
[5]New York University

As neural networks evolved for competitive behaviorally relevant tasks, to model their function, it is natural to take a normative approach, i.e. view network dynamics as optimizing a principled objective function. Such normative approach often leads to simplified and interpretable models. Neural networks with linear point neurons, for example, have been derived from various PCA cost functions, elucidating how unsupervised learning can arise from self-organization of simple synaptic updates. Here, we extend this approach to neurons with dendritic structure. We present a normative model of a simplified pyramidal neuron network, including activity dynamics and synaptic plasticity, as an implementation of an online canonical correlation analysis (CCA) algorithm, derived from a novel CCA objective function. Given two related datasets, CCA finds the subspace which maximizes a correlation between their projections onto that subspace. In the model, the two datasets are streamed to the network as activity vectors of the upstream neurons in feedforward and feedback pathways. At each time step, the two streamed activity vectors are projected onto the common subspace by multiplying by the synaptic weights and summing in basal and apical dendrites. Then, the pyramidal neurons add these projections and output the sum as their firing rate. In addition, synaptic weights are updated according to biologically-plausible Hebbian and anti-Hebbian learning rules.

II-54. Deciphering the neural coding of the motor cortex with a multidirectional reaching task in mice

Gregorio Galinanes GREGORIO.GALINANES@UNIGE.CH
Claudia Bonardi CLAUDIA.BONARDI@UNIGE.CH
Daniel Huber DANIEL.HUBER@UNIGE.CH

University of Geneva

In spite of being a longstanding neuroscience question, a unified vision of the role of the motor cortex is still debated. Although advances in rodent research have shed light on this subject, the majority of rodent tasks lag behind those of primates, limiting the potential of circuit dissection with recording, manipulating and genetic tools available for mice. Thus, developing behavioral paradigms for mice that offer well-controlled, tractable and versatile skilled movements becomes necessary. We have recently described a reaching task in which head-fixed mice were trained to reach and grasp water droplets presented at different locations. Mice learned the task rapidly and efficiently performed hundreds of trials to three target locations (left, center and right), demonstrating a higher degree of dexterity and flexibility than generally assumed. Optogenetic inactivation of the motor cortex impaired reaching supporting its involvement in motor control. Layer 2/3 neurons displayed reach-related activity with a high degree of selectivity for trial-types (left, center, right). However, what these neurons actually code for (e.g. motor commands or target-related information) remains unanswered. To address this, we designed an extended version of the task and trained mice to initiate reaching for multiple targets from two alternative starting points. This allows discerning whether the neuronal activity is preferentially related to movement parameters or to endpoint target locations. Two-photon imaging of cortical neurons revealed a region-specific, wide palette of activity occurring at the time of the target presentation, movement initiation or reward consumption. Surprisingly, while only a minority of the neurons was modulated when changing the reaching starting point, the majority of trial-type selective neurons seemed unaffected by this manipulation. Taken together, by dissociating trial type from movement trajectory these findings suggest that neurons in the motor cortex surprisingly overrepresent spatial or goal related features, as compared to actual movement parameters.

II-55. Opponency in basal ganglia circuits underlies action suppression

Bruno Filipe Pereira Cruz[1] BRUNO.CRUZ@NEURO.FCHAMPALIMAUD.ORG
Sofia Soares[2] SOFIA_SOARES@HMS.HARVARD.EDU
Joe Paton[3] JOE.PATON@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Harvard University
[3]Champalimaud Research

Adaptive behavior involves suppressing in addition to producing actions and imbalanced action production and suppression are associated with disorders involving the basal ganglia (BG) such as ADHD, Parkinson's, and Huntington's diseases. From striatum, two distinct neuronal populations of projection neurons, dMSNs and iMSNs, divide BG circuitry into two parallel pathways. Circuit architecture and optogenetic activation suggest opposing roles of these two pathways during behavior: dMSN activation promotes, whereas iMSN activation suppresses movement. However, dMSN and iMSN activity appear largely correlated during movement, in apparent contradiction to the aforementioned opponent functionality. A conciliating theory[1] posits that the direct pathway energizes motor programs while the indirect pathway suppresses competing ones and that both processes are recruited during movement. If this were the case, then a state of active immobility should reveal opposite patterns of activity: dMSN and iMSN activity should be suppressed and sustained, respectively. To test this hypothesis we trained mice in a decision-making task where reward is contingent on immobility. Using fiber photometry we recorded activity from the two pathways in the dorsolateral striatum (DLS). While both dMSNs and iMSNs displayed phasic activation during movements, during active immobility iMSN activity was sustained, while dMSN activity fell to or below baseline. Furthermore, bilateral optogenetic inhibition of iMSNs resulted in an almost complete inability to suppress movements. Unilateral inactivation disrupted animals' ability to suppress contralateral actions, strikingly, during those trial epochs when that movement would be potentially rewarded and thus tempting. These experiments reveal for the first time multiple forms of opponency in the endogenous activity of the direct and indirect pathways during behavior. In addition, we demonstrate the necessity of indirect pathway activity for proper suppression of specific actions, findings with far-reaching implications for understanding BG contributions to behavioral control in health and disease.

II-56. Prediction of intra-operative speech arrest by means of functional magnetic resonance imaging and graph theory

Gino Del Ferraro[1] GINO.DELFERRARO@GMAIL.COM
Hernan Makse[1] HMAKSE@CCNY.CUNY.EDU
Luca Pasquini[2] LUCAPASQUINI3@GMAIL.COM
Kyung Peck[3] PECKK@MSKCC.ORG
Andrei Holodny[3] HOLODNYA@MSKCC.ORG

[1]The City College of New York
[2]La Sapienza University
[3]Memorial Sloan Kettering Cancer Center

Task-based functional magnetic resonance imaging (fMRI) is largely employed as a pre-surgical method to highlight those brain areas involved in a given cognitive performance and guide surgical intervention. Yet, this imaging technique is incapable of differentiating among primary and non-primary cognitive areas.

Intraoperative direct cortical stimulation (DCS) is used to find out this differentiation, in order to prevent post-surgical deficits which could considerably reduce life quality. However, DCS is limited to the cortical brain areas exposed to the surgical intervention and wide mappings are often hard to accomplish for safety reasons.

In this work we combine task-based fMRI signal acquired during a language task together with network theory to predict the primary language areas and help guiding surgical intervention of tumor resection. We construct the functional language network using refined statistical inference techniques and we further apply graph theory to numerically study the effects of targeted perturbations on this architecture. We finally validate our predictions in the operating room by means of DCS.

We analyze the fMRI signal of twenty patients with brain tumor invading the inferior frontal gyrus who have been electro-cortically stimulated in this region and, additionally, in the middle frontal gyrus. The DCS produced speech arrest in half of the patients. From our theoretical analysis we find that the patients who experienced speech arrest due to DCS have the predicted primary areas in the inferior or middle frontal gyrus. On the contrary, those patients who did not exhibit any speech deficit with DCS have the predicted primary areas outside the gyrus that have been intra-operatively stimulated.

These results may have a big impact on surgical planning, providing valuable information about the location of the area at risk for speech arrest in order to increase the sensitivity of the intra-operative mapping.

II-57. Identifying behavioral and neural correlates of dynamic decision formation during navigation

Shih-Yi Tseng SHIHYI_TSENG@G.HARVARD.EDU
Selmaan Chettih SELMAAN.CHETTIH@GMAIL.COM
Christopher Harvey HARVEY@HMS.HARVARD.EDU
Harvard University

Decisions reflect an interaction of external stimuli with internal cognitive processes, and are therefore difficult to control experimentally. Thus even with tightly controlled external conditions, the temporal dynamics of decision formation may vary from trial-to-trial, including moments of hesitation or changes in mind. This variability poses a challenge for identifying the neural substrates of decision-making. We reasoned that in a navigation-based task, as opposed to tasks restricting an animal's actions, ongoing movements might reflect the internal dynamics of decision formation. We established a flexible decision-making task for mice navigating a virtual maze, where rules determining cue-action-reward contingencies changed unexpectedly within sessions. This encouraged animals to continuously adapt their decisions, and introduced transient periods of uncertainty following a rule switch. We used a local history of cues and choices to estimate strategies, uncovering periods of rule-following, exploratory behavior, and choice bias. After rule switches, animals transiently made choices independent of cue, before adapting their behavior to fit the new rule. We asked whether within-trial dynamics of decision formation estimated from movement related to strategies. We used recurrent neural networks (RNNs) to derive time-varying, probabilistic estimates of a mouse's final choice from its movement trajectories. Interestingly, we found strong relationships between strategy and choice dynamics. The RNN decoder identified decisions early in a trial at times of rule-following, but during exploratory periods decisions were unclear until mice were forced to report their choice. Using two-photon calcium imaging, we also estimated choice dynamics by decoding activity in posterior cortex. Neural- and movement-derived choice signals exhibited similar dynamics and modulation by strategy. Our methods for behavioral readout thus permit inference of internal processes, both gradually evolving strategies and single-trial dynamics of decision formation, providing a foundation for further study of the underlying neural mechanisms.

II-58. Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites

Jakob Jordan[1] JORDAN@PYL.UNIBE.CH
Joao Sacramento[2] SACRAMENTO@INI.ETHZ.CH
Mihai A Petrovici[1,3] MPEDRO@KIP.UNI-HEIDELBERG.DE
Walter Senn[1] SENN@PYL.UNIBE.CH
[1]University of Bern
[2]ETH Zurich
[3]University of Heidelberg

Animals collect information about their environment through a variety of senses that need to be integrated into a coherent perspective. Since all sensory information is both incomplete and corrupted by noise, this integration has the main goal of increasing the information otherwise obtained from only a single sense. The Bayes-optimal estimate for the most likely stimulus under the assumption of independent Gaussian noise is obtained by av-

eraging estimates from different modalities while weighting each with its respective reliability. It was previously demonstrated in behavioral experiments that animals and humans combine multisensory stimuli in this optimal manner (Ernst and Banks, 2002; Fetsch et al., 2009; Nikbakht et al., 2018). What type of neuronal circuitry is able to perform such sensory integration? We present a neuron model capable of implementing the required computations by exploiting the biophysical dynamics of conductance-based neurons with dendritic compartments. Furthermore, a plausible error-driven plasticity rule enables neurons to learn not only input-output mappings, but to also simultaneously represent the respective reliabilities of each input that are necessary for a Bayes-optimal integration. In addition, the model supports dynamic reweighting of modalities and can thereby react to changes in stimulus reliabilities on a much shorter time scale than the one of synaptic plasticity. While both neuron and synapse dynamics are derived from a probabilistic description of neuronal processing, the model does not require a Bayes-optimal teacher but only input-output samples, allowing efficient learning. To illustrate our model, we present a feed-forward circuit receiving input from two different modalities with different associated reliabilities and show that after learning, the circuit optimally takes into account the respective reliabilities when processing new information. Finally, we discuss extensions of our model to non-linear dendritic compartments and to multi-layered cortical circuits that learn continuous input-output mappings (Dold et al., 2018).

II-59. Reconciling the cognitive and neural architecture of multisensory processing in the autistic brain

Michael Crosse[1] MICHAEL.CROSSE@EINSTEIN.YU.EDU
Cristiano Cuppini[2] CRISTIANO.CUPPINI@UNIBO.IT
John J Foxe[3] JOHN_FOXE@URMC.ROCHESTER.EDU
Sophie Molholm[1] SOPHIE.MOLHOLM@EINSTEIN.YU.EDU

[1]Albert Einstein College of Medicine
[2]University of Bologna
[3]University of Rochester

Individuals with autism spectrum disorder (ASD) are often impaired in their ability to process multisensory information, which may contribute to some of the social and communicative deficits that are so prevalent in this population. Despite the increasing empirical evidence, our understanding of the cognitive and neural architecture that underpins the development of this ability is lacking. Recent work has shown that response times (RTs) to multisensory stimuli can be modeled using an extension of the classic race model framework (Otto and Mamassian, 2012). Here, we demonstrate that the basic race model architecture is predictive of multisensory benefits in adult participants with ASD (R2(23) = 0.25, p = 0.017), but not in children with ASD (R2(34) = 0.012, p = 0.531). We find that by modelling an alternative processing strategy whereby multisensory RTs are determined by the preceding modality is a better predictor of behavior in autistic children than the race model (R2(34) = 0.14, p = 0.029). To understand the neural basis of this cognitive framework, we developed a neuro-computational model presenting two levels of multisensory interactions: inhibition and cooperation. This model is based on a previous neural network implementation used to explain acquisition of multisensory integrative abilities at the neuronal level in the superior colliculus (Cuppini et al., 2011; Cuppini et al., 2018). The model suggests that in the absence of substantial experience with multisensory stimuli (i.e., at an early stage of development), the main interaction between sensory modalities is competition, with the preceding modality the stronger competitor. At a later stage of development, experience with crossmodal events appears to promote positive interactions between modalities, thus enhancing behavior. These findings link our cognitive framework to a plausible neural implementation and provide an explanation for the multisensory deficits commonly reported in children with ASD.

II-60. Motor context dynamically tunes walking performance in exploratory Drosophila

Tomas Cruz TOMAS.CRUZ@NEURO.FCHAMPALIMAUD.ORG
Eugenia Chiappe EUGENIA.CHIAPPE@NEURO.FCHAMPALIMAUD.ORG
Terufumi Fujiwara TERUFUMI.FUJIWARA@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Foundation

Exploration is thought to rely on neural circuits that combine an animal's locomotive movements with ensuing sensory signals to extract information about the state of the body and the world. Here we show that in walking Drosophila an internal representation of body movements, which is based on self-generated visual signals (visual feedback) and congruent non-visual information, contributes to the performance control of walking in a motor-context specific manner. Using virtual reality technology in freely walking flies, we show that the forward movement of the fly, characterized by high translational speed, defines a specific motor context upon which visual feedback controls the straightness of the fly course. Specific perturbations of the virtual world, together with simulations reveal that a multimodal control system consisting of multiple channels integrating visual motion cues with rotation-related, non-visual signals can robustly explain the observed straightness performance. Furthermore, simultaneous analysis of head-fixed behavior and neural activity reveals a concerted multimodal activity of distinct populations of visuomotor neurons well posed to process such visual feedback. This activity accurately represents deviations of the fly from a straight course, and can steer the fly's walking trajectory. However, its effective contribution to straight course control changes dynamically as a function of the forward speed of the fly. Our data suggest a general circuit mechanism for a dynamic contribution of visual feedback onto locomotion performance in which self-generated visual signals associate with non-visual locomotive related signals in a motor-context specific manner to transform a faithful representation of ongoing movement into corrective locomotor commands.

II-61. Co-refinement of network interactions and neural response properties in visual cortex

Sigrid Tragenap[1] TRAEGENAP@FIAS.UNI-FRANKFURT.DE
Bettina Hein[1] HEIN@FIAS.UNI-FRANKFURT.DE
David E Whitney[2] DAVID.WHITNEY@MPFI.ORG
Gordon B Smith[3] GBSMITH@UMN.EDU
David Fitzpatrick[2] DAVID.FITZPATRICK@MPFI.ORG
Matthias Kaschube[1] KASCHUBE@FIAS.UNI-FRANKFURT.DE

[1]Frankfurt Institute for Advanced Studies
[2]Max Planck Florida Institute for Neuroscience
[3]University of Minnesota

In the mature visual cortex, local tuning properties are linked through distributed network interactions with a remarkable degree of specificity (Smith et al., 2018 Nat Neurosci). However, it remains unknown whether the tight linkage between functional tuning and network structure is an intrinsic feature of cortical circuits, or instead gradually emerges in development. Combining virally-mediated expression of GCAMP6s in pyramidal neurons with wide-field epifluorescence imaging in ferret visual cortex, we longitudinally monitored the spontaneous activity correlation structure - our proxy for intrinsic network interactions – and the emergence of orientation tuning around eye-opening. We find that prior to eye-opening, the layout of emerging iso-orientation domains is only weakly similar to the spontaneous correlation structure. Nonetheless within one week of visual experience, the layout of iso-orientation domains and the spontaneous correlation structure become rapidly matched. Motivated by these observations, we developed dynamical equations to describe how tuning and network correlations co-refine to become matched with age. Here we propose an objective function capturing the degree of consistency between orientation tuning and network correlations. Then by gradient descent of this objective function, we derive dynamical equations that predict an interdependent refinement of orientation tuning and network correlations. To

first approximation, these equations predict that correlated neurons become more similar in orientation tuning over time, while network correlations follow a relaxation process increasing the degree of self-consistency in their link to tuning properties. Empirically, we indeed observe a refinement with age in both orientation tuning and spontaneous correlations. Furthermore, we find that this framework can utilize early measurements of orientation tuning and correlation structure to predict aspects of the future refinement in orientation tuning and spontaneous correlations. We conclude that visual response properties and network interactions show a considerable degree of coordinated and interdependent refinement towards a self-consistent configuration in the developing visual cortex.

II-62. Decoding value from human amygdala neurons in observational learning

Tomas Aquino[1] TAQUINO@CALTECH.EDU
Juri Minxha[1] JMINXHA@CALTECH.EDU
Ian Ross[2] IANROSSMD@GMAIL.COM
Adam Mamelak[3] ADAM.MAMELAK@CSHS.ORG
Ueli Rutishauser[1] URUT@CALTECH.EDU
John O'Doherty[1] JDOHERTY@CALTECH.EDU

[1]California Institute of Technology
[2]Huntington Memorial Hospital
[3]Cedars-Sinai Medical Center

Humans can learn how to represent and predict rewards from their own experience, but also from observing others in a social context. Previous studies have demonstrated the importance of the amygdala in creating predictive associations between stimuli and outcomes, but the extent to which these factors are represented at a single neuron level in observational learning is less understood. Can we find reward predictive signals and reward prediction error signals encoded at the single-unit level or at the population in the amygdala? Are these values encoded in the same way during learning through observation? To answer these questions, we designed a reward learning paradigm with self-experienced and observational learning, and administered it to epilepsy patients implanted with depth electrodes. We modeled behavior with a simple reinforcement learning model and compared the neural representation of self-experienced and observed expected values and reward prediction errors. We performed amygdala population decoding and found representations of expected value in observational trials, outcome value in self-experienced trials, and unsigned reward prediction error in both trial types. This provides evidence to implicate human amygdala in encoding of predictive value signals during observational learning, further suggesting that value encoding happens in a distributed manner in populations of amygdala neurons. We compared self-experienced and observational learning by training and testing decoders across trial types and found no evidence for a shared code for the value regressors, suggesting a distinct code for self-experienced and observational learning.

II-63. A neural circuit for processing social motion cues in the juvenile zebrafish brain

Johannes Kappel[1,2] JKAPPEL@NEURO.MPG.DE
Johannes Larsch[1] JLARSCH@NEURO.MPG.DE
Joseph Donovan[1] JOE@NEURO.MPG.DE
Herwig Baier[1] HBAIER@NEURO.MPG.DE

[1]Max Planck Institute of Neurobiology
[2]Baier

Social interactions rely on the ability of each animal to recognize conspecifics based on sensory cues such as visual appearance. However, the perceptual mechanisms by which the vertebrate brain processes visual conspecific cues to regulate social behavior are largely unknown. In a recent study, we found that juvenile zebrafish are specifically attracted to the characteristic motion dynamics underlying discrete swim bouts of same-age conspecifics (Larsch & Baier, 2018). Here, we used this well-defined visual cue to identify neural correlates of social recognition across the juvenile zebrafish brain. By performing functional two-photon calcium imaging in the head-immobilized animal, we recorded neuronal responses in the forebrain and midbrain to a moving black dot. To observe which neurons preferentially respond to physiologically relevant stimulus conditions, we systematically changed the motion from continuous to bout-like dynamics. By fitting a flexible model to capture response characteristics of each neuron to the different stimulus conditions, we extracted neuronal clusters whose tuning matched behavioral preference in the virtual reality assay. We discovered that neurons in specific areas of the tectum and forebrain encode bout-like motion frequencies, with spatial clusters exhibiting varying degrees of tuning. In summary, we functionally identified candidate brain areas and neurons that might be involved in the perception and processing of visual conspecific cues. Future studies including targeted ablations and holographic optogenetic stimulation combined with behavioral readout will reveal which neurons are crucial for social recognition and social affiliation.

II-64. Geometric interpretation of robustness in spike coding networks

Nuno Calaim[1] NUNO.CALAIM@NEURO.FCHAMPALIMAUD.ORG
Florian Alexander Dehmelt[2] FLORIAN.DEHMELT@UNI-TUEBINGEN.DE
Pedro J Goncalves[3] PPJGONCALVES@GMAIL.COM
Christian Machens[4] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Research
[2]Tuebingen University Hospital
[3]Research Center Caesar
[4]Champalimaud Foundation

Neural circuits are remarkably robust against a variety of perturbations, including their partial destruction. Whether such robustness relies on some unifying principles of neural circuit design or whether each system has found its own solution, remains a largely unexplored question in systems neuroscience. Spiking networks based on (predictive) spike coding have been shown to be remarkably robust against neuron loss or spike failures, while also capturing key aspects of neural circuit dynamics, such as irregular and unreliable firing, as well as balance of excitation and inhibition. Here, we develop a simple geometric framework that immediately illustrates the principles underlying this robustness against a variety of perturbations. Specifically, the dynamics of a spike coding network is confined to a polyhedron in an abstract error space, which we call the 'error bounding box', and perturbations (such as noise, spike failures, synaptic failures, parameter mistuning, or neuron loss) can all be represented as deformations of this bounding box. Using these insights, we demonstrate the precise regimes for which the systems maintain or lose robustness to these perturbations. Furthermore, new predictions and explanations for previously puzzling effects can be obtained with simple geometric reasoning. Namely, we show that the partial (optogenetic) inhibition or excitation of a neural circuit are fundamentally different perturbations of the dynamics, and do not simply lead to opposing effects. Specifically, spike coding networks will compensate partial inhibitory perturbations, whereas they can generally not compensate partial excitatory perturbations. Moreover, for inhibitory perturbations, a subset of the inhibited neurons might in fact end up being more excited. Our work provides an intuitive understanding of spike coding networks, and suggests that the robustness of neural circuits may be based on a single design principle.

II-65. A model of active sensing via grid cells for visual recognition memory

Andrej Bicanski A.BICANSKI@UCL.AC.UK
Neil Burgess N.BURGESS@UCL.AC.UK
University College London

Models of face, object, and scene recognition traditionally focus on massively parallel processing of low- level features, with higher order representations emerging at later processing stages [1,2,3]. However, visual perception relies on eye movements, which are necessarily sequential. Neurons in entorhinal cortex with grid cell-like firing have recently been reported in response to eye movements, i.e. in visual space [4,5,6]. A functional explanation for these ‘visual grid cells’ is so far lacking. Grid cells (GCs) [7] are predominantly known from spatial memory and navigation studies and exhibit regularly arranged firing fields in the navigation plan. Ensembles of GCs are thought to underlie path integration and the calculation of goal- directed movement vectors [8,9,10]. Following the presumed role of GCs in vector navigation, we propose a model of recognition memory for familiar faces, objects, and scenes, in which GCs encode translation vectors between salient stimulus features. A sequence of saccadic eye movement vectors, moving from one salient feature to the expected location of the next, potentially confirms an initial hypothesis (accumulating evidence towards a threshold) about stimulus identity. This identification is based on the relative feature layout (going beyond recognition of individual features), and implements an relational active sensing strategy to infer the stimulus identity of exemplars within a stimulus category. Category identification is hypothesised to occur in earlier occipito-temporal areas, likely via parallel processing. The model constitutes the first quantitative proposal for a role of GCs in visual recognition. The variance of GC activity along saccade trajectories exhibits 6-fold symmetry across 360 degrees akin to recently reported fMRI data [5,6]. The mechanism is robust with regard to partial visual occlusion, can accommodate size and position invariance, and suggests a functional explanation for medial temporal lobe involvement in visual memory for relational information and memory guided attention [14].

II-66. Neurotransmitter spillover redresses the information loss caused by synaptic depression

Mehrdad Salmasi[1] MEHRDAD.SALMASI@LRZ.UNI-MUENCHEN.DE
Stefan Glasauer[2] S.GLASAUER@LRZ.UNI-MUENCHEN.DE
Martin Stemmler[1] STEMMLER@BIO.LMU.DE
[1]Ludwig-Maximilians-Universitat Munchen
[2]Brandenburg University of Technology

Through diffusion, neurotransmitters can reach more receptors than just the ones that lie directly opposite the presynaptic release site. In the presence of neurotransmitter spillover, the evoked postsynaptic potential is no longer a linear function of the number of released synaptic vesicles. In particular, spillover amplifies weak synaptic events. It is, however, as of yet unclear to what degree spillover modulates synaptic information transmission. We here present a tractable model of synaptic transmission to quantify the number of additional bits of information provided by spillover. We model synaptic release sites as communication channels that incorporate both synchronous spike-evoked and asynchronous release. These channels are subject to short-term depression, reflecting the recovery process from previous synaptic releases, and interact with each other through spillover. We first show that asynchronous release, rather than being a nuisance, can enhance neuronal communication through stochastic resonance. Our analysis reveals a synergy between neurotransmitter spillover and asynchronous release and determines the optimal level of asynchronous release that leads to the maximum rate of information transfer. Short-term depression reduces the responsiveness of the synapse and typically degrades information transmission across the neurons. We show that neurotransmitter spillover can switch the functional role of synaptic depression—in neurons with spillover, the rate of information transmission increases through short-term depression. These findings have important implications for dendritic computation and provide evidence for the critical role of neurotransmitter spillover in modulating information flow in neuronal networks.

II-67. Neural dynamics during changes of mind

Alex Piet ALEXPIET@GMAIL.COM
Ahmed El Hady AHADY@PRINCETON.EDU
Tyler Boyd-Meredith JBOYDMEREDITH@GMAIL.COM
Carlos Brody BRODY@PRINCETON.EDU

Princeton University

How are choices made within constantly-changing noisy environments? The gradual accumulation of noisy evidence is considered to be a fundamental component of decision making. In a static environment, the optimal strategy involves equally weighting all evidence over time effectively integrating with a long timescale. In a changing environment, the optimal evidence accumulation strategy involves the additional task of discounting old evidence that may no longer inform the current state of the world, effectively integrating with a short timescale. Previous work has characterized the evidence accumulation process in both static (Brunton, 2013) and dynamic (Piet, 2018) environments, finding rats adopt the optimal timescale in both cases. Discounting old evidence, via leaky integration, in a dynamic environment causes the subject to change their minds during each trial. We investigated the neural dynamics underlying changes of mind using a dynamic auditory decision-making task. We recorded neural activity from rats using microwire arrays from the Frontal Orienting Fields (FOF), to assess if and how neurons encode changes of mind. We developed a behavioral model to predict, on a moment-by-moment basis, the internal accumulated evidence variable used by each rat. This model makes quantitative predictions about the timing of changes of mind. Average firing rates conditioned on the changes of mind reveal single units in the FOF encode changes of mind. In order to develop a more complete mapping between neural activity and the accumulated evidence we applied and extended a cognitive tuning curve method to our data (Hanks, 2015). We found neural activity in the FOF can be described by a stereotyped temporal trajectory modulated by the strength of the accumulated evidence. Our results suggest FOF neurons encode the upcoming choice of the animal with a stable code, despite changes of mind which alter the tentative choice of the rat.

II-68. Emergence of stable sensory and dynamic time representations in the hippocampus

Jiannis Taxidis[1] JTAXIDIS@G.UCLA.EDU
Eftychios Pnevmatikakis[2] EPNEVMATIKAKIS@GMAIL.COM
Apoorva Mylavarapu[1] AMYLAVARAPU@UCLA.EDU
Peyman Golshani[1] PGOLSHANI@MEDNET.UCLA.EDU

[1]University of California, Los Angeles
[2]Simons Foundation

Hippocampal networks form maps of experience through spiking sequences that can encode external sensory cues as well as internal representations of time. Yet, the principles underlying the emergence and stability of such diverse representations remain largely unknown. Do both sensory and temporal codes exhibit the same reliability, plasticity or long-timescale stability? Do they both pre-exist in the network or do they emerge as a context is learned? To address these questions, we trained mice to perform an olfactory delayed non-match-to-sample task (DNMS), requiring working-memory activation during the delay period. Using in vivo two-photon calcium imaging, we recorded activity from thousands of dorsal CA1 neurons over multiple consecutive days, while mice learned and performed the task. We observed anatomically intermixed spiking sequences, comprised of 'odor-cells' encoding olfactory cues, followed by 'time-cells' encoding odor-specific delay timepoints. Odor-cells were reliably activated across trials and retained strikingly stable fields over multiple days and extended delays. In contrast, time-cells exhibited sparse, unreliable activation and labile fields that remapped across days or delays. During learning of the task, the number of odor-cells remained stable, whereas time-cells increased over days, as performance improved. This increase was not observed during passive exposure to the task in untrained animals. Therefore, multi-modal representational regimes, with different learning-related dynamics and stability, can co-

exist within the same CA1 network and are likely driven by distinct neurophysiological and plasticity mechanisms.

II-69. A face familiarity detection system with complex synapses

Ji-An Li[1] JIAN.LI.ACAD@GMAIL.COM
Fabio Stefanini[2] FS2545@COLUMBIA.EDU
Marcus K Benna[2] MKB2162@COLUMBIA.EDU
Stefano Fusi[2] SF2237@COLUMBIA.EDU

[1]University of Science and Technology of China
[2]Columbia University

Neural systems with plastic synapses that have low precision have a limited memory capacity. We recently showed that it is possible to overcome these limitations by making the synapses more complex, imitating the intricate networks of interactions of biochemical processes that characterize the biological synapses. However, these types of synapses have been tested only on simple memory retrieval problems involving random uncorrelated patterns. Here we consider a real-world problem, face familiarity detection, which is a particularly challenging task that the biological brain can solve. The task is difficult because faces are presented only once, memorized in one shot and have to be remembered for a long time. We built a multi-modular neural system that can fully take advantage of complex biological synapses. The input (embedding) module is a deep convolutional network pre-trained on a face recognition task. This module projects to the memory module, the only one containing complex plastic synapses which are continuously updated using a simple Hebbian rule. Finally, a readout (detection) module compares the output of the memory module to the face representations of the input module to assess the level of familiarity. We analyzed the memory lifetime of this system in the case of structured face patterns using a standard machine learning dataset of faces in real-life scenarios. We showed that the performance increases by orders of magnitude when the complexity of the synapses increases. The use of complex synapses allows the system to learn quickly (in one shot) and retain the information about face familiarity for a long time.

II-70. Dynamic causal modelling (dcm) of the attentional modulation of V1-V4 coherence

Christini Katsanevaki[1] TINAKATSANEV@GMAIL.COM
Andre Moraes Bastos[2] ANDREMBASTOS@GMAIL.COM
Hayriye Cagnan[3] HCAGNAN@GMAIL.COM
Conrado Arturo Bosman[4] C.A.BOSMANVITTINI@UVA.NL
Karl Friston[5] K.FRISTON@UCL.AC.UK
Pascal Fries[6] PASCAL.FRIES@ESI-FRANKFURT.DE

[1]Ernst Strungmann Institute (ESI) for Neuroscience in Cooperation with Max Planck Society
[2]The Picower Institute for Learning and Memory
[3]University of Oxford
[4]Swammerdam Institute for Life Sciences
[5]Wellcome Trust Centre for Neuroimaging
[6]ESI Ernst Struengmann Institute for Neuroscience, Donders Institute Nijmegen

The Communication-through-Coherence hypothesis proposes that flexible effective communication among neurons depends on their coherence [1]. Indeed, when two visual stimuli induce two local gamma rhythms in macaque V1, only the gamma induced by the attended stimulus entrains gamma in V4 and establishes coherence [2]. Additionally, attended V1 gamma shows higher peak frequency, potentially causing the selective V1-V4 coherence. Here, we investigate whether and which local modulations within V1 can result in the observed attentional effects. We used DCM for Cross-Spectral Densities [3] to fit a canonical microcircuit [4] to macaque electrocorticographic

data [2]. DCM allows investigating the causes of the attentional effects at the microcircuit level, which is challenging to access experimentally. To fit the complex, non-linear models to the dataset and overcome local minima, we used Variational Laplace [5], introducing multiple re-initializations with stochastic updates of parameters. This new scheme allowed us to 1) identify models with higher evidence than achieved with the conventional method, 2) show that the nature of the attentional effects in the winning model is robust across neighboring local minima and 3) show that there is an optimal level of noise that both overcomes local minima and avoids poor fitting. The accuracy of the obtained model fits suggests that modulation of activity within V1 is sufficient to reproduce the observed attentional effects. We used Parametric Empirical Bayes [6] to identify attentional modulations conserved over the selected V1-V4 pairs. The results point to inhibitory pathways within V1 that are likely to be regulated for attention. We find 1) an increase in the inhibition from a common pool of interneurons to both superficial and deep excitatory neurons and 2) a decrease in the local inhibition (potentially PV+ activity) to superficial excitatory neurons, allowing for a higher gain (or precision) of inputs from the attended population to V4.

II-71. A novel computational model capturing dynamic changes in the perisaccadic response of visual neurons

Kaiser Niknam KAISER.NIKNAM@UTAH.EDU
Amir Akbarian AMIR.AKBARIAN@UTAH.EDU
Behrad Noudoost BEHRAD.NOUDOOST@UTAH.EDU
Neda Nategh NEDA.NATEGH@UTAH.EDU

University of Utah

The spatiotemporal sensitivity of visual neurons changes dynamically and dramatically around the time of rapid eye movements, called saccades. For example, many studies have shown that neurons in the visual cortex lose their sensitivity to stimuli appearing in their receptive fields (RFs) shortly before a saccade (saccadic suppression); or, even before a saccade is initiated, visual neurons' RFs may preemptively shift to the post-saccadic RF location (future field remapping, or FF-remapping), or to the saccade target (saccade target remapping, or ST-remapping). Taken together, these findings indicate that the perisaccadic responses evoked in visual neurons are modulated by stimuli presented perisaccadically at multiple locations in the visual field, reflecting multiple modulation sources contributing to perisaccadic response generation. The fast timescale of perisaccadic response dynamics creates challenges both for experimental approaches and for computational models aiming to quantitatively characterize the relationship between various extrinsic or intrinsic system covariates and the changes in sensitivity to visual input across space and time. No existing computational approach is able to reproduce changes in sensitivity on the millisecond-scale at which they occur across eye movements. In order to capture these changes, we developed a novel extension of the generalized linear models (GLMs), which are widely used to describe neural responses in early sensory areas, along with an efficient model estimation procedure. The fitted model, with biophysically plausible components, is able to (1) reproduce the neuron's single-trial responses on a millisecond timescale and capture single-trial perisaccadic response changes and neural events, and (2) decompose responses and trace the dynamics of neuron's sensitivity to different parts of space (RF, FF and ST) across saccades. This millisecond-scale, quantitative decomposition of a neuron's perisaccadic responses offers unique opportunities for quantifying perisaccadic modulations and testing the independent contribution of various perisaccadic modulatory sources to perceptual effects induced by saccadic eye movements.

II-72. Between-subject prediction reveals a shared representational geometry in the rodent hippocampus

Hung-Tu Chen HTCHEN.GR@DARTMOUTH.EDU
Jeremy Manning JEREMY.R.MANNING@DARTMOUTH.EDU
Matthijs van der Meer MVDM@DARTMOUTH.EDU

Dartmouth College

A fundamental challenge faced by any memory system is how related experiences should be organized: storing the details of individual experiences preserves potentially valuable details, but is storage-inefficient and hampers generalization, whereas treating all experiences as the same risks ignoring potentially important differences.

The rodent hippocampus is a model system for studying the neural basis of these computations. It can construct statistically independent representations across environments ("global remapping"; Alme et al. 2014) and assigns individual neuron firing fields to locations in an apparently random fashion (Rich et al. 2014). Similarly, "engram" studies suggest that the population of neurons allocated to a given experience is determined by a competition based on slowly fluctuating excitability levels among eligible neurons (Josselyn & Frankland, 2018).

This relatively random mapping between hippocampal neurons and their coding properties implies that it should be challenging to predict hippocampal encoding of a novel experience in a given subject based on the neural encoding of experience in another subject. Contrary to this prediction, we find that by constructing a common representational space across rats ("hyperalignment", commonly applied to fMRI data; Haxby et al. 2011), we can consistently predict withheld data ("right" trials on a T-maze) in a target rat based on 1) the "left" trials of the target rat, and 2) the relationship between left and right trials from a different source rat. This cross-subject prediction outperformed a number of control mappings, such as those based on permuted data that broke the relationship between "left" and "right" activity for individual neurons, and those based on within-subject prediction alone.

These results suggest a representational geometry in the rodent hippocampus that is shared across subjects, using a novel approach with the potential to reveal which aspects of neural ensemble activity for related experiences are shared vs. unique to an individual.

II-73. Gain modulation arises naturally from spike coding networks representing a background variable

Sander Keemink[1,2] SWKEEMINK@SCIMAIL.EU
Nuno Calaim[3] NUNO.CALAIM@NEURO.FCHAMPALIMAUD.ORG
Christian Machens[1] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Champalimaud Centre for the Unknown
[3]Champalimaud Research

Brains are able to represent real-world information with incredible precision. Although the principles underlying such coding are not fully understood, a recurring mechanism across the brain is response normalization (Carandini and Heeger, 2012). In response normalization a neuron's response is normalized by some multiplicative factor, usually dependent on the population firing rate. Such normalization allows for adjusting the sensitivity to changes in stimuli, and is also thought to be involved in attention-based response changes (Reynolds and Heeger, 2009). Neural networks based on (predictive) spike coding have been shown to be remarkably efficient in their representation of stimuli (Boerlin et al., 2013; Deneve and Machens, 2016), but they do not currently have a natural way to implement such response normalization. We propose that this can be done, without needing to adjust the underlying theory, by making each neuron represent an additional background variable with equal weight. We show that the network's firing rate can then be controlled directly by the background variable. As the background variable increases, both the population firing rates and the decoding precision increase, mimicking attention-based response normalization. As stimulus strength increases, population firing rates increase but the decoding precision decreases, mimicking feed-forward based response normalization. We have thus effectively implemented response normalization, with a plain-vanilla network of integrate-and-fire neurons, simply by representing an additional background variable. Moreover, we find that the adjusted spike-coding networks are more stable (and no longer prone to epilepsy-like seizures), that they balance excitation with inhibition (but no longer inhibition by excitation), and that all recurrent connections are inhibitory (as opposed to mixing inhibition and excitation from the same neuron, a violation of Dale's law). The adjusted networks present an intriguing hypothesis for brain function, and allow for new predictions on the effects of stimulus strength, attention, and response normalization.

II-74. Efficient coding of numbers explains biased judgments

Arthur Prat-Carrabin ARTHUR.P@COLUMBIA.EDU
Brian Ho BLH2136@COLUMBIA.EDU
Michael Woodford MW2230@COLUMBIA.EDU

Columbia University

Decisions often require the aggregation of multiple sources of information, a simple example of which is the computation of an average. Human subjects, when averaging numerical or visual stimuli in comparison tasks, have been reported to unequally weight different numbers, that should be equally relevant to a correct decision. This selective weighting has been interpreted as resulting from a nonlinear transformation of the presented numbers, that optimally compensates for the presence of internal noise arising later in the decision process. A natural alternative hypothesis is that numbers are encoded with noise at the time of their presentation. Assuming this coding is efficient, the amount of noise should vary across the stimulus space, and the coding should depend on the prior distribution from which numbers are sampled. Moreover, if ultimate judgments make optimal use of the information in the coded values, the decoding rule should also depend on the prior. We investigate these predictions through a task in which participants are asked to compare the averages of two series of numbers, each sampled from a prior distribution which differs across blocks of trials. We show that, in addition to a nonlinear transformation, subjects encode the numbers with a degree of noise that varies with the magnitude of the number; in particular, less likely values (under the prior) are encoded with greater noise. Our results are also consistent with optimal decoding given the prior-dependent coding rule. Together, our results shed both experimental and theoretical light on an encoding-decoding paradigm, in which efficient coding is combined with optimal inference from noisy representations.

II-75. Entorhinal cells as generalisable basis functions for transition statistics

James Whittington[1] JAMES.WHITTINGTON@MAGD.OX.AC.UK
Shirley Mark[2] S.MARK@UCL.AC.UK
Timothy Muller[1] TIMOTHYMULLER127@GMAIL.COM
Tim EJ Behrens[1] BEHRENS@FMRIB.OX.AC.UK

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

Since the discovery of grid cells, numerous other entorhinal cells with distinct spatial representations have emerged. Here we propose that many of these cells (grid cells, border cells, band cells, object vector cells) can be unified under a common framework, with each cell serving as a basis function describing a particular feature of learned transition structures. With particular weightings of these bases (cells), one can describe transition structures of arbitrary novel environments. We implement our proposal in an artificial neural network tasked with predicting sensory observations while traversing 2D graphs. The subsequent predictive representations that emerge mirror spatial representations found in the brain. With random walks through the graph, grid and band cells emerge.

Spending more time near boundaries leads to the emergence of border cells. Biasing towards particular sensory experiences leads to the emergence of object vector cells. These entorhinal representations (bases) serve as an abstract code that generalises across environments. We simultaneously observe environment/object specific codes in the form of place and landmark cells that map the current environment in a conjunctive representation, which models the hippocampus. Our results suggest that the 'zoo' of different cell types found in entorhinal cortex may be viewed under a unified framework - summarising the common statistics of tasks into basis functions that can be flexibly combined depending on the particular structural constraints of the environment that the animal/agent faces.

II-76. Ancestral computations constrain the evolution of new neural computations

Stephanie Palmer SEPALMER@UCHICAGO.EDU
Nathan Buerkle NBUERKLE@UCHICAGO.EDU

University of Chicago

Evolving new behaviors or perceptual abilities requires modifying the complex organization of neural circuits. This ancestral circuitry might impose significant constraints on which behaviors can feasibly evolve as well as the specific computations implemented. To address questions about the evolution of neural computation, we used machine learning to model the repeated evolution of red-sensitive, tetrachromatic color vision in butterflies from a UV/blue/green trichromatic ancestor. We first trained networks with trichromatic inputs to discriminate color stimuli. To simulate color vision evolution, these trichromatic networks were then retrained with 'mutated', tetrachromatic inputs. We compared evolved networks to tetrachromatic networks trained de novo from random starting weights to test two hypotheses: 1) trichromatic networks have a latent capacity to evolve tetrachromacy, and 2) trichromatic starting points constrain how tetrachromatic networks implement novel computations. As expected, tetrachromatic networks performed well over a larger range of stimuli than trichromatic networks. The initial mutation, which results in a mismatched tetrachromatic input and trichromatic hidden layer severely impaired performance, revealing a potentially strong barrier to evolution. Subsequent training improved performance, but networks trained de novo from random weights performed slightly better than evolved networks. We then asked how starting weights affected the computational structure of the hidden layer. PCA reduced the dimensionality of the hidden layer to 3 hidden units that were directly comparable between networks. Using hierarchical clustering, we showed that networks with identical trichromatic starting weights typically converged on one or two computational patterns. In contrast, networks with different starting weights or identical random starting weights used a diverse set of computations. Overall, our results suggest that trichromatic circuits impede tetrachromacy evolution, but when evolution does occur, the underlying computational mechanism is funneled into specific evolutionary trajectories.

II-77. Separable codes for behavior in mouse primary visual cortex across attentional states

Ashley Wilson WILSON@NEURO.DUKE.EDU
Jeffrey Beck JEFF.BECK@DUKE.EDU
Lindsey Glickfeld GLICKFELD@NEURO.DUKE.EDU

Duke University

Fluctuations in neuronal activity in sensory cortices are tightly linked to perception and can be used to identify the contribution of specific neuronal populations to behavior. However, since flexible behavior likely requires engaging distinct sensory circuits, this relationship between neuronal activity and behavior may depend on context. Thus, we developed a multi-modal attention task for head-fixed mice in order to better understand how distinct sensory

neuron populations are engaged by moment-to-moment changes in behavioral context. In this task, mice are cued on a trial-by-trial basis to detect visual or auditory targets, but receive the same visual input across trial types as they anticipate targets. To understand how sensory cortex is modulated by attention we recorded from populations of neurons in primary visual cortex (V1) as mice performed the task. We then used a population of tuned, task-responsive neurons to predict either the presence of a target or, in a separate analysis, the mouse's behavior on visual trials or auditory trials. As expected, small populations of V1 neurons could accurately predict both visual targets and behavior. Surprisingly, when this analysis was performed on auditory trials, V1 neurons could also predict auditory targets and decisions above chance. This suggests two possibilities – that populations of V1 represent decisions in general when the mouse is engaged in a task or that there is mixed selectivity in the population across different contexts. We find that the weights in the visual and auditory models were highly uncorrelated, and could not be used to predict targets or behavior across modalities, indicating that representation of behavior in V1 neurons is unique across attention states. Thus, to investigate the neuronal specificity of this attentional modulation, we are currently analyzing the relationship of a neuron's weight, attention selectivity, and tuning during the task.

II-78. Dimensionality expansion via chaotic dynamics facilitates classification in a trained network

Matthew Farrell[1] MSF9@UW.EDU
Stefano Recanatesi[1] STEFANO.RECANATESI@GMAIL.COM
Timothy Moore[1] MTIMOTHY984@GMAIL.COM
Guillaume Lajoie[2] LAJOIE@DMS.UMONTREAL.CA
Eric Shea-Brown[1] ETSB@UW.EDU

[1]University of Washington
[2]Universite de Montreal

Chaos plays an important role both in interpreting experiments and in models of cortical dynamics. More recently, chaos has been seen to play an important role in training recurrent neural networks (RNN), allowing networks to learn to do tasks using dynamics that resemble those used by the brain. Here we introduce a new role for chaos in the context of training an RNN to classify inputs using working memory driven by recurrent dynamics.

We consider a standard RNN model trained to classify static inputs after a delay period, and find that the fully trained network performs the task with either stable attractor dynamics or chaotic dynamics, depending on initialization of the weights before training. We compare these two dynamical strategies and find that, surprisingly, the network is able to more easily classify certain inputs with chaotic attractors than it is with stable attractors. We explain this phenomenon through the lens of dimensionality, and propose chaotic dynamics as one mechanism by which an RNN can expand the dimensionality of its representation, facilitating classification of inputs with highly nonlinear class separation boundaries. Our findings suggest a new functional role for the chaos observed in cortical networks and provide a new addition to dynamical systems perspectives on choosing optimal parameters when training neural networks.

II-79. Dissociating navigation from spatial representation

Omri Barak OMRI.BARAK@GMAIL.COM
Tie Xu FEXUTIE@GMAIL.COM

Technion

Navigation in partially observable environments is often assumed to involve an estimation of current position. This view is strengthened by tuning properties of place and grid cells. Mechanistically, position can be estimated by attractor dynamics implementing path integration with corrections from landmarks. The goal of navigation, however, is to reach a desired location – not to estimate position. The various uncertainties and constraints of navigation may render a metric representation suboptimal. To explore a possible dissociation between navigation and spatial representation, we use reinforcement learning to train recurrent neural networks (RNNs) on a navigation task inspired by the Morris water maze. Surprisingly, we find that an echo state RNN with random internal connectivity is able to perform well in multiple contexts. A weak spatial representation does develop. This representation emerges from the coupling of the RNN dynamics to the environment, giving rise to a composite dynamical system that is no longer random. In this case, landmarks do not reset (a non-existing) path integrator, but rather modulate a composite dynamical system. The internal connectivity can be pre-trained to have a better or worse spatial representation. This is done through supervised training of a path integrator or adversary learning, respectively. We find that a minimal representation is needed to support reinforcement learning. The representation itself arises from effective timescales of the network as reflected by its eigenvalues. While a network pre-trained to have a path integrator performs better, relying on exact spatial coordinates can be detrimental. We compare generalization of the networks to new environments with similar topology and different metrics. In this case the echo state RNN with random internal connectivity shows better performance than the path integrator. Our results show that an explicit representation of position might be neither optimal, nor necessary for navigation.

II-80. Criticality between cortical states

Pedro Carelli[1] PVCARELLI@GMAIL.COM
Antonio Fontenele[1] FONTENELENETO.A.J@GMAIL.COM
Nivaldo Vasconcelos[1] NAPVASCONCELOS@GMAIL.COM
Thais Feliciano[1] THAISFELICIANO.SILVA@GMAIL.COM
Leandro Aguiar[2] LEANDROAGUIAR@BIOLOGO.BIO.BR
Carina Soares-Cunha[3] CARINACUNHA@MED.UMINHO.PT
Barbara Coimbra[3] ID5309@ALUNOS.UMINHO.PT
Leonardo Dalla Porta[4] LEONARDODALLAPORTA@HOTMAIL.COM
Sidarta Ribeiro[5] SIDARTARIBEIRO@GMAIL.COM
Ana Joao Rodrigues[3] AJRODRIGUES@MED.UMINHO.PT
Nuno Sousa[3] NJCSOUSA@MED.UMINHO.PT
Mauro Copelli[1] MCOPELLI@DF.UFPE.BR

[1]Universidade Federal de Pernambuco
[2]Universidade Federal Rural de Pernambuco
[3]University of Minho
[4]IDIBAPS
[5]Universidade Federal do Rio Grande do Norte

Since the first measurements of neuronal avalanches, the critical brain hypothesis has gained traction. However, if the brain is critical, what is the phase transition? For several decades it has been known that the cerebral cortex operates in a diversity of regimes, ranging from highly synchronous states (e.g. slow wave sleep, with higher spiking variability) to desynchronized states (e.g. alert waking, with lower spiking variability). Here, using independent signatures of criticality, we show that a phase transition occurs in an intermediate value of spiking variability. The critical exponents point to a universality class different from mean-field directed percolation (MF-DP). Importantly, as the cortex hovers around this critical point, it follows a linear relation between the avalanche exponents that encompasses previous experimental results from different setups and is reproduced by a model.

II-81. The significance of nominally non-responsive activity in frontal and sensory cortex

Michele Insanally[1] MNI1@NYU.EDU
Ioana Carcea[1] IOANA.CARCEA@RUTGERS.EDU
Rachel Field[1] RACHEL.FIELD2@NYUMC.ORG
Chris Rodgers[2] XRODGERS@GMAIL.COM
Brian DePasquale[3] DEPASQUALE@PRINCETON.EDU
Kanaka Rajan[4] KANAKA.RAJAN@MSSM.EDU
Michael DeWeese[5] DEWEESE@BERKELEY.EDU
Badr Albanna[6] BALBANNA@FORDHAM.EDU
Robert Froemke[1] ROBERT.FROEMKE@MED.NYU.EDU

[1]New York University
[2]Columbia University
[3]Princeton University
[4]Icahn School of Medicine at Mount Sinai
[5]University of California, Berkeley
[6]Fordham University

Spike trains recorded from the cortex of behaving animals can be complex, highly variable from trial-to-trial, and therefore challenging to interpret. A fraction of cells exhibit trial-averaged responses with obvious task-related features such as pure tone frequency tuning in auditory cortex. However, a substantial number of cells (including cells in primary sensory cortex) do not appear to fire in a task-related manner1 and are often neglected from analysis. Even classically-responsive cells lose their stimulus representation during task-engagement without impairing behavior2,3. These results suggest that nominally non-responsive cells may play an underappreciated role in sensory processing and cognition. At Cosyne 2018, we presented a novel single-trial, spike-timing-based analysis to evaluate whether population activity recorded from auditory and frontal cortex encode task variables in behaving rats. Here we expand our investigation to explore the potential utility and necessity of nominally non-responsive cells and demonstrate: 1) Nominally non-responsive cells reveal hidden task information comparable to responsive cells. The activity of cells that seem unresponsive when trial-averaged often encode additional task-relevant information at levels comparable to responsive cells. 2) Stimulus information is more prevalent and pervasive in frontal cortical ensembles. When tones become behaviorally significant, stimulus information is encoded more accurately in frontal cortex suggesting it is critical for extracting task-relevant stimulus information. 3) Nominally non-responsive ensembles predict behavioral errors. Historically, the capacity to predict behavioral errors on single trials using responsive cells has been limited. Using our novel decoder, we demonstrate that including non-responsive cells significantly improved predictions of behavioral errors in both auditory and frontal cortex indicating that these neglected cells can provide missing insights into behavioral variability. 4) Nominally non-responsive activity is necessary for task performance and interacts synergistically with the activity of responsive cells. Using a recurrent neural network model trained to perform our frequency recognition task, we establish that nominally non-responsive activity is necessary for successful task performance. Furthermore, we demonstrate that responsive and non-responsive subpopulations have a synergistic effect on network behavior.

II-82. The successor representation in plastic networks of spiking neurons

Sara Zannone S.ZANNONE14@IMPERIAL.AC.UK
Jacopo Bono BONO.JACOPO@GMAIL.COM
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK

Imperial College London

It has been recently proposed that the hippocampus learns a predictive representation of the environment, known in reinforcement learning as the Successor Representation. The Successor Representation accounts for numer-

ous experimental findings, while also providing a useful representation for reward maximization. Nonetheless, the details and plausibility of a biological implementation remain unclear. Here, we show that the successor representation can be learned in a spiking neural network through a biologically plausible plasticity rule, with no need for a separate prediction error. Our plasticity rule results in updates equivalent to $TD(\lambda)$ and can learn over different timescales. We show that fast-spiking sequences, reminiscent of hippocampal replays, lead to $TD(\lambda \approx 1)$, while activations over longer behavioural timescales approximate $TD(\lambda \approx 0)$. Thanks to the bootstrapping property of $TD(0)$, our model can bridge the gap between neuronal and behavioural timescales of plasticity. Moreover, we show that our framework implicitly defines the notion of 'state' without the need to discretise time. It also allows state-dependent discounting, which is useful to encode information like terminal or salient states. Finally, because of the different properties of $TD(1)$ and $TD(0)$, we predict an additional functional role for replays, namely controlling the bias-variance trade-off when learning a cognitive map.

II-83. Multi-scale calcium imaging of functional maps in the primate primary visual cortex

Declan Rowley[1] DECLAN.ROWLEY@MONASH.EDU
Osamu Sadakane[2] OSAMU.SADAKANE@RIKEN.JP
Akiya Watakabe[2] AKIYA.WATAKABE@RIKEN.JP
Toshiki Tani[2] TOSHIKI.TANI@RIKEN.JP
Hiroshi Abe[2] HIROSHI.ABE@RIKEN.JP
Noritaka Ichinohe[2] NORITAKA.CHINOHE@RIKEN.JP
Hiroaki Mizukami[3] MIZ@JICHI.AC.JP
Elizabeth Zavitz[1] ELIZABETH.ZAVITZ@MONASH.EDU
Hsin-Hao Yu[1] HHYU00@GMAIL.COM
Marcello Rosa[1] MARCELLO.ROSA@MONASH.EDU
Tetsuo Yamamori[2] TETSUO.YAMAMORI@RIKEN.JP

[1]Monash University
[2]RIKEN CBS
[3]Jichi Medical University

Multiple overlapping functional maps have been reported in the primate visual cortex. However, it has been difficult to characterise these maps in the same animal over a wide field of view at high spatial resolutions. To study the precise cross-scale spatial relationships among functional maps, we conducted in vivo calcium imaging of neuronal activities in the primary visual cortex (V1) of the marmoset monkey (Callithrix jacchus). An AAV viral vector carrying the GCaMP6s calcium indicator (amplified with the Tet-Off system) was injected into V1 of three adult marmosets. We imaged the calcium activities at columnar resolution with wide-field (3mm diameter) one-photon imaging, and at single-neuron resolution with two-photon imaging in sufentanil-anaesthetised monkeys. A suite of visual stimuli was used to quantify orientation, direction, spatial frequency and colour selectivity. In addition, we mapped the retinotopy of the imaged regions with bars at randomised orientations and positions. We were able to visualise multiple pinwheel centres and orientation domains using wide-field imaging, and observed that these corresponded tightly to the preferred orientations of the individual neurons in each domain. We also observed domains with different preferred spatial frequencies that overlapped, but were organised orthogonally to the orientation domains. Furthermore, blob-like clusters were revealed by isoluminant colour stimuli. By mapping preferences to multiple visual features we hope to better understand how the complex multidimensional space of visual information can be projected smoothly onto the flat cortical surface and encoded by the cells within it.

II-84. Laminar dependence of the effect of task engagement in the auditory cortex

Raphael Steinfeld[1,2] RAPHAEL.STEINFELD@NEURO.FCHAMPALIMAUD.ORG
Andre Tacao-Monteiro[1] ANDRE.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG
Alfonso Renart[3] ALFONSO.RENART@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Champalimaud Research
[3]Champalimaud Centre for the Unknown

Although a primary sensory area, representations in auditory cortex have been shown to display strong context dependence. Here, we used an auditory delayed-response task to study this problem in head-fixed mice. We train animals to discriminate the frequencies of pure tones to receive water rewards. Importantly, after sound presentation mice have to withhold their response until presentation of a visual go-cue. This structure allows temporal decoupling of sensory and decision-making processes from the movement-related activity. Pharmacological silencing using muscimol confirmed the functional involvement of ACx in task performance. We recorded the activity of ACx during execution of the task as well as during passive listening to replays of the same sessions in acute recordings using multi-shank 64-channel silicone probes. Across our datasets, the most salient effect of task engagement is twofold: Stimulus information decays quicker after the initial sound response of the neurones during task engagement, but is re-enhanced at the end of the delay period, while it steadily decays during replay conditions. We demonstrate that these two effects are mediated by distinct, localised, modulations of stimulus information across cortical layers. Across-trial variability in activity was constantly lower during behaviour, an effect mediated by a reduction in the magnitude of positive noise correlations. Finally, while our data shows little difference in sound responses and frequency discriminability during sound presentation at a single cell level, population analysis using regularised logistic regression suggests a less redundant, population-level representation of frequency during task engagement. Our results clarify the changes induced by task engagement in the representation of sounds by cortical circuits.

II-85. Adjusting for batch effects in two photon imaging recordings of the retinal inner plexiform layer

David Klindt[1] KLINDT.DAVID@GMAIL.COM
Luke E Rogerson[2] LUKE.E.ROGERSON@GMAIL.COM
Zhijian Zhao[2] ZHIJIAN.ZHAO@UNI-TUEBINGEN.DE
Katrin Franke[2] KATRIN.FRANKE@UNI-TUEBINGEN.DE
Dmitry Kobak[2] DMITRY.KOBAK@UNI-TUEBINGEN.DE
Alexander Ecker[2] ALEXANDER.ECKER@UNI-TUEBINGEN.DE
Matthias Bethge[2] MATTHIAS.BETHGE@UNI-TUEBINGEN.DE
Philipp Berens[2] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Thomas Euler[2] THOMAS.EULER@UNI-TUEBINGEN.DE
Klaudia Szatko[2] KLAUDIA.SZATKO6@GMAIL.COM

[1]National Research University Higher School of Economics, University of Tuebingen
[2]University of Tubingen

The retina decomposes visual stimuli into parallel channels which encode different features of the visual environment, including light increments, and decrements, edge orientation and direction and colour. Central to this parallelisation is the excitatory feedforward network composed sequentially of photoreceptors, bipolar cells and ganglion cells. Two-photon imaging with an electrically tuneable lens allows us to scan vertical cross-sections of excised intact retina, and to observe glutamate release at bipolar cell terminals across the inner plexiform layer (IPL), where different bipolar cell types stratify at different depths. To integrate recordings across retinae, and jointly disentangle the feature channels from the observed data, it is necessary to remove any batch effects, i.e.

inter-experimental variation caused by temperature, jitter, bleaching, indicator concentration or laser intensity. We wanted to robustly identify parallel retinal feature channels under two light stimulus conditions: 1) a 1,000s oscillating light spot with variable frequency, contrast and spot size; 2) a sequence of 108 5s natural movie clips. For each, PCA of glutamate recordings from the IPL reveals that most variance is driven by differences in response polarity (On vs. Off), however a large proportion of the variance is due to batch effects. We systematically evaluated whether popular batch correction algorithms from the RNA-Seq literature were able to remove this effect, in addition to a batch-specific transform on the temporal kernels of a Convolutional Neural Network (CNN). We found that Matching Mutual Nearest Neighbours (MMNN) [1] outperformed classical linear methods (ComBat [2], Limma [3]), and could be used with the batch-specific convolution transform to improve performance, both increasing the predictive power and better distributing the recordings across clusters [4]. Batch effect problems have been largely unaddressed in the two-photon imaging literature, but are likely to be increasingly important as the size, complexity and ambition of imaging datasets grow.

II-86. A slow drift in macaque visual and prefrontal cortex: Implications for perceptual decision-making

Benjamin Cowley[1,2] BCOWLEY@PRINCETON.EDU
Adam Snyder[3] ADAM.SNYDER@ROCHESTER.EDU
Katerina Acar[4] KAC216@PITT.EDU
Ryan Williamson[5] RCW30@PITT.EDU
Byron Yu[5] BYRONYU@CMU.EDU
Matthew Smith[4] SMITHMA@PITT.EDU

[1]Princeton University
[2]Princeton Neuroscience Institute
[3]University of Rochester
[4]University of Pittsburgh
[5]Carnegie Mellon University

A subject's ability to detect a stimulus change can be influenced by cognitive strategies (e.g., motivation, effort, arousal, etc.) that vary over the course of a task. It is largely unknown how these strategies exert control over the neural signals that govern the decision-making process. To study this, we trained macaque monkeys to perform a perceptual decision-making task, and observed that subjects slowly varied their likelihood of indicating that a stimulus change occurred. One possible neural mechanism could be a slowly-varying cognitive factor that biases sensory evidence and then propagates to the decision. Under this model, we would expect to observe this factor in sensory evidence (e.g., activity in visual area V4) as well as in activity of downstream, decision-related areas (e.g., pre-frontal cortex, PFC). To test for this, we simultaneously recorded from macaque V4 and PFC (about 20 to 60 neurons in each area), and found that V4 and PFC population activity slowly drifted together. This slow drift also covaried with observed changes in behavior. Thus, these results were consistent with the proposed model. However, we further found, inconsistent with the model, that V4 activity without the slow drift better predicted within-trial behavior. This suggests a more complicated picture: The slow drift is broadcast throughout the brain but is removed from readouts of sensory evidence. Still, the slow drift can influence the final decision, likely after integration of sensory evidence but before execution of a motor command. Overall, this work uncovers a cognitive factor that can be observed across multiple brain areas and suggests a role this factor plays in the decision-making process.

II-87. Decoding time and neural dimensionality as tools for understanding delay activity dynamics

Christopher Cueva[1] CCUEVA@GMAIL.COM
Alex Saez[1] CJC2197@COLUMBIA.EDU
Encarni Marcos[2] ENCARNI.MARCOS@UNIROMA1.IT
Aldo Genovesio[2] ALDO.GENOVESIO@UNIROMA1.IT
Mehrdad Jazayeri[3] MJAZ@MIT.EDU
Ranulfo Romo[4] RROMO@CORREO.IFC.UNAM.MX
C Daniel Salzman[1] CDS2005@COLUMBIA.EDU
Michael Shadlen[1] MS4497@CUMC.COLUMBIA.EDU
Stefano Fusi[1] SF2237@COLUMBIA.EDU

[1]Columbia University
[2]University of Rome
[3]Massachusetts Institute of Technology
[4]Universidad Nacional Autonoma de Mexico

Our decisions often depend on multiple sensory experiences separated by time delays. The brain can remember these experiences (working memory) and, at the same time, it can easily estimate the timing between events, which plays a fundamental role in anticipating stimuli and planning future actions. To understand the neural mechanisms underlying working memory and time encoding we analyze neural activity recorded during delays in four different experiments on non-human primates, and we consider three classes of neural network models to explain the data: attractor neural networks, chaotic reservoir networks and recurrent neural networks trained with backpropagation through time. To disambiguate these models we propose two analyses: 1) decoding the passage of time from neural data, and 2) computing the cumulative dimensionality of the neural trajectory as it evolves over time. Our analyses reveal that time can be decoded with high precision in tasks where timing information is relevant and with lower precision in tasks where it is irrelevant to perform the task. The neural trajectories are low dimensional for all datasets. Consistent with this, we find that the linear "ramping" component of each neuron's firing rate strongly contributes to the slow timescale variations that make decoding time possible. We show that these low dimensional ramping trajectories are beneficial as they allow computations learned at one point in time to generalize across time. Our observations constrain the possible models that explain the data, ruling out simple attractor models and randomly connected recurrent networks (chaotic reservoir networks) that vary on relatively fast timescales, but agree with recurrent neural network models trained with backpropagation through time. Our results demonstrate a powerful new tool for studying the interplay of temporal processing and working memory by objective classification of electrophysiological activity.

II-88. Representation of uncertainty during hippocampal theta sequences

Gergo Orban[1] ORGERGO@GMAIL.COM
Balazs Ujfalussy[2] BALAZS.UJFALUSSY@GMAIL.COM
Marton Kis[3] KIS.MARTON@KOKI.MTA.HU
Lajos Vago[1] VAGOLALA@GMAIL.COM

[1]MTA Wigner RCP
[2]Institute of Experimental Medicine
[3]Institute for Experimental Medicine

Behavioural studies suggest that both humans and other animals are able to perform probabilistic computations. Such computations imply that the nervous system of biological agents is capable of the representation and manipulation of probability distributions. However, the way encoded distributions are related to population activity of neurons remains hotly debated because measures that could dissociate alternative models based on experimental data are remarkably lacking.

Here, we focus on hippocampal activity in the context of exploratory behavior to derive contrasting predictions for the alternative models. Place cells, selective for specific locations in the environment, become sequentially activated during each theta cycle, thus neurons encoding past, present, and future locations outline the trajectory of the animal. We interpret this activity pattern as the result of repeatedly performing probabilistic inference about possible trajectories in a dynamical generative model. Critically, during a single theta sequence the uncertainty is expected to change systematically, thus providing a chance to identify how it is encoded in the population activity. Specifically, we consider four alternative encoding models: (a) encoding the most likely trajectory; (b) sampling from the posterior distribution; (c) standard probabilistic population coding and (d) convolutional encoding.

We create synthetic datasets in which place cells encode trajectories according to the four alternative encoding strategies, all bearing the key features of place cell dynamics in the hippocampus. By decoding the population activity during theta sequences, we derive three measures that can distinguish between competing models and which can be assessed on neuronal activity patterns recorded during exploratory behavior. These results are directly applicable to experimental data to identify if and how uncertainty of spatial trajectories are represented in the hippocampus. Our analysis is an important step towards elucidating the strategies used by the brain to encode probability distributions and to understand the computational role of neuronal variability.

II-89. Optimal multisensory decision making with invariant probabilistic population codes

Han Hou[1] HOUHAN@GMAIL.COM
Yuchen Zhao[1] YCZHAO7@GMAIL.COM
Qihao Zheng[1] QHZHENG@ION.AC.CN
Alexandre Pouget[2] ALEX.POUGET@GMAIL.COM
Yong Gu[1] GUYONG@ION.AC.CN

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

Most of our decisions are based on multiple sources of evidence whose reliabilities vary over time. This problem is conceptually complex, because the optimal solution requires that sensory evidence be weighted in proportion to their respective time-varying reliabilities, which are typically not known in advance. Despite of this complexity, theoretical studies have shown that the optimal solution can be achieved by simply summing neural activity across modalities and time without the need for any form of reliability-dependent synaptic reweighting, as long as sensory inputs are represented with a type of code known as invariant linear probabilistic population codes (ilPPC). Here we address this problem in the context of visual-vestibular heading perception. We first asked whether the visual and vestibular momentary evidence in cortex are encoded with ilPPCs. Although the responses of visual and vestibular cortical neurons were found to deviate significantly from the assumptions of ilPPC, we found that the information loss due to this deviation was around 5%. In other words, the heading perception system is 95% ilPPC-compatible. Next, we recorded from single neurons in the lateral intraparietal area (LIP) while the animals were optimally performing a visual-vestibular heading discrimination task. We report that LIP population activities are consistent with the prediction of an ilPPC neural network which performs optimal evidence integration via random, but time-invariant, linear combination of visual and vestibular inputs across both modalities and time. Importantly, this conclusion did not hold for MSTd neurons, suggesting that MSTd provides the visual momentary evidence but lies upstream from the site of temporal and cross-modality integration. Taken together, our results provide the first neural evidence in support of a remarkably simple solution to a very general, and ecologically relevant, form of decision making involving multiple sources of evidence.

II-90. Using local plasticity rules to train recurrent neural networks

Owen Marschall OEM214@NYU.EDU
Kyunghyun Cho KYUNGHYUN.CHO@NYU.EDU
Cristina Savin CS5360@NYU.EDU

New York University

To learn useful dynamics on long time scales, neurons must use plasticity rules that account for long- term, circuit-wide effects of synaptic changes. In other words, neural circuits must solve a credit assignment problem to appropriately assign responsibility for global network behavior to individual circuit components. Furthermore, biological constraints demand that plasticity rules are spatially and temporally local; that is, synaptic changes can depend only on variables accessible to the pre- and postsynaptic neurons. While artificial intelligence offers a computational solution for credit assignment, namely backpropagation through time (BPTT), this solution is wildly biologically implausible. It requires both nonlocal computations and unlimited memory capacity, as any synaptic change is a complicated function of the entire history of network activity. Similar nonlocality issues plague other approaches such as FORCE (Sussillo et al. 2009). Overall, we are still missing a model for learning in recurrent circuits that both works computationally and uses only local updates. Leveraging recent advances in machine learning on approximating gradients for BPTT, we derive biologically plausible plasticity rules that enable recurrent networks to accurately learn long-term dependencies in sequential data. The solution takes the form of neurons with segregated voltage compartments, with several synaptic sub-populations that have different functional properties. The network operates in distinct phases during which each synaptic sub-population is updated by its own local plasticity rule. Our results provide new insights into the potential roles of segregated dendritic compartments, branch-specific inhibition, and global circuit phases in learning.

II-91. Auditory representation in cortex during perceptual learning

Kathleen Martin KAM953@NYU.EDU
Robert Froemke ROBERT.FROEMKE@NYULANGONE.ORG

New York University

Auditory learning reliably enhances cortical representations of task-relevant stimuli and improves behavioral discrimination. However, these neural changes have predominantly been observed outside of the behavioral context. Our lab recently showed found that auditory cortical neurons display different degrees of facilitation and suppression to behaviorally-relevant tones when comparing responses inside and outside of the behavioral context. However, how sensory representations change, specifically in a behavioral context, over auditory perceptual learning and how this impacts behavioral performance remains under-studied. In this study, we assessed neural dynamics over learning during behavior in the same animal. We developed a head-fixed behavioral paradigm in which mice learn to lick a left lick port for a chosen tone and right for tones of all other frequencies for water rewards. This experimental set up allows us to probe both auditory acuity and categorical responses over perceptual learning. Animals improve their discrimination between center and surround frequencies over the course of weeks, but continue to make errors on frequencies close to the center frequency. We performed two photon calcium imaging of excitatory neurons in auditory cortex throughout learning, to assess neural activity both during this behavior and passive listening. In the behavioral context, many excitatory neurons exhibit a categorical response to the auditory stimuli, not encoding the frequency of the stimulus, but rather the behavioral meaning. This categorical response is present early in behavioral learning, but is broader, mirroring the behavioral performance. During passive tone presentation, the percentage of responsive neurons to behaviorally-relevant tones increases, corresponding to the behavioral categorical boundaries. We are currently imaging excitatory and inhibitory neurons simultaneously and cholinergic axons during behavior to investigate how cholinergic-mediated changes in inhibition may shape this categorical response.

II-92. Functional segregation of the ferret auditory cortex probed with natural and model-matched sounds

Agnes Landemard[1] AGNES.LANDEMARD@ENS.FR
Sam Norman-Haignere[1] SN2776@COLUMBIA.EDU
Shihab Shamma[2] SAS@ISR.UMD.EDU
Yves Boubenec[1] YVES.BOUBENEC@ENS.FR
Celian Bimbard[1] CELIAN.BIMBARD@ENS.FR

[1]Ecole Normale Superieure
[2]University of Maryland

Sensory systems are adapted to represent natural stimuli despite their complexity. How auditory cortex encodes this richness of acoustic features into spatially organized patterns of activity remains poorly understood. Here, we combined a novel computational approach contrasting the brain responses to synthetic sounds matching either part or all of natural acoustic features (panel a) with a cutting-edge high-resolution neuroimaging technique, functional UltraSound. Using this unique combination, we set out to explore functional cortical domains at the basis of natural sound processing. We first mapped the classical tonotopy of ferret auditory cortex (panel b), highlighting core and belt regions. We then used Independent Component Analysis on sound-evoked neural responses to reveal in an unsupervised manner independent spatial sources with specific functional properties. We illustrate a high-frequency component localized in secondary auditory cortex (panel c) and consistent throughout the cortical volume (see inset in panel b). This component displayed enhanced projections for stimuli endowed with natural sound-like spectral modulations compared to cochlear-matched stimuli (panel c). In order to pinpoint the respective contribution of each modulation type, we contrasted responses between different model-matched stimuli in individual voxels (panel d). We show enhanced responses to full model-matched sounds compared to cochlear-matched sounds in different subregions of core and belt areas (panel e1-f1). Interestingly, primary regions showed enhanced responses to temporally modulated stimuli, whereas distinct subregions of the belt areas had more diverse selectivity (temporal in green; spectral in red; both in yellow regions for panels e2-f2), consistent with functional segregation. Inclusion of an additional modulation feature increased responses in the belt (green and red zones in panels e3-f3), suggesting mixed selectivity. We found that low-frequency regions were completely explained by the full model, while high-frequency regions in primary and secondary regions had larger responses to original than full model-matched stimuli (panel f4), as expected from cortical areas sensitive to correlations between modulation filters.

II-93. Evolving selectivity in macaque visual cortex

Peter Schade PETER_SCHADE@HMS.HARVARD.EDU
Margaret Livingstone MLIVINGSTONE@HMS.HARVARD.EDU

Harvard University

Neuronal stimulus selectivity arises as a consequence of inputs from other neurons: near and far. These short and long-range interactions create complex tuning and dynamic response profiles. From the earliest stages of the macaque visual system, neurons have complex spatiotemporal response patterns, with the stimulus selectivity of transient responses differing from later sustained selectivity, reflecting the slight delay of lateral and feedback contributions. However, an understanding of temporal patterns of activity in later stages of the visual hierarchy, such as inferotemporal cortex (IT), has been elusive. To investigate the temporal dynamics of responses in macaque visual cortex, we used temporally constrained response windows to guide a genetic algorithm that searched an image space encoded in a deep generative network for images that could drive neuronal responses with distinct temporal patterns. The temporal constraints were predetermined and under investigator control. We searched in parallel for images that could cause individual neurons and population of neurons to generate responses that were selectively transient, transient and sustained, or only delayed. Images evolved via the genetic algorithm that differentially activated transient and sustained responses in V1, V4, posterior IT, and central IT.

Thus, we could control the temporal response pattern of populations of neurons across the ventral visual stream. Evolutions revealed known properties of V1, such as surround suppression, which is thought to arise from local lateral connections and inter-areal recurrent feedback. Similar temporally distinct response profiles were evoked in inferotemporal cortex. Since the visual system is a processing hierarchy, and V1 and IT are only stages within the hierarchy, the kinds of computations that take place in these areas are likely not qualitatively different. Given this, we may have a glimpse of how recurrent activity in higher order cortex arising from lateral and inter-areal feedback contributes to object representations.

II-94. Categorical perception: probing top-down signaling

Bin Min[1] MINBMATH@GMAIL.COM
Daniel Bliss[2] DPB6@NYU.EDU
Yang Zhou[3] YANGZHOUION@OUTLOOK.COM
David Freedman[3] DFREEDMAN@UCHICAGO.EDU
Xiao-Jing Wang[2] XJWANG@NYU.EDU

[1]Shanghai Researcher Center for Brain Science and Brain-Inspired Intelligence
[2]New York University
[3]University of Chicago

A fruitful theme of cognitive science is the interplay between analog feature-based perception and discrete categorization. There is early evidence of interactions between the two, called the “categorical perception” effect, in which category learning warps perception such that differences between objects that belong to different categories are exaggerated (expansion) while differences within the same category are deemphasized (compression). This suggests a top-down influence from category-selective to feature-selective representations, but the underlying neural mechanisms have not been established. To gain insight into this question, we examined data from behavioral categorization experiments in non-human primates. In the experiments, monkeys performed the same visual motion discrimination task before and after visual motion categorization training. Data analysis shows that, after categorization training, stimuli within the same category were more difficult to discriminate than before categorization training, while the change for stimuli that belong to different categories was less pronounced, supporting compression without clear expansion. To explain this result, we built a neural circuit model that incorporates key existing experimental findings and makes new predictions, including: (1) learned categories are encoded in the spiking activities of neurons in the lateral intraparietal (LIP) area, (2) neurons in the middle temporal area (MT) show graded encoding of stimulus motion directions and (3) neurons in the medial superior temporal (MST) area integrate top-down category and bottom-up motion direction information. This model proposes that it is mainly through the feedback projections from LIP to MST that learned categories induce categorical perception. We find that this prediction is largely consistent with recent single neuron recordings in the MST and LIP areas. Collectively, we showed monkey experimental evidence for compression in visual motion and developed a biological neural circuit model that allows us to make experimentally testable predictions, thereby potentially elucidating the underlying neural mechanisms of categorical perception.

II-95. Learning efficient, task-dependent representations with synaptic plasticity

Colin Bredenberg CJB617@NYU.EDU
Eero Simoncelli EERO.SIMONCELLI@NYU.EDU
Cristina Savin CS5360@NYU.EDU

New York University

Neural circuits cannot perfectly represent the sensory world: their capacity is limited by the number of available neurons, metabolic and other biophysical resources, and response variability (noise). It has been hypothesized that the brain responds to these limitations by constructing efficient representations, discarding some aspects of incoming sensory streams while preserving or enhancing more probable or behaviorally-relevant information. How the brain learns such representations remains an open question. Here we construct a recurrent neural circuit model that can learn efficient, task-specific sensory codes using simple forms of synaptic plasticity. We leverage a tractable stochastic recurrent network model to derive a broad class of reward-modulated Hebbian plasticity rules for optimizing task-specific cost functions. The stochasticity of the dynamics is key for this result. Although naively one would expect internal noise to be strictly detrimental for high-fidelity coding, here it allows the network to sample activity states that improve performance. Plasticity then increases the likelihood of those stimulus-dependent responses occurring in the future. We illustrate the flexibility of our approach by training an initially unstructured recurrent neural circuit to solve two distinct tasks, either faithfully representing its inputs, or classifying them into two categories. The emerging neural representation reflects task structure with inhomogeneously distributed tuning functions, which preferentially encode probable stimuli and the decision boundary for classi?fication. Overall our results reveal how simple synaptic plasticity allows noisy circuits to exploit environmental statistics to maximize their performance within tight biological constraints.

II-96. Mixed pattern generator for forward locomotion in C. elegans

Erick Olivares[1] ERICKOLIVARESB@GMAIL.COM
Eduardo Izquierdo[1] EDIZQUIE@INDIANA.EDU
Randall D Beer[2] RDBEER@INDIANA.EDU

[1]Indiana University Bloomington
[2]Indiana University

C. elegans locomotes in an undulatory fashion, generating thrust by propagating dorsoventral bends along its body. Despite the relative simplicity of the worm, how locomotion is generated is not yet well understood. Using the available connectome, we integrated a neuroanatomically-grounded model of its nervous system with a biomechanical model of its body. We then used an evolutionary algorithm to determine the unknown physiological parameters of each neuron and connection so that the complete system reproduces the locomotive behavior of the worm. We analyzed the ensemble of solutions as a way to generate novel hypotheses about the neuromechanical basis for locomotion. Specifically, we first demonstrate the feasibility of a network CPG in the worm's ventral nerve cord in the neural model alone. Our model suggests specific roles for key motoneurons and connections involved in the generation of the oscillation. Building on the first result, we integrate the neural model to the biomechanical model of the body and environment. Using this model, we demonstrate that a chain of CPGs, connected through a set of chemical and gap junctions, can drive forward locomotion on agar in the absence of proprioception. In addition to matching the speed of the worm, the model worms match several other key features of locomotion and of experimental manipulations that the model worms were not trained to match. Analysis of the solutions reveals three different strategies for how to achieve the anterior-posterior coordination of the oscillators. Each of the strategies constitutes a testable hypothesis. Finally, we added stretch receptors to the brain-body-environment model, and we re-evolved them to move forward efficiently. Analysis of the results suggests a mixed pattern generator is feasible, along with specific mechanisms for how the stretch-receptor feedback is used to help coordinate the intrinsic oscillators.

II-97. Quantifying route planning in rodent spatial navigation

Michael Pereira[1] MICHAEL.PEREIRA@NEURO.FCHAMPALIMAUD.ORG
Christian Machens[2] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Thomas Akam[3] THOMAS.AKAM@NEURO.FCHAMPALIMAUD.ORG
Rui M Costa[1,4] RUI.COSTA@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Research
[2]Champalimaud Foundation
[3]University of Oxford
[4]Columbia University

A longstanding question in neuroscience is how animals and humans select actions in complex decision trees. Planning, the evaluation of action sequences by anticipating action outcomes, is thought to coexist in the brain with simpler decision-making strategies, such as habit learning and heuristics. Though planning is often required for optimal choice, in many situations simpler strategies may result in similar action selection, making them difficult to disambiguate in behavioral data. As an ethologically ubiquitous behavior requiring sequential decision making, spatial navigation provides a compelling context for studying planning in rodents. However, no behavioral tasks currently exist which can dissociate route planning from vector navigation and habitual route selection, while generating large decision datasets well suited to neurophysiology. This lack of suitable behavioral paradigms has been a limiting factor for studying planning mechanisms in the rodent brain. To quantitatively isolate the contribution of planning to spatial navigation, we developed a novel navigation task in which mice navigate to visually cued goal locations in a large elevated maze with complex topology. We show how a principled search through the large space of possible maze configurations can produce layouts optimized to simultaneously isolate planning from vector navigation and habit learning. Mice perform hundreds of trials per session, each comprising several decisions informative of planning use. Trajectories shorten with training, approaching shortest paths to goals. Analysis of choices at decision points shows that after learning, mice preferentially choose actions favored by planning over those favored by vector navigation. Modeling the data with a mixture-of-strategies model including vector navigation, habit learning and planning reveals a significant planning component, which increases with training. These results quantitatively demonstrate the importance of route planning in rodent navigation and open the door to neurophysiological investigation of planning mechanisms.

II-98. Characterization of hippocampal neural activity in evidence accumulation and decision-making

Edward Nieh ENIEH36@GMAIL.COM
Carlos Brody BRODY@PRINCETON.EDU
Manuel Schottdorf MSCHOTTDORF@PRINCETON.EDU
Nicolas W Freeman NWF2@PRINCETON.EDU
Jeff Gauthier JEFF.L.GAUTHIER@GMAIL.COM
Mark L Ioffe MIOFFE@PRINCETON.EDU
Sue Ann Koay KOAY@PRINCETON.EDU
Lucas Pinto LPINTO@PRINCETON.EDU
David W Tank DWTANK@PRINCETON.EDU

Princeton University

In an ever-changing environment, the ability to accumulate and evaluate evidence is crucial for optimal decision-making. The hippocampus is thought to embody a cognitive map and has been well-studied in navigation and foraging tasks. However, it is unknown how the hippocampus behaves when evidence for decision-making must be accumulated over time. We performed 2-photon recordings in head-fixed mice performing a virtual reality T-maze task, where noisy evidence indicating a left or right turn is accumulated over time and imaged the activity of CA1 neurons (n=3557 neurons, n=5 animals). We show that while cells form choice-specific place sequences,

the consistency of firing in each cell's preferred place was surprisingly unreliable, suggesting that other cognitive dimensions may be affecting the activity of these neurons. When we analyzed the Skaggs spatial information between cellular activity and evidence, we found that CA1 neurons formed a sequence through evidence space as well. In fact, spatial information of neural activity in a joint 2-dimensional space comprised of position and evidence together was significantly greater than the same metric measured in either dimension alone. These results indicate that CA1 neurons can encode cognitive space in multiple dimensions in cognitively-demanding tasks. In addition, we discovered ordered sequences of neurons that were sequentially active, but variable in place or time, and found that these sequences appeared significantly more often than chance and were also more predictive of left or right turns than independent single cells. Premature ending of these right- or left-choice preferring sequences also significantly indicated a choice in the opposite direction. These results indicate two important findings: (1) CA1 neurons can encode multidimensional cognitive space and (2) form ordered yet time-variable sequences that are predictive of animal behavior.

II-99. Rats optimally manage learning during a decision-making task

Javier Masis[1] JMASIS@FAS.HARVARD.EDU
David Cox[1] DAVIDCOX@FAS.HARVARD.EDU
Andrew Saxe[2] ANDREW.SAXE@PSY.OX.AC.UK

[1]Harvard University
[2]University of Oxford

Optimally managing speed and accuracy during decision-making is crucial for survival in the animal kingdom and the subject of intense research. However, it is still unclear how an agent learns to manage this trade-off efficiently. Here, we show that rats learn to approach optimal behavior by simultaneously optimizing both instantaneous reward rate, and on a longer timescale, learning speed in a visual object recognition 2-AFC task. According to a theory for learning making use of deep linear neural networks, we show that this strategy leads to a higher reward rate faster, and a higher total reward than just maximizing instantaneous reward rate. We behaviorally test and confirm predictions from this theory: when required to learn a new stimulus-pair, well-trained rats slow down their reaction times during learning and these return to baseline upon asymptotic performance. Importantly, there is a strong correlation between how much each animal slows down and how fast it learns. We causally link the slow-down in reaction time with learning speed by showing that animals forced to respond above their average reaction time while learning a new stimulus-pair learn faster than those forced to respond below their average reaction time. Additionally, rats speed up their reaction times when placed in a setting where there are no prospects for learning. To our knowledge, ours is the first examination in this context in rats and our theory is the first to directly incorporate the learning process into free response binary choice models. Our results suggest that rats exhibit cognitive control of the learning process itself, and quantitatively demonstrate that their strategy can be a more favorable strategy during learning for decision-making agents in general.

II-100. Perceiving the duration of a vibration: psychophysics, leaky integrator model, optogenetic testing

*Sebastian Reinartz[1] REINARTZ.SEBASTIAN@GMAIL.COM
*Alessandro Toso[1] ALTOSO@SISSA.IT
*Arash Fassihi[2] ARFASSIHI@GMAIL.COM
Luciano Paz[1] LPAZ@SISSA.IT
Francesca Pulecchi[1] FRANCESCA.PULECCHI@GMAIL.COM
Mathew E Diamond[1] DIAMOND@SISSA.IT

[1]International School for Advanced Studies (SISSA)
[2]University of California, San Diego

The perceptual confound between stimulus duration and intensity, ‘a stronger stimulus feels longer’, has been described for several sensory modalities in animal and human behavioral studies. We recently found that the confound is complementary: ‘longer feels stronger’. Here we exploited this confound to study the computations and circuit dynamics underlying the perception of the duration and intensity of vibro-tactile stimuli.

In every trial, subjects received two “noisy” vibrations separated by a fixed delay; rats received stimuli on their whiskers and humans on the index fingertip. Subjects compared either the intensities (mean speeds) or the durations of the two vibrations. Rat/human results were similar. First, we observed psychometric curve shifts that quantify the reciprocal biases: stronger feels longer and longer feels stronger. Because this symmetry of biases suggested partially overlapping brain computations underlying the two percepts, we constructed a model whereby intensity and duration perception are mediated through leaky integration of tactile input, with task-specific levels of leak. The model replicated the observed intensity and duration psychophysics, suggesting that the networks involved in both perceptual tasks share broad architectural features (leaky integration of sensory drive) but differ in the setting of their time constants.

The model predicts that, in rats, perceived vibration duration, and the stimulus speed-dependent bias inherent in the perceived passage of time, are mediated through barrel cortex firing rate. We injected AAV5-CaMKIIa-hChR2(H134R)-EYFP or AAV5-CaMKIIa-eNpHR3.0-EYFP into barrel cortex and positioned an optical fiber accompanied by an array of moveable tungsten electrodes within the injection sites. Optogenetic excitation in barrel cortex led to overestimation of the duration of the simultaneous vibration; conversely, optogenetic suppression led to underestimation of the duration of the simultaneous vibration. In sum, psychophysics, modelling, and optogenetic testing combine to support a framework where time perception involves the accumulation of activity directly from sensory cortex.

II-101. Motion detection model predicts direction-reversing neurons as observed in the mouse visual cortex

Yazan Billeh YAZANB@ALLENINSTITUTE.ORG
Ramakrishnan Iyer RAMI@ALLENINSTITUTE.ORG
Gabriel Ocker GABEO@ALLENINSTITUTE.ORG
Michael A Buice MICHAELBU@ALLENINSTITUTE.ORG
Stefan Mihalas STEFANM@ALLENINSTITUTE.ORG
Anton Arkhipov ANTONA@ALLENINSTITUTE.ORG
Saskia de Vries SASKIAD@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

Studies of motion detection have been pursued for almost a century in multiple model organisms and it has long been established that many neurons in the visual system exhibit strong preference for direction of motion. Using calcium imaging data from the Allen Brain Observatory [De Vries et. al, 2018], we observe that approximately 7% of direction selective neurons in the mouse visual cortex reverse their preferred direction when exposed to drifting gratings with different temporal parameters. We observe these Direction Reversing Neurons (DRNs) in all visual areas recorded, with increased prevalence in higher visual areas, namely PM and AM, and with no difference in prevalence among different layer-specific Cre-defined populations. We also observed DRNs in electrophysiological recordings made in mouse V1 and LGN (data from Durand et. al).

Using spatiotemporal asymmetric filter models inspired by the receptive field properties of simple cells, we were able to recapitulate our experimental observations in response to drifting gratings. This mechanistic model also predicts that neurons reverse direction at different spatial frequencies, a phenomenon we observed in preliminary additional calcium imaging data, increasing our DRN estimate to 18% of direction selective cells. We emphasize that this observation of approximately a fifth of direction selective cells being DRNs, with strict criteria for selections, uses only a limited stimulus set.

This work naturally poses the question of whether DRNs are advantageous in visual processing or are an epiphenomenon. Basic theoretical analysis of binary neurons demonstrates that DRNs increase the encoding capacity of a network population compared to the same sized network without DRNs. Alternatively, DRNs could be an epiphenomenon elicited by cyclical motion stimuli that could explain illusory motion reversals such as the "wagon wheel" illusion.

II-102. Inferring interpretable nonlinear stochastic dynamics from population spike trains

Lea Duncker[1] DUNCKER@GATSBY.UCL.AC.UK
Julien Boussard[2] JULIENB@STANFORD.EDU
Gergo Bohner[3] GBOHNER@GATSBY.UCL.AC.UK
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

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

Computations in the brain are thought to be implemented by the dynamical evolution of distributed activity in recurrent circuits. A key goal is to characterise from population recordings how this computational activity evolves, and how the dynamical properties of the circuit give rise to this evolution and thus the computation. However, even today's technologies measure only a fraction of the neurons within a relevant population. This fact has driven the development of dimensionality reduction techniques, which leverage the available sample to estimate the trajectories of underlying computational state variables under the assumption that, in a controlled behaviour, these explore fewer degrees of freedom than the number of neurons recorded. But such methods do not directly characterise the circuit dynamics. Furthermore, by lacking dynamical structure they may fail to properly distinguish signal from noise in trajectories. Here, we develop a novel method to simultaneously infer the system's computational state and learn the underlying nonlinear dynamical properties of the circuit from the spike times of a sample of simultaneously recorded neurons. We model the continuous-time dynamical evolution of the population state and its variability using a stochastic differential equation (SDE), in which the nonlinear evolution is characterised in terms of its fixed points and local Jacobians, interpolated by a Gaussian Process. Thus, interpretable features of the dynamics enter our model directly as parameters, and are readily accessible for further analysis. We develop a novel algorithm for approximate posterior inference of the latent SDE trajectories, and use these posterior estimates to characterise the SDE parameters. We show that our method can successfully recover nonlinear dynamics and the location and stability of their fixed points on simulated examples.

II-103. Functional organization of excitatory and inhibitory presynaptic networks in primary visual cortex

Luigi Federico Rossi FEDERICO.ROSSI@UCL.AC.UK
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

Orientation selectivity emerges in the primary visual cortex (V1) from the convergence of retinotopically aligned inputs; yet, how this selectivity is maintained and amplified during further processing remains unclear. Here we show how the orientation selectivity of L2/3 neurons is computed from excitatory and inhibitory presynaptic neurons distributed across layers. We targeted individual pyramidal L2/3 neurons in mouse V1, and identified their presynaptic inputs through single-neuron initiated monosynaptic tracing with an EnvA-dG-DsRed rabies virus. Mice expressed GCaMP6 in all excitatory neurons, so that we could perform two-photon microscopy to reconstruct the postsynaptic neuron and its local presynaptic network, and to measure their retinotopy and orientation

tuning during passive wakefulness. For each postsynaptic neuron, we identified 132±41 (s.e., n = 13 mice) presynaptic partners across L1-5. The majority of these presynaptic partners were excitatory, providing sparse recurrent excitation from L2/3, and dense feedforward excitation from L4. Inhibitory inputs, which did not express GCaMP6, predominantly clustered in L2/3. Presynaptic networks formed elongated distributions in cortical space: in retinotopic space, this elongation matched the orientation preference of the postsynaptic neuron. Within this arrangement, there were striking differences between excitatory and inhibitory neurons: inhibitory neurons were concentrated locally around the postsynaptic cell, while excitatory neurons spanned a wider area. In retinotopic space, the excitatory inputs favored locations coaxially aligned with the orientation preference of the postsynaptic neuron, while inhibitory inputs were more uniformly scattered. While excitatory presynaptics represented the full range of orientation selectivity, the total excitation from every layer was tuned to the orientation preference of the postsynaptic neuron. In addition, excitatory presynaptic neurons formed ensembles sharing high pairwise total, signal, and noise correlations compared to the surrounding population. These results demonstrate that the orientation preference of L2/3 neurons arises from the convergence of retinotopically elongated, co-tuned excitation balanced by local, dense, inhibition.

II-104. Deep Dendrite: Bayesian inference of synaptic inputs from dendritic calcium imaging

Jinyao Yan YANJ11@JANELIA.HHMI.ORG
Aaron Kerlin KERLINA@JANELIA.HHMI.ORG
Laurence Aitchison LAURENCE.AITCHISON@GMAIL.COM
Boaz Mohar MOHARB@JANELIA.HHMI.ORG
Karel Svoboda SVOBODAK@JANELIA.HHMI.ORG
Srinivas C Turaga STURAGA@GMAIL.COM

HHMI Janelia Research Campus

In vivo calcium imaging can be used to probe the functional organization of synaptic activity across the dendritic arbor. Synaptic input onto spines can produce calcium transients that are largely isolated to the spine head. However – in otherwise unperturbed cells – the back-propagating action potential (bAP) also contributes strongly to the change in fluorescence measured at individual spines. To address this problem, we propose to perform Bayesian inference in a statistical model to separate these sources and infer the probability of both pre- and post-synaptic spiking activity. Our model is a simplified nonlinear approximation of the biophysical processes by which synaptic input and the bAP contribute to the fluorescence measurements at different sites. We use the framework of variational autoencoders (VAE) – a recent advance in machine learning – training a deep neural network (DNN) as part of the VAE to perform approximate Bayesian posterior inference over spike trains from fluorescence traces. In simulations, our approach successfully recovers correlations between simulated spine and soma activity from fluorescence signals, whereas conventional methods fail. We also processed data from in vivo imaging of the basal dendrites of L2/3 neurons in mouse frontal cortex and compare our method to conventional methods. Our method is a crucial step towards measuring the transformation of synaptic input to somatic output in vivo.

II-105. The geometry of abstraction in hippocampus and pre-frontal cortex

Silvia Bernardi[1] SILVIA.BERNARDI@GMAIL.COM
Marcus K Benna[1] MKB2162@COLUMBIA.EDU
Mattia Rigotti[2] MR2666@COLUMBIA.EDU
Jerome Munuera[3] JEROME.MUNUERA@GMAIL.COM
Stefano Fusi[1] SF2237@COLUMBIA.EDU
C Daniel Salzman[1] CDS2005@COLUMBIA.EDU

[1]Columbia University
[2]IBM Research AI
[3]Columbia University, Centre National de la Recherche Scientifique (CNRS), Ecole Normale Superieure

Abstraction can be defined as a cognitive process that finds a common feature - an abstract variable, or concept - shared by a number of examples. Knowledge of an abstract variable enables generalization, which in turn allows one to apply inference to new examples based upon old ones. Neuronal ensembles could represent abstract variables by discarding all information about specific examples, but this allows for representation of only one variable. Here we show how to construct neural representations that encode multiple abstract variables simultaneously, and we characterize their geometry. Representations conforming to this geometry were observed in dorsolateral pre-frontal cortex, anterior cingulate cortex, and the hippocampus in monkeys performing a serial reversal-learning task. These neural representations allow for generalization, a signature of abstraction, and similar representations are observed in a simulated multi-layer neural network trained with back-propagation. These findings provide a novel framework for characterizing how different brain areas represent abstract variables, which is critical for flexible conceptual generalization and deductive reasoning.

II-106. The hippocampus facilitates model-based inference by predictive memory sampling

Helen C Barron HELEN.C.BARRON@GMAIL.COM
Pavel Perestenko PAVEL.PERESTENKO@PHARM.OX.AC.UK
Renee S Koolschijn RENEE.KOOLSCHIJN@KEBLE.OX.AC.UK
Anna Shpektor ANNA.SHPEKTOR@MAGD.OX.AC.UK
Hamed Nili HAMED.NILI@PSY.OX.AC.UK
David Bannerman DAVID.BANNERMAN@PSY.OX.AC.UK
Tim EJ Behrens BEHRENS@FMRIB.OX.AC.UK
David Dupret DAVID.DUPRET@PHARM.OX.AC.UK

University of Oxford

Inferential reasoning involves the ability to form relationships between items or events that have not been directly experienced together. Such inferred novel relationships support flexible model-based decisions, where knowledge of the structure of the world is used to predict the value of a novel action. While inferential reasoning can readily be observed in the behavioural repertoire of both animals and humans, the underlying physiological mechanisms remain unclear. One possibility is that inference involves binding memories for disparate events 'offline', to facilitate representation of predicted outcomes at the time of choice. Here we adopt a parallel cross-species approach to reveal evidence for this mechanism, using ultra-high field MRI (7T) in humans and multi-unit electrophysiology and optogenetic manipulations in mice. During inferential decisions, in both humans and mice, we show evidence for sampling and reinstatement of predicted memory elements in the hippocampus. At a microscopic level, in dorsal CA1, this sampling and reinstatement of memory constitutes a temporal sequence in the firing of principal cells. In 'offline' periods, during sharp-wave ripple (SWR) oscillations, cells representing events that have not been directly experienced together cofire. This suggests that memories can be linked during SWRs, to facilitate predictive memory sampling during inferential choice. Finally, to test the necessary role of hippocampus in this mechanism, we show that inference is impaired when dorsal CA1 is optogenetically silenced during choice, and

mice instead revert to a model-free behavior where actions are determined by the running average of expected reward. Together, these findings suggest that the hippocampus supports inferential reasoning by binding disparate event elements during SWRs, before sampling the relevant predicted memory during choice. This mechanism may explain how novel links are inferred between events that have not been experienced together before.

II-107. Low- (25-35Hz) and mid-range gamma activity (35-55Hz) reflect distinct dynamics of cortical processing

Quentin Perrenoud QUENTIN.PERRENOUD@YALE.EDU
Austin Airhart AUSTIN.AIRHART@GMAIL.COM
Rong Mao RONG.M.MAO@GMAIL.COM
Jessica Cardin JESS.CARDIN@YALE.EDU

Yale University

Cortical activity in the gamma range is a hallmark of active wakefulness and is speculated to play an important role in neural circuit computations. Recent work has suggested that gamma-range activity in mouse cortex might occur in several distinct frequency bands, but the circuit-level impact and behavioral relevance of these patterns remains unclear. We recorded activity across the layers of the primary visual cortex of awake head-fixed mice using chronically implanted multi-contact silicon probes. Mice were either passively habituated to visual stimulation over several days or trained to associate visual stimuli with the availability of a reward in a Go/No Go task. We found that gamma activity in the local field potential occurs in two sub-bands reflecting distinct circuit dynamics. Rhythmic activity in the low gamma range (25-35Hz) strongly entrained units in layer 2/3 and was evoked mainly by large, high-contrast visual gratings. While nearly absent in naive animals, it strengthened over days of repeated visual stimulation, suggesting that its occurrence may be linked to passive learning of predictable visual regularities. In contrast, activity in the mid gamma range (35-55Hz) occurred in short bouts associated with a typical propagation of current sinks from layer 4 to layer 2/3 and then layer 5, suggesting that activity in this band reflects classic feedforward thalamo-cortical processing. Mid-range gamma activity entrained RS and FS units across layers 2/3, 4 and 5a and was strongly modulated by arousal and locomotion. Finally, each gamma band was differentially related to behavioral outcomes in the Go/NoGo task, with mid-range positively predicting and low-range gamma negatively predicting successful response. Our results thus suggest that these two types of patterned activity are associated with distinct motifs of interactions between cortical neurons and differentially impact local cortical circuit operations during active behaviors.

II-108. Efficient replay memory through sensory adaptation

Raymond Chua[1,2] RAY.R.CHUA@GMAIL.COM
David Jia[3] DAVID.W.JIA@GMAIL.COM
Rui Ponte Costa[1] RUI.COSTA@BRISTOL.AC.UK

[1]University of Bristol
[2]Technical University of Munich
[3]University of Oxford

The cortex adapts quickly to repetitive stimuli. Such adaptation suggests that dissimilar (or novel) experiences are more likely to be retained in memory. In contrast, recent deep reinforcement learning models rely on storing every single input into a hippocampal-like episodic memory. Here, inspired by rapid forms of synaptic plasticity – a key neural basis of sensory adaptation – we propose a reinforcement learning algorithm in which only dissimilar enough inputs are stored into the replay memory. We show that our method leads to a more efficient memory representation (reduced memory load), as similar inputs tend to be discarded. In addition, a model in which less experiences are discarded as the agent gradually learns to explore its environment performs similarly to standard

replay memory methods. This gradual change in adaptation is akin to the experimentally observed modifications of short-term plasticity over development and learning, and suggests an important role for this phenomenon in systems-level learning. Overall, our work shows how systems models of memory and learning can shed light on the function of synaptic plasticity and sensory adaptation.

II-109. Temporal integration windows for statistical estimation

Richard McWalter MCWALTER@MIT.EDU
Josh McDermott JHM@MIT.EDU

Massachusetts Institute of Technology

Stationary signals can be characterized by their statistics and are believed to be represented as such in sensory systems. Statistics must be measured from samples. Larger samples could produce more reliable estimates but run the risk of mixing signals with distinct statistics. Moreover, statistics might differ in the sample size needed for a criterion level of reliability. Here we ask whether sensory systems adopt different sampling strategies for different types of statistical measurements. One domain where statistical estimation is critical is sound texture – signals composed of many similar acoustic events (rain, fire, insects etc.). Sound textures appear to be represented with statistics averaged over windows that extend for several seconds [McWalter and McDermott, 2018]. However, many different statistics are needed to account for texture perception. It remains unclear whether there is a single global integration window for all statistics, or whether integration windows vary across statistics, potentially tailored to their variability. We measured integration windows for individual classes of statistics from a standard auditory texture model [McDermott and Simoncelli, 2011]. In a psychophysical experiment, listeners judged which of two sound textures was most similar to a reference texture. We measured performance for different stimulus durations, using textures synthesized to vary in individual classes of statistic. In all cases, texture discrimination improved with stimulus duration but then leveled off, presumably signaling the extent of a finite obligatory averaging window used to estimate the statistics mediating the decision. However, the performance asymptote occurred at different durations for different statistics, ranging from 150 milliseconds for the cochlear envelope mean (capturing the spectrum) to a few seconds for the power of slow modulation filters. The results suggest that the extent of time-averaging varies across texture statistics, as might be required to obtain stable estimates of functions differing in intrinsic variability.

II-110. Autoregressive Hidden Markov Models of language with intracranial human electrophysiology

Kiefer Forseth[1] KJFORSETH@GMAIL.COM
Xaq Pitkow[2] XAQ@RICE.EDU
Aram Giahi Saravani[2] ARAM.GI@GMAIL.COM
Nitin Tandon[3] NITIN.TANDON@UTH.TMC.EDU

[1]University of Texas Health Science Center
[2]Baylor College of Medicine
[3]University of Texas

Human speech production is an integrated multistage process that seamlessly translates appropriate conceptual representations in the brain to acoustic output. Instantiating psycholinguistic language models in neurobiology requires a sophisticated understanding of the dynamic interactions between functionally distinct regions. This can be accomplished using high-resolution recordings paired with an analytic approach capable of modeling discrete neural states. We used a novel implementation of Autoregressive Hidden Markov Models (ARHMMs) on direct intracranial recordings of human cortex to reveal a consistent and interpretable evolution of neural states during overt speech production.

Intracranial electrodes (n = 21,042; 126 patients), including both surface grid and penetrating depth electrodes, were implanted as part of an evaluation for epilepsy surgery. Patients performed common object naming cued by either pictures or spoken descriptions. A surface-based mixed-effects multilevel analysis of broadband gamma activity was used to identify loci with significant activity. These loci seeded network analysis with ARHMMs in 7 individuals with broad coverage of the language-dominant hemisphere using subdural grid electrodes and in 9 patients with dense focal coverage of perisylvian cortex using stereotactic depth electrodes.

ARHMMs revealed a consistent progression through 3 states for both visual- and auditory-cued naming. The first state corresponded to sensory processing: a fixed-length state for picture naming driven by the ventral visual stream and a stimulus-length state for spoken descriptions driven by the ventral auditory stream. This was followed by a language processing state – corresponding to semantic, lexical, and phonologic processes – that lasted until articulation began and featured broad connectivity emanating from inferior frontal gyrus. Finally, the third state was observed during articulation with network dynamics dominated by sensorimotor and acoustic cortex.

This new tool has great potential to drive isolation and analysis of the network states that drive complex cognitive processes essential for an expansive set of human behaviors.

II-111. Learning variability in the neural code of the retina

Nishal Shah[1] NISHALPS@STANFORD.EDU
Nora Brackbill[1] NBRACK@STANFORD.EDU
Ryan Samarakoon[1] SAMARAKOON.RYAN@GMAIL.COM
Colleen Rhoades[1] RHOADES@STANFORD.EDU
Alexandra Kling[1] ALEXTH.NN@GMAIL.COM
Alexander Sher[2] SASHAKE3@UCSC.EDU
Alan Litke[2] ALAN.LITKE@CERN.CH
Yoram Singer[3] SINGER@GOOGLE.COM
Jonathon Shlens[3] SHLENS@GOOGLE.COM
EJ Chichilnisky[1] EJ@STANFORD.EDU

[1]Stanford University
[2]University of California, Santa Cruz
[3]Google Brain

Motivation: Many studies of neural coding focus on understanding average properties, or properties that are invariant across experiments. However, inter-animal variability can be significant, and its structure is generally unknown. A fuller understanding of neural codes would involve a flexible, tractable and low-dimensional characterization of variation across individuals.

Approach: Light responses of complete populations of hundreds of ON and OFF parasol retinal ganglion cells (RGCs) were recorded in 84 macaque retinas. A single deep convolutional neural network, reflecting translational invariance of RGCs, was trained to predict responses from multiple retinas simultaneously. Retina-specific parameters that controlled the mean and variance of activation in different CNN layers were learned in a two-dimensional space.

Results: 1) A single shared model explained responses of multiple retinas simultaneously, with minimal loss of accuracy compared to a model trained for each retina separately. 2) Retinas with similar properties (response dynamics, nonlinearities, spike interval statistics) were embedded nearby in the parameter manifold. 3) Variations in ON and OFF parasol cell properties were captured by nearly orthogonal directions in the manifold, indicating independent variability. 4) Points on the manifold displayed known invariances, such as asymmetries between ON and OFF parasol cells. 5) Fixing all parameters except those in the manifold enabled estimation of retina-specific properties with minimal data, useful for closed-loop experiments. The model could be tuned to each retina efficiently merely by comparing pairs of points in the manifold, an approach that could permit estimation of individual retinal properties in behavioral experiments and biomedical applications.

Conclusion: A low dimensional representation captured the variation in neural codes across many retinas. This may be useful in the development of personalized retinal implants for treating blindness.

II-112. Latent structure in populations of truly continuous spike trains: inference with deep random splines

Gabriel Loaiza-Ganem GL2480@COLUMBIA.EDU
John Cunningham JPC2181@COLUMBIA.EDU
Columbia University

Whether via a state-space model, factor analyzer, or other dimensionality reduction technique, methods that infer low-dimensional structure from population level spiking activity have been a major theme in computational neuroscience over the last 5-10 years. While effective, most of these approaches bring with them two inconveniences. First, the consequence of increasing either trial length or temporal resolution is longer hidden trajectories, which effectively results in higher dimensional latent representations of trials. Second, time is typically discretized in both the latent structure (e.g. a linear dynamical system) and in the subsequent observation model (often a Poisson count distribution within a time bin, not the continuous process). In this work, instead of considering the spike counts on time bins to follow Poisson distributions, we directly model the spike trains as (truly continuous) Poisson Processes. The low-dimensional representation of a trial is mapped through an artificial neural network to obtain the parameters of the intensity function (which is given by a spline) of each spike train. By placing a distribution on these low-dimensional representations, we obtain a novel probability distribution over spline functions, which we call Deep Random Splines. By considering time as being continuous, we recover low-dimensional representations that are not trajectories and whose dimensionality does thus not depend on temporal resolution. We compare our model to the PfLDS model in both synthetic and real data and show that our model outperforms it, both in terms of data fit and quality of learned latent representations.

II-113. Stimulus statistics restructure correlated variability within and between visual areas

Elizabeth Zavitz ELIZABETH.ZAVITZ@MONASH.EDU
Maureen A Hagan MAUREEN.HAGAN@MONASH.EDU
Brian H Oakley BRIAN.OAKLEY@MONASH.EDU
Yan T Wong YAN.WONG@MONASH.EDU
Nicholas SC Price NICHOLAS.PRICE@MONASH.EDU
Monash University

Populations of neurons in the middle temporal area (MT) represent motion direction based on the input they receive from direction-selective neurons in primary visual cortex (V1). Although we have strong models for how these populations should behave based on single-unit recordings to carefully parameterised stimuli, the visual world contains information at a range of spatial scales, and this information is often phase aligned. These properties will produce different population activity, and thus representations, of stimulus motion. Here, we measure noise correlations to understand how the structure of visual information impacts the way it is represented across neuronal populations. We made simultaneous population recordings with arrays placed in V1 and MT in five anaesthetised marmosets and measured inter-trial and inter-neuron variability under stimulus conditions that varied in spatial frequency content and phase alignment. Overall, we found that correlated variability (noise correlation) within and between cortical areas was reduced when stimuli were most naturalistic. Sine wave gratings induced the highest shared variability within and between visual areas; adding more spatial frequency information at random phases reduced the shared variability; and it was further reduced when the cross-frequency energy was phase-aligned as a square wave. Between areas in particular, we found that the relationship between the representation in a V1 neuron and its paired MT neuron (direction selective or not, how much receptive fields

overlapped) were also related to the strength of the noise correlations we observed. These results demonstrate that shared variability depends on the functional network, and that naturalistic stimuli engage these networks in a different way to sine wave gratings. They also demonstrate how changes in shared variability in one area can propagate throughout the visual hierarchy.

II-114. Low-dimensional hindbrain responses during optokinetic behavior in larval zebrafish

Claudia Feierstein[1,2] CLAUDIA.FEIERSTEIN@NEURO.FCHAMPALIMAUD.ORG
Michelle de Goeij[2] M.H.M.DEGOEIJ@GMAIL.COM
Ruben Portugues[3] RPORTUGUES@NEURO.MPG.DE
Christian Machens[4] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Michael Orger MICHAEL.ORGER@NEURO.FCHAMPALIMAUD.ORG

[1]Chamaplimaud Foundation
[2]Champalimaud Research
[3]Max Planck Institute of Neurobiology
[4]Champalimaud Foundation

How animals convert sensory information into motor commands is a central question in neuroscience. We address this question in larval zebrafish using the optokinetic response (OKR), a rotating movement of the eyes in response to whole-field motion. Despite its apparent simplicity, populations of neurons throughout the brain are active during the OKR, as shown by whole-brain imaging of zebrafish larvae (1). Here, we use dimensionality reduction techniques to relate sensory and motor variables to population activity in the hindbrain, an area enriched in eye-related activity. Specifically, we imaged neuronal activity, while monitoring the behavioral response to gratings that rotated with different speeds. We identified putative neurons (regions of interest or ROIs, up to 2000 for individual fish) and generated a set of more than thirty regressors related to the visual stimuli, the eyes, and the tail. To elucidate the representation of information at the population level, we used reduced-rank regression, a method that condenses the regressors into a smaller number of ‘features’ in order to predict the fluorescence traces of all ROIs. Each feature is simply a linear combination of the original regressors. Despite the potential complexity of the visuomotor transformation, five to seven features sufficed to account for the population activity. The majority of the activity could be explained by the first two features, which represented vergent and rotational motion, respectively. Each feature combined information about eye and stimulus motion. These findings implicate the respective hindbrain areas in the transformation of visual motion stimuli into eye movements. Most of the ROIs were explained exclusively by one of these two features, indicating that vergent and rotational eye motion are represented in separate networks. Just as described in many cortical areas, our analysis revealed that population activity in the hindbrain, as related to our regressors, lies in a low-dimensional subspace.

III-1. Place cells in the marmoset monkey hippocampus during free navigation

Cory Miller CORYMILLER@UCSD.EDU
Hristos Courellis HCOURELL@ENG.UCSD.EDU

University of California, San Diego

The role of mammalian medial temporal lobe structures– such as hippocampus – for spatial navigation has been amongst the most thoroughly investigated topics in Neuroscience for several decades. The discovery of place cells that encode self-position during active navigation of real-world environments in rats has been replicated in several other mammalian species, but notably not in primates. Studies of spatial navigation in primates have focused on how hippocampus encodes space through visual exploration and shown compelling evidence of both place cells and grid cells in the medial temporal lobe in this context - including in virtual environments. During these latter experiments, subjects are physically restrained and explore their respective virtual spaces using a joystick, receiving visual feedback through a monitor to simulate locomotion yet lacking the proprioceptive cues critical to place field formation in rodents. Despite these experimental differences, it has been argued that primate spatial representations differ fundamentally from rodents. Rather than encode self-position in space, the evolution of foveal vision in primates is conjectured to drive specializations for representations of scenes through visual - rather than physical - exploration. Remarkably, there have been no previous studies examining primate hippocampus as individuals physically explored space. To address this fundamental gap in our knowledge, we recorded the activity of isolated single neurons in the hippocampus of freely-moving marmoset monkeys (Callithrix jacchus) as they actively explored real-world environments. Here we report the discovery of place cells in primate hippocampus for self-position during physical exploration that share many of the characteristics observed in other mammals. Although theta oscillations were prevalent in marmoset hippocampus, it was not closely coupled to spike timing or locomotor velocity in the monkeys. These results suggest that space is not solely represented by visual exploration in primates and have significant implications for models of spatial encoding in hippocampus.

III-2. Inference, validation and predictions about statistics and propagation of cortical spiking in vivo

Jens Wilting JENS.WILTING@DS.MPG.DE
Viola Priesemann VIOLA.PRIESEMANN@DS.MPG.DE

Max Planck Institute for Dynamics and Self-Organization

Electrophysiological recordings of spiking activity can only access a small fraction of all neurons at the same time. This spatial subsampling has hindered characterizing even basic properties of collective spiking in cortical networks. In particular, two contradictory hypotheses have prevailed for more than a decade: the first proposed an asynchronous irregular, the second a critical state. While distinguishing them is straightforward in models, we show that in experiments classical approaches fail to correctly infer network dynamics, because subsampling can bias measures as basic as correlation strengths. Deploying a novel, subsampling-invariant estimator, we find that in vivo dynamics do not comply with either hypothesis, but instead occupy a narrow “reverberating regime" (RR), consistently across multiple mammalian species and cortical areas. We validate this finding by showing that a generic spiking model in the RR reproduces numerous properties of in vivo recordings, e.g. spike count cross correlations between neurons, Fano factors, and inter-spike-interval, rate, and avalanche distributions.

We then use the model to predict experimentally inaccessible network properties: the sizes and durations of responses to minimal perturbations, network timescales and network Fano factors, and the fraction of external vs. recurrent input to cortical networks. These basic response properties strongly constrain how networks can process and integrate stimuli. For example, recent results indicate that network timescales increase along the visual and auditory pathways, potentially providing a dynamical substrate for increasing temporal receptive fields. As network function and dynamics are mutually tied to eachother, our results provide insight and clearcut constraints on how cortical networks achieve their performance.

III-3. Generalized phase locking analysis of electrophysiology data

Michel Besserve MICHEL.BESSERVE@TUEBINGEN.MPG.DE
Shervin Safavi RESEARCH@SHERVINSAFAVI.IO
Vishal Kapoor VISHAL.KAPOOR@TUEBINGEN.MPG.DE
Fanis Panagiotaropoulos THEOFANIS.PANAGIOTAROPOULOS@TUEBINGEN.MPG.DE
Nikos Logothetis NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE

Max Planck Institute for Biological Cybernetics

Brain information processing likely relies on cooperative interactions between neural populations at multiple scales. Growing evidence suggests that network oscillations, as observed in Local Field Potentials (LFP), are instrumental to the spatiotemporal coordination of these interactions. Therefore, investigating the coupling between spatiotemporal patterns of LFP and spiking activity is instrumental to understand distributed neural information processing. Common approaches to investigate this coupling are restricted to pairwise spike-LFP interactions, which are suboptimal for modern datasets with hundreds of simultaneous recording sites. Capturing efficiently the overall spike-LFP coupling structure in this high dimensional setting is of paramount importance to exploit the full potential of modern electrophysiology recording techniques. We develop a Generalized Phase Locking Analysis (GPLA), a multivariate extension of phase locking analysis, by gathering pairwise complex phase locking information in a rectangular matrix and summarize its structure with the largest singular value and the corresponding singular vectors. Singular vectors represent the dominant LFP and spiking patterns and the singular value, called generalized Phase Locking Value (gPLV), characterizes the strength of the coupling between LFP and spike patterns. We further investigate statistical properties of the gPLV and develop a statistical testing framework. Compared to pairwise approaches, simulations with networks of Leaky Integrate and Fire (LIF) neurons show that GPLA: (1) can reliably retrieve the coupling between spikes and LFP with lesser amount of data and (2) exploits optimally the activity of multiple units to increase the statistical power while preserving individual coupling properties. Application to recordings from Utah arrays in macaque prefrontal cortex reveals a previously undetected large-scale coupling through an LFP traveling wave in the beta band synchronized with an array-wide synchronous spiking event. These results illustrate the interest of GPLA to assess global relationships between spatiotemporal patterns of spikes and network oscillations.

III-4. Remembered reward locations restructure entorhinal spatial maps

Kiah Hardcastle KIAHHARDCASTLE@GMAIL.COM
William Butler WBUTLER@STANFORD.EDU
Lisa Giocomo GIOCOMO@STANFORD.EDU

Stanford University

The ability to navigate to a remembered reward location is critical to animal survival. In order to execute such a task, animals must form an internal map of their surrounding environment, while simultaneously identifying the location of the reward within it. Decades of research have shown that neural representations of self-localization reside in the hippocampus and medial entorhinal cortex (MEC). In particular, hippocampal place cells, entorhinal grid cells, and entorhinal non-grid position-encoding cells all fire in specific locations within an environment and can be used to decode an animal's current location. Recent work found that a subset of hippocampal place cells specifically encode the location of a reward, thus providing a potential neural substrate for reward-localization and reward-driven navigation (Gautier et al., 2018). However, it remains unclear whether MEC also integrates information regarding learned reward locations. In addition, recent work has shown that MEC signals show remarkable flexibility in their spatial response features (Hardcastle et al., 2017), in contrast to initial work suggesting these signals are inflexible and context-independent. To examine whether MEC neurons are also capable of adapting their activity in a reward-specific manner, we recorded MEC neurons in two separate environments while rats performed either a free foraging task or a spatial memory task. We found that the majority of spatially-modulated neurons, including grid cells, re-mapped their activity to represent the reward location. This resulted in improved accuracy of position decoding near the reward location in the spatial task environment. Further, we observed that re-structuring of spatial maps occurred throughout the task environment recordings, not only while the animal actively performed the spatial task. Combined, this indicates MEC can encode reward location, and that the presence of a remembered reward location is sufficient to drive the emergence of discrete entorhinal maps of space.

III-5. Sequential activity in ramping neurons during evidence accumulation

Jonathan Boyd-Meredith JTB3@PRINCETON.EDU
David W Tank DWTANK@PRINCETON.EDU
Carlos Brody BRODY@PRINCETON.EDU

Princeton University

During gradual accumulation of evidence for decision making, neurons across many brain regions have been observed to exhibit trial-averaged firing rates that ramp with slopes proportional to the strength of evidence in favor of a preferred decision option [4]. These neural signatures are thought to support working memory across a wide variety of decision tasks. A recent study from Hanks and colleagues found that in rats the average population activity in parietal and frontal cortical regions stably encodes the magnitude of evidence in favor of each alternative throughout the trial [2]. However, results from mice performing cue-guided navigation have suggested that working memory may rely on sequential activation of neurons throughout behavior, such that information is passed from one set of neurons to the next [3]. Here, we ask whether the population of neurons analyzed by Hanks and colleagues exhibits sequential activity. We find that the majority of the neurons in this dataset, including those modulated by accumulated evidence, exhibit reliably timed peaks of activity that span all periods of the task including trial initiation, stimulus presentation, and movement. We develop simple tests for identifying sequential activity and describe the pitfalls of alternative methods. Our results suggest a hybrid model in which working memory is supported by diverging sequences of neural activity where relative peak timing is reliable across neurons and modulated by the strength of evidence for the decision.

III-6. A data-constrained model for operant visual learning behavior in Drosophila

Ann Hermundstad HERMUNDSTADA@JANELIA.HHMI.ORG
Jason Wittenbach JASON.WITTENBACH@GMAIL.COM
Vivek Jayaraman VIVEK@JANELIA.HHMI.ORG
Chuntao Dan DANC@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

Many animals navigate their visual environment to find food and conspecifics while avoiding predators. Some can adjust their navigational strategies based on past experience, for example by remembering the consequences of being in specific visual surroundings. Drosophila melanogaster has been shown to associate visual patterns with heat punishment in a flight simulator [1], remember the location of a cool spot in an otherwise hot environment using visual cues [2], and associate color or light intensity with a reward or punishment [3]. However, visual conditioning in flies has typically been quantified via performance indices [1-3]. Here we focused on the precise pattern of movements that these animals use to sample their environment, and the algorithms by which they associate these movements with positive or negative consequences and thereby modify future behavior. We trained tethered flies in a flight simulator in which they have closed-loop control of a one-dimensional visual scene [1,4]. A subset of visual patterns is paired with an aversive heat punishment. Flies learn within minutes to avoid dangerous sectors and direct their movements towards safer sectors. This learning is accompanied by a modulation of the size and directionality of fictive turns, and the duration of periods of straight flight, that depends on heading relative to safety; this modulation is not observed in negative control flies. We developed an actor-critic learning algorithm that is constrained by the distributions of movements we observe in the data, and we show that this algorithm captures key modifications in behavior during and after training. Together, these results constrain mechanistic models of operant learning, and they provide clues about sources of individual variability observed across animals.

[1] Wolf, Heisenberg (1991), J. Compar. Physiol. A. [2] Ofstad, Zuker, Reiser (2011), Nature. [3] Schnaitmann et al. (2010), Front. Behav. Neurosci. [4] Reiser, Dickinson (2008), J. Neurosci. Methods.

III-7. Recovery of consciousness after traumatic brain injury: Biomarkers and a mechanistic model

Sima Mofakham[1] S.MOFAKHAM@GMAIL.COM
Adam Fry[2] ADAM.FRY@MOUNTSINAI.ORG
Joseph Adachi[3] JOSEPH.ADACHI@STONYBROOKMEDICINE.EDU
Bradley Ashcroft[3] BRADLEY.ASHCROFT@STONYBROOKMEDICINE.EDU
Nathan Winans[3] NATHAN.WINANS@STONYBROOKMEDICINE.EDU
Justine Liang[3] JUSTINE.LIANG@STONYBROOKMEDICINE.EDU
Himanshu Sharma[3] HIMANSHU.SHARMA@STONYBROOKMEDICINE.EDU
Susan Fiore[3] SUSAN.FIORE@STONYBROOKMEDICINE.EDU
Memming Park[3] MEMMING.PARK@STONYBROOK.EDU
Charles Mikell[3] CHARLES.MIKELL@STONYBROOKMEDICINE.EDU

[1]Stony Brook University Hospital
[2]Icahn School of Medicine at Mount Sinai
[3]Stony Brook University

Although the underlying mechanism supporting consciousness remains unknown, recent evidence suggests that thalamic coordination of frontoparietal network (FPN) activity is critical. To better understand the relevance of the functional integrity of FPN, we directly probed FPN in recovering comatose patients following traumatic brain injury (TBI), using depth electrodes inserted into the anterior cingulate gyrus (ACC) and dorsolateral prefrontal cortex (DLPFC). We also interrogated the FPN using single-pulse electrical stimulation. Using spontaneous and stimulation evoked-responses, we identified features that correlate with reversibility of cortical dysfunction, and the eventual recovery of consciousness. At a short timescale ($<$1s), we found that evoked responses in those who recovered were multiphasic, complex, and showed greater variability. In contrast, in patients with irreversible cortical dysfunction, we found that the evoked response had low inter-trial variability, was monophasic, and exhibited a slower return to baseline. Moreover, over longer timescales (up to 30s), we found that only in patients with irreversible cortical dysfunction, stimulation trains reliably produced a temporal pattern that extended long after the stimulation. MRI of these patients revealed injury to thalamocortical projections. To investigate the underlying mechanism behind these observations, we built a thalamocortical spiking model, in which thalamic input to the cortex modulates excitability and effective connectivity across cortical populations. Our in silico study suggests that the lack of thalamic input disrupts long-range cortical communication, leading to unbalanced inhibition/excitation, resulting in the characteristic simple, monophasic evoked-responses observed in coma patients. These findings suggest that it may be possible to predict whether cortical damage will lead to reversible or irreversible impairment based on measured FPN integrity inferred from evoked-responses, which could potentially serve as a prognostic indicator in clinical treatment and management of comatose patients.

III-8. Geometry of neural computations unifies working memory and planning

Daniel Ehrlich DANIEL.EHRLICH@YALE.EDU
John D Murray JOHN.MURRAY@YALE.EDU

Yale University

The ability to maintain information in working memory forms one of the key building blocks of cognition. Traditionally models of working memory have been conceived as a form of extended sensory perception that preserves prior information to inform current computation. In this study, we recast working memory representations as an active part of ongoing goal-oriented computations. Using both Recurrent Neural Networks (RNNs) and human subjects performing a set of logical tasks, encompassing several known working memory paradigms, we demonstrate how such working memory representations may help to minimize computational complexity of cognitive operations. The contextual delayed logic (CDL) task, is a novel delayed decision task in which subjects use a rule context to flexibly modulate their response to a pair of binary stimulus inputs separated by a mnemonic delay. On

each trial the subject sees one of ten Boolean logical operations (e.g. 'and'), must hold the first stimulus identity across a delay, and then upon presentation of the second stimulus, must apply the stated logical operation. The different retrospective and prospective demands between rules allow us to disentangle memory, planning, computational or rule specific influences on representations and behavior. We trained RNNs to perform the CDL task, analyzing resulting optimized networks and state trajectories. We found that delay epoch representations were organized computationally and prospectively with overlapping sub-problems sharing network resources (e.g. fixed points, transitions). Further, delay states were geometrically arranged within a stable mnemonic subspace to linearize the overall stimulus-output classification problem. We demonstrate that resulting representations recapitulate known properties of prefrontal delay task firing patterns such as mixed selectivity and tuning reversals. Lastly we designed a human version of our CDL task paradigm to identify behavioral correlates of the three strategic elements identified in our RNNs memory representations: precomputation during mnemonic delays, prospective representational organization, and sub-problem linearization.

III-9. A cellular mechanism makes switches in brain states compatible with synaptic plasticity

Kathleen Coutisse KATHLEEN.COUTISSE@ULIEGE.BE
Guillaume Drion GDRION@ULIEGE.BE

University of Liege

Animal performance relies on their ability to quickly process, analyze and react to incoming events, as well as to learn from experience to constantly increase their knowledge about the environment. Learning and memory are attributed to the ability of neurons to modify their connections with other cells based on experience, a property called synaptic plasticity. Synaptic plasticity mechanisms often exploit the level of correlation in the activity of connected neurons, and can therefore be affected by abrupt changes in neuronal excitability. On the other hand, brain information processing is constantly shaped by fluctuations in neuronal rhythmic activities at the cellular and population levels, each defining distinctive brain states. Switches between these brain states can be fast and localized, such as e.g. those observed in different brain areas prior to movement initiation, or global and long lasting, such as those observed during the wake-sleep transition. The coexistence of these two mechanisms raises challenging questions: how can switches in brain states remain reliable despite of constant rewiring of neuron connectivity, and how is synaptic plasticity affected by switches in brain states? In this work, we highlight the critical role of a cellular dynamical property in the generation of switches in brain states that are compatible with changes in network connectivity and cellular heterogeneity. This dynamical property, called slow regenerativity, is accessible to all neurons that embed slowly activating voltage-gated calcium channels or slowly inactivating potassium channels in their membrane, yet it is largely overlooked in computational and mathematical neuron models and absent from all available hybrid models. To demonstrate this point, we compare the robustness of 5 published thalamic neuron models at the cellular, circuit and population levels. We show that the robustness of rhythms at the population level correlates with the presence or absence of slow regenerativity at the cellular level.

III-10. Orbitofrontal cortex promotes dynamic updating of risk preferences

Christine Constantinople[1] CONSTANTINOPLE@NYU.EDU
Charles Kopec[2] CKOPEC@PRINCETON.EDU
Carlos Brody[2] BRODY@PRINCETON.EDU
Alex Piet[2] ALEX.PIET@GMAIL.COM
Peter Bibawi[2] PBIBAWI@PRINCETON.EDU
Athena Akrami[3] ATHENA.AKRAMI@GMAIL.COM

[1]New York University
[2]Princeton University

[3]Sainsbury Wellcome Centre

A core assumption of most economic and psychological theories is that individual risk preferences are stable traits. However, documented violations of this assumption indicate that preferences can be flexible. Here, we leveraged a novel behavioral paradigm and circuit-based tools in rodents to identify the orbitofrontal cortex (OFC) as required for this flexibility. We exploited a phenomenon called the "hot-hand fallacy," which refers to an increased willingness to take risks following wins, even when outcomes are independent. This pervasive phenomenon has wide-ranging consequences, including in finance4, and can be easily studied in trial-based paradigms with repeated sequential choices. We trained rats to repeatedly choose between guaranteed and probabilistic rewards. Rats exhibited the hot-hand bias: they were more likely to gamble following risky wins. Neurons in OFC encoded reward history and whether choices were risky or safe. Optogenetic inhibition of OFC at the time of the choice, but not during the trial, eliminated the hot-hand bias, sparing other trial history effects. We have therefore identified the neural basis of the hot-hand fallacy, a ubiquitous decision bias that is a famous exemplar of irrational behavior in humans and an experimentally tractable entrypoint into flexible risk preferences generally.

III-11. Temporal dictionary learning for calcium imaging analysis

Gal Mishne[1] GAL.MISHNE@YALE.EDU
Benjamin Scott[2] BBS@BU.EDU
Stephan Thiberge[3] THIBERGE@PRINCETON.EDU
Nathan Cermak[4] CERMAN07@GMAIL.COM
Jackie Schiller[4] JACKIE@TECHNION.AC.IL
Carlos Brody[3] BRODY@PRINCETON.EDU
David W Tank[3] DWTANK@PRINCETON.EDU
Adam Charles[3] ADAMSC@PRINCETON.EDU

[1]Yale University
[2]Boston University
[3]Princeton University
[4]Technion

Optical calcium imaging is a versatile imaging modality that permits the recording of neural activity, including single dendrites and spines, deep neural populations using two-photon microscopy, and wide-field recordings of entire cortical surfaces. To utilize calcium imaging, the temporal fluorescence fluctuations of each component (e.g., spines, neurons or brain regions) must be extracted from the full video. Current methods typically use a combination of both a-priori spatial and temporal statistics to isolate each fluorescing source in the data, along with the corresponding time-traces (Mukamel et al. 2009, Pachitariu et al. 2013, Pnevmatikakis et al. 2016). Such methods often rely on strong spatial regularization and temporal priors that can bias time-trace estimation and that do not translate well across imaging scales. We propose to instead emphasize time-trace inference, using only weak spatial information to relate per-pixel generative models across a field-of-view. Our method, based on spatially-filtered Laplacian-scale mixture models (Garrigues & Olshausen 2010, Charles & Rozell 2014), introduces additional regularization and spatial-smoothing hyper-priors to the dictionary-learning framework, where the learned dictionary consists of the fluorescing components' time-traces. We demonstrate our approach on calcium imaging at different scales, including population, dendritic and widefield data. First we compare our method to a current state-of-the-art algorithm, Suite2p, on the publicly available Neurofinder dataset. The lack of strong spatial contiguity constraints allows our model to isolate both disconnected portions of the same neuron, as well as small components that would otherwise be over-shadowed by larger components. In the latter case, this is important as such configurations can easily cause false transients which can be scientifically misleading (Gauthier et al. 2018). On dendritic data our method isolates spines and dendritic firing modes. Finally, our method can partition widefield data in to a small number of components that capture the scientifically relevant neural activity.

III-12. Harnessing noise and disorder to rescue optimal predictive coding from transmission delays

Jonathan Timcheck TIMCHECK@STANFORD.EDU
Jonathan Kadmon KADMONJ@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Kwabena Boahen BOAHEN@STANFORD.EDU

Stanford University

Both neuronal noise and disordered connectivity are prominent features of cortical circuitry. However, it is unclear whether such heterogeneity is biologically unavoidable, or whether it benefits computation. Indeed, in classical models of networks with weakly correlated noise, decoding errors decrease as the square-root of network size. In contrast, a recent predictive coding framework, developed by Deneve and colleagues, achieves superclassical scaling in which decoding error decreases linearly with network size. The key is that each neuron spikes to reduce error only when the error exceeds a threshold. Importantly, each spike must inhibit similarly encoding neurons with zero latency to prevent overcorrection. This overall dynamics leads to asynchronous activity, which enables high coding fidelity. However, it is unclear whether superclassical coding fidelity can be maintained in the presence of biologically realistic spike transmission delays, as well as noise and disordered neural connectivity. Indeed transmission delays tend to synchronize neurons thereby impairing coding fidelity.

In this work, we derive an analytic theory and simulations for predictive coding that simultaneously considers axonal propagation delays, neuronal noise, and disordered connectivity. We find, for propagation delays less than the network interspike interval, there exists an optimal noise level or weight disorder that minimizes error. This optimal level arises from two competing effects: the need to introduce some heterogeneity to prevent synchronization, but not so much to destroy coding fidelity. Moreover, we show, even with propagation delays, that at the optimal noise level, the superclassical performance of predictive coding remains intact, with error scaling inversely with network size. We find also that disordered connectivity can achieve the same beneficial effects as neuronal noise. Thus overall our theory of predictive coding under transmission delays demonstrates that either neuronal noise or deterministic chaos induced by disordered connectivity can play a beneficial role by facilitating asynchronous cortical computation.

III-13. A fundamental difference between primate and mouse object vision

Francisco Luongo FLUONGO@GMAIL.COM
Lu Liu LIULU@CALTECH.EDU
Doris Tsao DORTSAO@CALTECH.EDU

California Institute of Technology

Studies of the rodent visual system have exploded in recent years, revealing tuning to elementary image features similar to primates. How are these features further transformed? In the primate, researchers have argued that a key transformation is “figure-ground segmentation.” Behaviorally, humans can delineate figures from backgrounds using only kinetic borders independent of low-level texture cues. Neurally, long-latency enhancement of responses in early visual areas to regions within figure compared to background has been reported. Do mice also show behavioral and neural signatures of figure-ground segmentation? We trained mice on a figure-ground segmentation task where figures were defined by gratings and naturalistic textures, moving counterphase to the background. Unlike primates, mice were incapable of figure-ground segmentation using kinetic borders alone. While mice could report the location of gratings, they were at chance for figures defined by naturalistic textures. Remarkably, mice could learn to localize texture-defined figures after many weeks of training, but when tested on new textures, were unable to generalize. This suggests a strategy of memorizing a lookup table of noise patterns, a cognitively impressive feat. We also recorded visual responses to the same stimuli in V1, RL, and LM using both 2-photon imaging and electrophysiology. Neural responses were consistent with the behavior, revealing robust position decoding from neural populations for gratings but not naturalistic textures. Lastly, we simulated neural

responses using a model of V1 incorporating orientation-dependent suppression, and a deep network, VGG16. While the former model failed to reproduce the neural results, the latter model was successful. Interestingly, the best model for mouse V1 neural data came from mid to late layers. Taken together, these findings reveal a fundamental difference between primate and mouse mechanisms for object segmentation, showing that orientation contrast is critical for object perception in mice, unlike in primates.

III-14. The temporal dimension of our memories

Georgia Christodoulou GEORGIA.CHRISTODOUL@GMAIL.COM
Everton J Agnes EVERTON.AGNES@CNCB.OX.AC.UK
Tim P Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Specific motor actions, like speaking or writing words, are learned and reproduced with great reliability over a lifetime. Their execution demands continuous and precise orchestration of many muscle groups and thus many neurons, but we don't know what the neural substrate of such a temporal memory looks like. In other words, what are the dynamics and architecture of a network that can support the formation of distinct and noise robust spatiotemporal dynamics, and what is the capacity of such networks? To answer these questions, we study memories as the dynamical outputs of a Stability Optimised Circuit (SOC) - a recurrent neuronal network with strong excitatory amplification that is perfectly balanced by inhibitory connections. SOCs have been shown to create multi-modal, amplified stimulus responses that can be interpreted as motor patterns. However, it is not clear how reliable and robust SOCs are, and what their capacity is. In other words, they have not been studied as memory networks.

Here, we study the memory capacity and noise robustness of SOCs. We analyse their dynamics with respect to encodability and retrievability, and test network architectures ranging from random Erdos-Renyi to Small World with both static and plastic connections. Next, we show that memories are more robust when Hebbian plasticity is employed to mould network architecture. Finally, we find that the performance of SOCs can be controlled by dynamic thresholds that depend on the context of a given task: high threshold for elevated capacity (e.g., to spot small differences) and low for robustness (e.g., to identify speech in noise). Our work is a formal analysis of highly dynamical recurrent neuronal networks as storage systems, and the development of tools to analyse neuronal activities as temporal memories in model and experiment.

III-15. Automated discovery of precise spike timing by simple time warping models

Alexander Williams[1] ALEX.H.WILLIA@GMAIL.COM
Ben Poole[2] POOLEB@GOOGLE.COM
Niru Maheswaranathan[2] NIRUM@GOOGLE.COM
Ashesh Dhawale[3] DHAWALE@FAS.HARVARD.EDU
Tucker Fisher[1] TGFISHER@STANFORD.EDU
Christopher Wilson[4] CDW291@NYU.EDU
David Brann[3] DAVID_BRANN@G.HARVARD.EDU
Eric Trautmann[1] ETRAUTMANN@GMAIL.COM
Stephen Ryu[1] SEOULMANMD@GMAIL.COM
Roman Shusterman[5] ROMAS@UOREGON.EDU
Dmitry Rinberg[4] RINBERG@NYU.EDU
Bence Olveczky[3] OLVECZKY@FAS.HARVARD.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU
Surya Ganguli[1] SGANGULI@STANFORD.EDU

[1]Stanford University
[2]Google Brain
[3]Harvard University
[4]New York University
[5]University of Oregon

Neural spike times are typically aligned to stimulus triggers or behavioral events on repeated trials. However, on each trial, neural activity is often shifted and/or skewed in time due to variable delays in sensation, decision-making, reaction times, and other unobserved processes. We developed a class of statistical models that capture these effects on a trial-by-trial basis with linear and piecewise-linear time warping functions. These models yield more interpretable results than classical time warping methods, are less prone to overfitting, and are computationally scalable to large datasets. On multielectrode recordings from different animal models (Rhesus monkey, mouse, and rat) and different tasks (discrete reaching, timed motor production, and olfaction), linear warping reveals precise firing patterns that are invisible in raw data. Despite being fit to neural activity, the learned time warps can be tightly correlated with measurable behaviors on each trial—e.g., explaining $\sim 80\%$ of the variability in reaction times in cued reaching in primates, and $\sim 90\%$ of the variability in inhalation onset in mouse olfaction. Intriguingly, the optimal time alignment is partly decoupled from behavior in other cases, suggesting that unmeasured, internal variables can strongly influence the timing of neural dynamics. In rat motor cortex during timed motor production, we identify precisely timed theta-band oscillations that are imperceptible in raw spike trains. In some neurons, these oscillations are precisely initiated at (but not phase-locked to) the first motor action, suggesting a potential mechanistic importance for the task. In primate motor cortex during movement preparation, we identify similar oscillations at the level of multiunit spikes, which are not phase-locked to behavior. Thus, even in systems close to the sensory/motor periphery which are thought to be tightly locked to experimental cues, simple time warping models can reveal striking spike patterns that are likely to be otherwise overlooked.

III-16. Flickering hope? Inferred hippocampal maps and splitter cells support multi-task learning

Tamas Madarasz TAMASMADARASZ1@GMAIL.COM
Tim EJ Behrens BEHRENS@FMRIB.OX.AC.UK

University of Oxford

Humans and animals show remarkable flexibility in adjusting their behaviour when their goals, or rewards in the environment change. While such flexibility is a hallmark of intelligent behaviour, these multi-task scenarios remain

an important challenge for ai algorithms and neurobiological models.

Factored representations enable flexible behaviour by abstracting away general aspects of an environment or task from those prone to change. The successor representation (SR) for example factors the value of actions into components representing expected outcomes and corresponding rewards, useful when rewarded locations in the environment can change. While the SR framework has been proposed to underlie a hippocampal predictive map, it also suffers from important limitations because of the representation's dependence on the behavioural policy, under which expected future states are calculated. A substantial change in the environment's rewards can require visiting vastly different parts of the state space, but the current policy does not map out routes to these rewards, resulting in a lack of flexibility and positive transfer.

Here we present a novel learning algorithm that combines the SR framework with nonparametric inference and the clustering of the reward-space, while explaining important neurobiological signatures of 'splitter' cells and hippocampal place cell representations. Our algorithm dynamically samples from a flexible number of distinct SR maps using inference about the current reward context, with inference enabled by appropriately stored memories of experienced rewards. It comfortably outperforms competing algorithms both in settings with known and unsignalled rewards. Importantly, it also reproduces the 'flickering' behaviour of hippocampal maps seen when rodents navigate to changing reward locations, and gives a quantitative account of trajectory-dependent hippocampal representations and their similarity structure in route-dependent navigation tasks. We thus provide a novel framework for the analysis of a growing number of experimental paradigms and real-world tasks with changing goals, and volatile reward environments.

III-17. Evidence of behavioral and neural interaction between task complexity and state-space uncertainty during reinforcement learning

Geon Young Park[1] PKY3436@KAIST.AC.KR
Dongjae Kim[1] KIM10481@KAIST.AC.KR
John O'Doherty[2] JDOHERTY@CALTECH.EDU
Sang Wan Lee[1] SANGWAN@KAIST.AC.KR

[1]Korea Advanced Institute of Science and Technology (KAIST)
[2]California Institute of Technology

A major open question is how does the brain allocate control over behavior to arbitrate between two distinctive learning and control systems: model-based and model-free reinforcement learning (RL). In a previous study we provided evidence to suggest that the reliability of the prediction of the two types of RL is a key variable for arbitration [1]. However, another key variable, the role of task complexity, has hitherto remained unexplored. Here we investigate the possibility that task complexity interacts with the prediction reliability of model-based and model-free RL to guide the arbitration process. We developed an augment computational model of arbitration to implements the hypothesis that arbitration is sensitive to both the complexity of the state-space and the degree of prediction reliability. To test this, we implemented a novel behavioral task in which state-transition uncertainty and state-space complexity are independently manipulated. By combining a behavioral analysis with comprehensive computational modelling, we found behavioral evidence to suggest that task complexity is utilized in the arbitration process alongside state-space uncertainty. Model-based RL control becomes more dominant as task complexity increases. However, we also found that when both uncertainty and complexity are high, humans tend to relinquish model-based control, resorting to model-free control. These findings support the view that uncertainty and task complexity interact to guide human behavior. In subsequent model-based fMRI analyses, we also found neural evidence for the interaction between these two variables in the bilateral inferior prefrontal cortex. These findings not only provide insight into how both uncertainty and task complexity variables contributes to arbitration control in a dynamically changing environment, but also significantly advance understanding of the nature of computations being implemented in the inferior prefrontal cortex during the arbitration process.

III-18. Multi-sensory integration in the mouse meso-scale connectome using a network diffusion model

Constantine Dovrolis[1] CONSTANTINE@GATECH.EDU
Eva Dyer[1,2] EVADYER@GATECH.EDU
Kamal Shadi[1] KAMAL.SHADI85@GATECH.EDU

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

Having a structural network representation of the mesoscale connectivity in the brain is instrumental in analyzing communication dynamics and information processing. Current methods for predicting information flow from structural connectomes rely on assumptions about whether a connection is feedforward or feedback. However, this type of approach is cortico-centric and is typically constrained to the study of flow in specific sensory systems. Thus, new approaches are needed to predict the flow of information across brain regions from structural networks.

In this work, we make steps towards understanding multi-sensory information flow using a network approach. In particular, we model the flow of evoked activity, initiated by stimuli at primary sensory regions, using the Asynchronous Linear Threshold (ALT) diffusion model. The ALT model captures how evoked activity that originates at a given region of the cortex ripples through other brain regions. Our preliminary results on the Allen Institute for Brain Science's Mouse Connectivity Atlas show that a small number of brain regions (the Claustrum being at the top of the list) integrate all sensory information streams, suggesting that the cortex relies on an "hourglass architecture" to integrate and compress multi-sensory information before utilizing that lower-dimensionality representation in association regions and tasks.

III-19. Rare rewards drive enhanced dopamine responses and learning

Kathryn Rothenhoefer KMR0602@PITT.EDU
Aydin Alikaya AYA25@PITT.EDU
William Stauffer WRS@PITT.EDU

University of Pittsburgh

Dopamine neurons code reward prediction errors (RPEs), the differences between received and predicted values. However, it is not known how estimated (predicted) uncertainty affects value updating or the neural coding of RPEs. To investigate this, we created two sets of rewards with the same expected value but different 'shapes.' One reward set consisted of three reward sizes that were delivered with equal probability (p = 1/3, 1/3, 1/3). The other reward set consisted of the same three reward sizes, but with the middle value more certain to be delivered (p = 2/15, 11/15, 2/15). Thus, the two sets coarsely simulated uniform and normal distributions, respectively. Distinct images predicted the reward sets. Monkeys made choices between one of the reward sets and a sure alternative option. At un-signaled intervals during a session, we shifted all the reward sizes in a reward set by a constant factor. This manipulation forced animals to learn the shifted mean in order to maximize reward value. The choice behavior indicated monkeys learned faster from RPEs generated by rare rewards drawn from the tails of normal-like distributions – compared to the more frequent, but same magnitude RPEs generated by rewards drawn from uniform-like distributions. Single-unit dopamine recordings showed that dopamine responses were magnified by infrequent positive and negative RPEs generated in the normal-like distribution, compared to the identical positive and negative RPEs generates with greater frequency from the uniform-like distribution. Crucially, we observed bi-directional modulations in the same neurons and therefore the modulations in the dopamine responses could not be explained by different subjective values assigned to the predictive stimuli. Thus, rare rewards enhance the dopamine RPE responses, even when the conventional RPE is identical. This neuronal effect corresponded to the faster behavioral learning and together these results indicate that reward updating is modulated by estimated uncertainty.

III-20. Resolving the neural mechanisms of reinforcement learning through new behavioral technologies

Jeffrey Markowitz[1] JEFFREY_MARKOWITZ@HMS.HARVARD.EDU
Winthrop Gillis[1] WGILLIS@G.HARVARD.EDU
Julie Murmann[1] MURMANN.J@HUSKY.NEU.EDU
Scott Linderman[2] SCOTT.LINDERMAN@COLUMBIA.EDU
Bernardo Sabatini[1] BERNARDO_SABATINI@HMS.HARVARD.EDU
Sandeep Datta[1] SRDATTA@HMS.HARVARD.EDU

[1]Harvard University
[2]Columbia University

The ability to learn through experimentation paired with reward (i.e. reinforcement learning) is central to how both organisms and machines learn which action to perform on a moment to moment basis (i.e. action selection). However, the neural mechanisms underlying the reinforcement of action selection remain unknown. A key technical limitation is the way we study behavior in the lab. Typically, animals are trained to repeat a single, simple action (e.g. reach for a food pellet) through weeks of repetition. This approach has dominated neuroscience for decades, perhaps since it provides a simple, quantitative readout of behavior, usually a single number (e.g. reaction time). However, this is drastically different from how animals behave in the wild, where they learn to flexibly deploy any number of complex actions–rearing, running, hiding, and freezing–in response to external cues. Unfortunately, there are currently no readily-available technologies for evaluating this in the mouse. To study reinforcement learning and action selection in a more naturalistic setting, we aim to track and quantify a freely behaving animal's 3D movements in real-time with minimal latency in order to drive reward. Here, we describe the development of a fully automated system for real-time, accurate, high-resolution detection of repeatable mouse behaviors, which is based on a recently developed platform for the unsupervised identification of brief behavioral motifs (i.e. syllables), MoSeq (Motion Sequencing). In preliminary tests of our system, which we call rt-MoSeq (real-time Motion Sequencing), syllable-contingent optogenetic stimulation of neurons in the ventral tegmental area (VTA) is sufficient to reinforce the probability of expressing experimenter defined behavioral syllables with high specificity, both validating the performance of our closed-loop system and enabling a new class of experiments to study how reinforcement learning circuits sculpt spontaneous action selection.

III-21. Macaque V1 population encodes trial-by-trial sensory uncertainty as likelihood function

Edgar Walker[1] EYWALKER@BCM.EDU
Andreas Tolias[1] ASTOLIAS@BCM.EDU
R James Cotton[2] RJCOTTON@BCM.EDU
Wei Ji Ma[3] WEIJIMA@NYU.EDU

[1]Baylor College of Medicine
[2]Shirley Ryan AbilityLab
[3]New York University

For more than a century, Bayesian-inspired models have been used to explain human and animal behavior, suggesting that organisms represent the uncertainty associated with sensory variables. Nevertheless, the neural code of uncertainty remains elusive. It has been hypothesized that a critical feature of the neural code of uncertainty, implemented throughout the sensory processing chain in the neocortex, is that the same neurons that encode a specific world state variable (e.g. stimulus orientation in V1) also encode the uncertainty about those variables. According to the probabilistic population coding (PPC) hypothesis, inference in the brain is performed by inverting a generative model of neural population activity. Specifically, according to PPC, a neural population encodes sensory uncertainty in the form of the sensory likelihood function – the probability of observing a given pattern of neural activity across hypothesized stimulus values.

Here, we studied the neural code of uncertainty by simultaneously recording the activity of V1 cortical populations while monkeys performed a visual classification task in which the trial-by-trial uncertainty information is beneficial to the animal. To decode the trial-by-trial likelihood functions from the V1 population responses, we developed a novel technique based on deep learning. Critically, we performed all analyses conditioned on the contrast – an overt driver of uncertainty – and performed further orientation-conditioned analyses to isolate the effect of random fluctuations in the decoded likelihood function on behavior. We found that the likelihood function from the trial-to-trial population activity predicted decisions better than using only a decoded point estimate of the orientation. Therefore, we provide the first evidence that in perceptual decision-making, the same cortical population that encodes a point estimate of a sensory variable also encodes its trial-by-trial sensory uncertainty information in the form of likelihood functions mediating perceptual decisions, consistent with the theory of PPC.

III-22. Different tasks engage distinct populations in mouse parietal cortex

Julie Lee JULIE.LEE.15@UCL.AC.UK
Michael Krumin M.KRUMIN@UCL.AC.UK
Kenneth D Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

A challenge for the brain is to represent a vast variety of potential situations and act accordingly. Posterior parietal cortex (PPC) has been shown to be active in tasks probing decision-making, navigation, reward, and motor planning. Are the same PPC neurons involved in all these diverse functions, and if so, how are these different demands encoded across the population? To address these questions, we used two-photon calcium imaging to record from the same population of PPC neurons while mice performed two visual detection tasks. In one task, the mouse turned a steering wheel to report whether a visual grating was on the left or right side. In the other task, the mouse ran on a ball to navigate a virtual T-maze and turned to report the location of a grating on the left or right wall. Mice were able to perform both tasks on the same day. Although both tasks involved vision, decision, movement and reward, many neurons were selectively active in only either task. Task selectivity was consistent across days and was spatially intermingled in PPC. Modulation of neural activity by running could not fully account for task selectivity. Instead, selectivity was related to context more generally; selectivity in either task context (ball or wheel apparatus) was preserved even in the absence of a task; that is, the same neurons were active in passive conditions on the ball or wheel apparatus as in the respective task. Finally, we related each neuron's modulation by task events to task selectivity, including fitting models of task-relevant parameters. We conclude that PPC seems to allocate distinct, but overlapping, sets of PPC neurons according to the behavioral apparatus employed in each task. Within each task, task selectivity of individual neurons may be governed by the multi-dimensional characteristics specific to each context.

III-23. Neural correlates of strategic cooperation in rhesus macaques

Seth Madlon-Kay SETH.MADLONKAY@GMAIL.COM
Weisong Ong WEISONG.O@GMAIL.COM
Michael Platt MPLATT@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Competitive opponents today might be cooperative allies tomorrow. The ability to strategically play mixed-motive games is important for animals that have to navigate complex social environments, which requires understanding the motivations and strategies of others, and how their actions may change across different situations. To study how the brain learns when and how to cooperate with others, we devised a variant of the "chicken" game for that encouraged players to alternate between competitive and cooperative behavior and trained rhesus macaques to

play the game in pairs.

To determine whether macaques learnt and acted upon the strategies of their opponents, we constructed a decision model in which the agent, M1, uses an internal model of his opponent, M2, to predict M2's action on each trial. M1 updates his model of M2 using a strategic prediction error, the mismatch between his expectation about M2's actions and what M2 did. This model outperformed simpler models that do not involve learning the strategies of other, namely a reinforcement learning model, heuristics such as tit-for-tat, and models where the animal's choice depended only on the payoffs (mean decrease in AIC = 774).

We collected single-unit data from two brain areas: one linked to vicarious reward (ACCg), and another linked to strategizing (mSTS). During reward delivery, firing rates of 38% of mSTS neurons were modulated by whether or not the animals had interacted cooperatively, compared to only 20% in ACCg. Roughly 20% of neurons in both areas carried information about the opponent's predicted strategy.

Our behavior results indicate that rhesus macaques predict and learn about the strategies of others. While ACCg and mSTS neurons play a role in the integration of social cues, actions, and outcomes, mSTS neurons appear to be selectively signal cooperatively obtained rewards, suggesting an important role in guiding strategic social decisions.

III-24. Temporal spiking patterns in somatosensory cortex carry texture information

Katie Long KTHLONG@UCHICAGO.EDU
Justin Lieber JUSTINLIEBER@GMAIL.COM
Sliman Bensmaia SLIMAN@UCHICAGO.EDU

University of Chicago

We have an exquisite sensitivity to the microstructure and material properties of surfaces, allowing us to differentiate satin from silk and detect differences in surface features across six orders of magnitude, from tens of nanometers to tens of millimeters. In the peripheral nerves, two separate mechanisms convey information about texture: coarse features are encoded in the spatial pattern of activation of slowly adapting (SA1) fibers, and information about fine features is represented in the precise spike timing of rapidly adapting (RA) and Pacinian corpuscle (PC) fibers. Having previously shown that this precise temporal patterning is informative about texture at the periphery1, we examined whether these temporal patterns are still present in somatosensory cortex and, if so, the extent to which they carry texture information. To this end, we scanned a diverse set of everyday textures at a controlled speed and contact force across the fingertip of awake macaques while collecting single unit recordings from somatosensory cortex2. We then assessed the degree to which neuronal responses convey information about texture identity. To this end, we classified textures based on cortical spiking patterns at various temporal resolutions, from millisecond precision to spike count. We found that precise temporal patterning carries texture information, and that the optimal temporal resolution varies across cortical cells and depends on the submodality composition of their peripheral inputs (SA1, RA, PC). At the population level, a combination of rate and timing is more informative than each neural code in isolation. We conclude that temporal spiking patterns do carry information about texture in somatosensory cortex and probably reflect a progressive conversion from temporal to rate code that is not yet complete at this stage along the neuraxis.

III-25. Representation of sensory variables in recurrent spiking networks with realistic connection probabilities and synaptic dynamics

Mirjana Maras[1] MARAS.MIRJANA@GMAIL.COM
Christian Machens[2] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Sophie Deneve[1] SOPHIE.DENEVE@GMAIL.COM

[1]Ecole Normale Superieure
[2]Champalimaud Foundation

Representation of sensory variables in biological spiking networks is limited by the cost of firing action potentials, synaptic dynamics, and the fixed volume that can be filled by synaptic connections. Under the constraints of biologically observed connection probabilities and synaptic dynamics, we ask whether a spiking neural network can learn to represent sensory variables accurately while firing a minimal number of spikes. In contrast to the predictive spike-coding framework [Boerlin et al, PLoS CB, 2013], in which the optimal signal representation is achieved by a network with symmetric all-to-all recurrent connections and instantaneous interactions, we address optimality in the context of sparser connections and realistic synaptic dynamics. In the spike-coding framework, the optimal network reconstructs sensory variables with a minimal number of spikes for a given level of precision, and outperforms the equivalent rate network with matched firing rates by an order of magnitude. This framework challenges the common assumption that neural populations represent signals via rate codes, i.e., firing rates extracted from noisy spike trains. To address optimality in more biologically constrained networks, we extend the local balance restoring learning rule [Brendel et al, 2017, arXiv] to include a sparseness constraint on the recurrent connections. We implement the learning rule with instantaneous synapses, as well as with realistic synaptic dynamics. We find that the synaptic strength decreases with lower connection probability, and longer synaptic delays. As expected, the accuracy of the reconstructed signal for a given number of spikes degrades as we diverge from the optimal network. However, the reconstruction error continues to be smaller than that of the equivalent rate network. Thus, we show that a spiking network can learn to represent sensory variables more optimally than a rate code under the constraints of biologically plausible connection probability and realistic synaptic delays.

III-26. Theta-phase dynamics of prospective position coding in grid cells

Richard Gardner RICHARD.GARDNER@NTNU.NO
Abraham Vollan ABRAHAM.Z.VOLLAN@NTNU.NO
May-Britt Moser MAY-BRITT.MOSER@NTNU.NO
Edvard I Moser EDVARD.MOSER@NTNU.NO

Kavli Institute for Systems Neuroscience and Centre for Neural Computation (NTNU)

Theta-phase precession is expressed by place-signalling neurons across the hippocampal-entorhinal circuit. One hypothesis postulates that each theta cycle represents a stereotyped sequence in which the cognitive map initially signals the animal's current location before shifting this representation smoothly forwards (Sanders et al., 2015).

According to this hypothesis, the unknown dynamic offset between the true and signalled location could profoundly confound studies of spatial coding. We were therefore motivated to develop a statistical approach that could measure dynamic offsetting effects, and hence yield a more faithful estimate of position tuning.

We used the established Poisson GLM framework to model the spiking of medial entorhinal grid cells (n = 793 cells, 14 animals) as a function of the animal's 2D position. From this basic model we created a modified "shift model", by adding a pre-processing step which parametrically shifted the position coordinates. Specifically, each coordinate was translated along a line parallel to the animal's current head-direction. The shift quantity for each position coordinate was modelled as a function of the current theta phase; this function was provided by a curve fitted by the model.

Across cells, the shift model reliably improved on the basic model's cross-validated fit quality ($p < 0.1$ x 10-36) and yielded sharper receptive fields. Furthermore, the curve relating shift quantity to theta phase was consistently aligned across cells, supporting the hypothesis that the theta rhythm coordinates prospective shifting across the hippocampal-entorhinal network.

We also searched for ensemble-level manifestations of prospective position coding. By decoding position from 239 simultaneously recorded medial entorhinal neurons, we observed a theta-entrained forward-sweeping pattern in the decoding error (see figure), which resembled previously described CA1 theta sequences (Wikenheiser and Redish, 2015). We are now investigating the relationship between theta-paced sequences in MEC and CA1, to elucidate the mechanistic origin of anticipatory position signalling.

III-27. Representational similarity analysis reveals population coding complexities in mouse visual cortex

Jianghong Shi[1] ZJJXSJH@GMAIL.COM
Gabriel Ocker[2] GABEO@ALLENINSTITUTE.ORG
Michael Oliver[2] MICHAELO@ALLENINSTITUTE.ORG
Nicholas Cain[2] NICHOLASC@ALLENINSTITUTE.ORG
Peter Ledochowitsch[2] PETERL@ALLENINSTITUTE.ORG
Daniel Millman[2] DANIELM@ALLENINSTITUTE.ORG
Saskia de Vries[2] SASKIAD@ALLENINSTITUTE.ORG
Eric Shea-Brown[1] ETSB@UW.EDU
Michael A Buice[2] MICHAELBU@ALLENINSTITUTE.ORG

[1]University of Washington
[2]Allen Institute for Brain Science

Recently, researchers have found that features from the last hidden layer of goal-driven deep artificial neural networks (DNNs) can predict primates' inferior temporal cortex (IT) neurons' responses quite accurately, and far more so than classical models. The success of this approach lies in our relatively good understanding of primates' hierarchical visual pathways. In contrast to primates and despite recent insights from anatomical analysis, we know little about the functional architecture of the mouse visual cortex, and models for predicting single-neuron responses remain quite inaccurate for the majority of the neurons recorded here.

In this work, we bring together the framework of DNNs with the comprehensive survey of the Allen Brain Observatory to ask two basic questions about visual representations in the mouse visual cortex. First, how complex are these representations, in terms of the number of nonlinear transformations that are computed in forming them? Second, is there a clear functional hierarchy in mouse visual cortex regions? We investigate the functional complexity of mouse visual regions by comparing the representations of the mouse visual cortex with pretrained DNNs. By applying two metrics that depend on distinct statistical assumptions, we consistently found that mouse visual cortical regions produce representations that are more similar to middle layers of DNNs, than the beginning or final layers. In addition, the "most similar" layers in DNNs have much higher similarity values than classical 2D Gabor filter models. This suggests a high level of complexity for computations in mouse visual cortex, consistent with multiple nonlinear processing stages. Moreover, all 6 visual cortex areas are most similar to middle DNN layers, suggesting a more parallel or recurrent architecture rather than simple feed-forward hierarchy in the mouse visual pathway.

III-28. The latent structure of signal and noise in multi-region neural recordings

Stephen Keeley[1] STEPHENLKEELEY@GMAIL.COM
Mikio Aoi[1] MAOI@PRINCETON.EDU
Yiyi Yu[2] YIYIY@UCSB.EDU
Spencer Smith[2] SLS@UCSB.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]University of California, Santa Barbara

A central problem in systems neuroscience is characterizing how neural populations in interconnected brain areas work in concert to process sensory input. Factor analytic models are a powerful tool for identifying structure in high-dimensional neural data, but they typically are not used to distinguish signal from noise, or to characterize how signal and noise are represented across neural populations. Here we present an extension of Gaussian process factor analysis to address these questions. In particular, we propose a multi-region Poisson noise Gaussian Process Factor Analysis (P-GPFA) model that dissociates signal and noise components of neural spike trains in multiple simultaneously recorded brain areas. We show that the model successfully captures both the stimulus-locked PSTH and trial-by-trial variability in recordings from rodent visual cortex, and outperforms models that include either only trial-averaged signal or only trial-varying noise latents. We then use the model to assess noise structure across multiple simultaneously recorded visual cortical regions. We compare models of shared and independent noise across these regions and find that shared noise-models out perform models of independent cortical noise. This suggests that trial-by-trial variability is common across cortical regions, and neural noise is not a predominantly local dynamic feature. We learn the parameters of our P-GPFA model using a versatile black box variational inference approach. Our model is a general and flexible tool that can characterize latent structure of both signal and noise in neural recordings.

III-29. The anterior cingulate cortex projection to the dorsomedial striatum regulates outcome-dependent motivation

Julia Cox JULIAMC@PRINCETON.EDU
Sharon Ornelas SOG@EXCHANGE.PRINCETON.EDU
Ellia Miller ELLIAM@EXCHANGE.PRINCETON.EDU
Ilana Witten IWITTEN@PRINCETON.EDU

Princeton University

Decision-making in a complex and changing environment requires not only selection of appropriate actions, but also regulation of the rate at which to engage in those actions. The dorsomedial striatum (DMS) has been implicated in both of these aspects of decision-making, but it is unknown how specific inputs to the DMS contribute to each aspect. One of the major cortical inputs to the DMS is the anterior cingulate cortex (ACC), a region that encodes decision-related information and is thought to be especially important for error evaluation. To investigate how ACC inputs to the DMS contribute to action selection and behavioral pacing, we trained mice to perform a self-paced probabilistic reinforcement learning task in which choices and trial initiation latencies are modulated by previous outcome. Cellular resolution imaging revealed that the activity of many ACC neurons projecting to the DMS (ACC-DMS) was modulated by both choice and outcome, with more neurons modulated by unrewarded than rewarded outcomes. Inhibition of ACC-DMS neurons decreased latencies to initiate the next trial, specifically following unrewarded outcomes. This suggests that the ACC-DMS projection is involved in mediating the decreased motivation that occurs following unrewarded trials. Finally, to understand the downstream effects of ACC-DMS activity, we asked how the two major cell-types within the DMS—D1R and D2R medium spiny neurons (MSNs)—contribute to this behavior by selectively inhibiting them. Transient inhibition of D2R MSNs, but not D1R MSNs, decreased trial initiation latency following unrewarded trials. Thus, ACC-DMS may contribute to

outcome-dependent regulation of motivation by acting on D2R neurons in the DMS.

III-30. Accurate estimation of neural population dynamics without spike sorting

Subhaneil Lahiri[1] SULAHIRI@STANFORD.EDU
Eric Trautmann[1] ETRAUTMANN@GMAIL.COM
Daniel O'Shea[1] DJOSHEA@STANFORD.EDU
Saurabh Vyas[1] SMVYAS@STANFORD.EDU
Xulu Sun[1] XULU2@STANFORD.EDU
Surya Ganguli[1] SGANGULI@STANFORD.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU
Sergey Stavisky[1] SERGEY.STAVISKY@STANFORD.EDU
Katherine C Ames[2] KCA2120@COLUMBIA.EDU
Matthew T Kaufman[3] ANTIMATT@GMAIL.COM
Stephen I Ryu[1] SEOULMAN@STANFORD.EDU

[1]Stanford University
[2]Columbia University
[3]Cold Spring Harbor Laboratory

A central goal of systems neuroscience is to relate an organism's neural activity to behavior. Such efforts often begin by reducing the dimensionality of data to extract dominant neural population dynamics that could underlie task performance. A major hurdle to these analyses involves spike sorting, or extracting individual units from mixed activity on recording electrodes. This hurdle becomes more severe as the number of recorded neurons increases. Here, we investigate whether spike sorting itself may even be a necessary first step to accurately estimating neural population dynamics.

We re-analyzed data from three previous electrophysiological studies involving primate motor cortical dynamics of hundreds of neurons during reaching behaviors, as well as conducted a new study involving neuropixel probe measurements. In all cases, we found that both neural population dynamics and the scientific conclusions reached are quite similar using multi-unit threshold crossings instead of sorted neurons. Moreover we developed a quantitative theory, based on random projections, to explain the finding that mixing individual neurons on the same electrode does not significantly impair the estimation accuracy of neural population dynamics. Our theory predicts quantitative scaling laws for the estimation error: it should scale inversely with the number of electrodes, and logarithmically with the length, number and curvature of neural population trajectories. We successfully verified these theoretically predicted scaling laws in all four experiments.

Overall, our combined theory and experiment provides a conceptual framework, based on random projection theory, to uncover experimental regimes in which spike-sorting need not be necessary for understanding neural population dynamics. In terms of future impact, these findings may unlock existing data for new analyses without time-consuming or error-prone sorting, inform the design and clinical use of new wireless, low power, limited bandwidth electrode arrays, and enable new science with arrays that do not afford high quality spike sorting.

III-31. Temporal dynamics of inter-area neuronal population interactions

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

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

Most brain functions rely on activity distributed across brain areas. Yet, little is known about how information is selectively routed across neuronal populations in distinct brain areas. Indeed, our understanding of inter-area interactions largely relies on the measurement of pairwise relations (e.g., spike count correlation or spike-LFP/LFP-LFP coherence). Here, we employed dimensionality reduction to study the interaction of neural populations across two brain areas. We simultaneously recorded from neuronal populations in the output layers of primary visual cortex (V1) and in the input layers of area V2 of 3 anesthetized monkeys, and in areas V1 and V4 of an awake monkey, while animals viewed either oriented gratings (evoked activity) or a blank screen (spontaneous activity). We used canonical correlation analysis (CCA) to characterize moment-by-moment inter-area interactions under these two conditions, at multiple time scales. We found that population-level correlations between areas were higher for spontaneous activity than for evoked activity. Inter-area interactions showed robust dynamics, with interactions changing from being feedforward dominated early in the stimulus period to feedback dominated during spontaneous activity, as suggested by prior work based on PSTHs. Importantly, we found that the dimensions of the V1 and V2/V4 population activity that were involved in the V1-V2/V4 interactions for spontaneous activity were distinct from those involved for evoked activity. These results suggest that the patterns of activity relayed across areas during "feedforward" interactions are distinct from those involved during "feedback" interactions. More broadly, this work demonstrates that the dynamic nature of inter-areal communication can be captured by population-based analyses of neuronal responses recorded simultaneously in multiple areas.

III-32. Functional circuits for spatial choice in the superior colliculus

Jaclyn Essig JACLYN.ESSIG@UCDENVER.EDU
Gidon Felsen GIDON.FELSEN@UCDENVER.EDU

University of Colorado

Decision making is a fundamental process of the nervous system for generating goal-directed behaviors. The midbrain superior colliculus (SC) contributes to sensorimotor decision making by integrating cortical and subcortical inputs to guide orienting movements of the eyes, head, and body towards spatial goals. The SC is topographically organized and encodes for specific regions in retinotopic space, however the underlying circuitry for how the SC selects where to orient is unknown. Multiple models of excitatory/inhibitory interactions have been proposed to describe SC function, but these are based on cellular anatomy, ex-vivo slice physiology, and in-vivo recordings in the absence of behavior, or on recordings during behavior from unknown cell types. Here, we record and manipulate the activity of GABAergic neurons in mice performing a spatial-choice task1,2 to determine the functional role of inhibition during spatial choice. We train mice to select a left or right reward port based on a binary odor mixture (Fig. 1A). Importantly, after odor delivery, mice wait for a 'go tone' before orienting to the reward port, giving us access to neural activity during the decision (i.e., the "delay epoch") (Fig. 1B). We hypothesized that GABAergic neurons would shape spatial choice locally by inhibiting SC motor output neurons promoting contralateral choice, and therefore predicted that these cells would be most active before an ipsilateral choice. However, optogenetic identification (i.e., "optotagging") and activation of channelrhodopsin-expressing GABAergic neurons revealed that GABAergic neurons are active before contralateral choices and driving their activity during the 'delay

epoch' biases mice to select the contralateral port. These findings support a role for long-range inhibitory interactions in the SC. We are incorporating these data into a bump attractor model as a framework for understanding how the dynamics of excitation and inhibition give rise to nodes of activity in the SC underlying spatial choice.

III-33. Wearable non-invasive human neural interface with action potential resolution

Patrick Kaifosh PK@CTRL-LABS.COM
Alexandre Barachant ALEX@CTRL-LABS.COM
Curtis Barbre CURTIS@CTRL-LABS.COM
Nathan Danielson NATHAN@CTRL-LABS.COM
Alan Du ALAN@CTRL-LABS.COM
Ning Guo NING@CTRL-LABS.COM
Carlos Hernandez CARLOS@CTRL-LABS.COM
Nadine Hussami NADINE@CTRL-LABS.COM
Peter Li PETER@CTRL-LABS.COM
Michael Mandel MIM@CTRL-LABS.COM
Anthony Moschella ANTHONY@CTRL-LABS.COM
Thomas Reardon REARDON@CTRL-LABS.COM
Jonathan Reid JON@CTRL-LABS.COM
Ran Rubin RAN@CTRL-LABS.COM
Jeffrey Seely JEFF@CTRL-LABS.COM
Zhuo Wang JIMMY@CTRL-LABS.COM
Akshay Yembarwar AKSHAY@CTRL-LABS.COM

CTRL-labs

As the nervous system's evolved output, spinal motor neuron activity is from an evolutionary perspective a natural source of signals for a neural interface. Furthermore, the amplification of these signals by muscle fibers allows them to be measured non-invasively with surface electromyography (sEMG). We have developed a novel wearable wireless system that records state-of-the-art sEMG signals from the human forearm with dry electrodes and without the need for shaving or other skin preparation. Using this system, we demonstrate real-time detection and identification of individual motor unit action potentials, each of which corresponds to an action potential of an individual spinal motor neuron. This ability to monitor spiking activity of individual neurons sets sEMG apart from other non-invasive neural recording paradigms such as electroencephalography, functional magnetic resonance imaging, and near infrared spectroscopy. From the recorded sEMG signals, we also compute real-time predictions of joint angles, muscle tensions, and forces of the wrist and hand. The wireless and unobtrusive form factor of this recording system allows for long-duration monitoring of human neuromotor activity suitable not only for research and clinical applications, but also for real-time control in terms of motor unit action potentials, aggregate sEMG signals, and/or estimates of joint angles, muscle tensions, and forces. Relative to traditional human-computer interfaces, neuromotor interfaces have the potential to increase bandwidth by eliminating information loss as neural signals are converted to muscle tensions and then to input device signals. Meaningful reductions in latency are also achievable because sEMG signals precede forces and movement by tens of milliseconds. As a non-invasive neural interface with action potential resolution and the capability to augment human capacity for real-time control, this system expands the applicability of neural interface technology beyond research and clinical domains.

III-34. Replay for mitigating catastrophic forgetting

David Rolnick[1] DROLNICK@SEAS.UPENN.EDU
Arun Ahuja[2] ARAHUJA@GOOGLE.COM
Jonathan Schwarz[2] SCHWARZJN@GOOGLE.COM
Timothy P Lillicrap[2] COUNTZERO@GOOGLE.COM
Greg Wayne[2] GREGWAYNE@GOOGLE.COM

[1]University of Pennsylvania
[2]DeepMind

Continual learning is the problem of learning new tasks or knowledge while protecting old knowledge and ideally generalizing from old experience to learn new tasks faster. A significant hurdle to continual learning is posed by catastrophic forgetting, in which new knowledge supplants previously acquired knowledge. One of the leading hypotheses for how catastrophic forgetting is handled by the brain is hippocampal replay and consolidation. Here we show that a simple replay method greatly reduces catastrophic forgetting in artificial neural networks during reinforcement learning. In fact, the replay method performs as well as, or better than, state-of-the-art deep learning techniques for mitigating forgetting, despite being significantly less complicated and not requiring any knowledge of the individual tasks being learned.

III-35. Nonlinear variational inference for neural data

Daniel Hernandez Diaz[1] DH2832@COLUMBIA.EDU
Antonio Khalil Moretti[1] AMORETTI@CS.COLUMBIA.EDU
Ziqiang Wei[2] WEIZ@JANELIA.HHMI.ORG
Shreya Saxena[1] SS5513@COLUMBIA.EDU
John Cunningham[1] JPCUNNI@GMAIL.COM
Liam Paninski[1] LIAM.PANINSKI@GMAIL.COM

[1]Columbia University
[2]HHMI Janelia Research Campus

Modern experiments, such as multi-electrode arrays or calcium imaging, often produce high-dimensional data that represent a challenge for data analysis. Various techniques for dimensionality reduction and, in particular, latent variable models have been widely applied. At the other extreme, there is a big body of electrophysiological data coming from voltage measurements in single cells. In this setting it is understood that the underlying dynamics is in fact highly nonlinear and multidimensional. From 1D recordings, the task is to approximately recover the complete hidden latent space paths and dynamics.

In this work we propose a novel variational inference framework for the modeling of time series, Variational Inference for Nonlinear Dynamics (VIND), that is able to uncover nonlinear observation and transition functions from sequential data. The framework includes a structured approximate posterior and an algorithm that relies on the fixed-point iteration method to find the best estimate for latent trajectories. We apply VIND to multielectrode, single-cell and calcium datasets and show that it is able to accurately infer the underlying dynamics of these systems, in some cases substantially outperforming state-of-the-art methods.

III-36. Accurate and adaptive neural recognition in dynamical environments

Li K Wenliang KEVINWLI@OUTLOOK.COM
Eszter Vertes ESZTER@GATSBY.UCL.AC.UK
Maneesh Sahani MANEESH@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, University College London

Survival often depends on the ability to accurately track and learn the behaviour of a latent feature in the world—for example, the movement of a predator or prey—from noisy and incomplete sensory data. The statistical computations that would guarantee accuracy are well understood. As new noisy sensory information arrive, an agent or organism must integrate this information with its current belief about the latent quantity to form an updated representation. The form of this update depends on an internal model of how the quantity being tracked is likely to change in time. In general this model will itself be uncertain, and thus must be adapted on the basis of the incoming sensory data (and the updated beliefs inferred therefrom) to ensure consistency with external world. In all but the most trivial cases, accurate internal models and beliefs must be probabilistic, and cannot be summarised by the evolution of simple point estimates, such as the mean, alone. Here, we propose a biologically plausible inference and learning algorithm that relies on a deterministic representations of probabilistic beliefs using the distributed distributional code (DDC; Vertes and Sahani, 2018). We demonstrate empirically that the flexibility of the DDC makes possible accurate inference and learning in nonlinear settings; that synaptic weights for computation can be trained using a biologically plausible delta rule; that no explicit distributional assumptions are required; that the statistics of the latent variables are readily accessible from the representation; and finally that recognition can adapt to new problems through learning a flexible generative model of the observations.

III-37. Manipulating a cognitive variable using temperature

Tiago Monteiro[1] TIAGO.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG
Margarida Pexirra[1] MARGARIDA.PEXIRRA@RESEARCH.FCHAMPALIMAUD.ORG
Filipe S Rodrigues[1,2] FILIPE.RODRIGUES@NEURO.FCHAMPALIMAUD.ORG
Bruno Filipe Pereira Cruz[3,3] BRUNO.CRUZ@NEURO.FCHAMPALIMAUD.ORG
Ana Goncalves[1] ANA.GONCALVES@NEURO.FCHAMPALIMAUD.ORG
Joe Paton[2] JOE.PATON@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]Champalimaud Research
[3]Champalimaud Foundation

Time, like space, is a fundamental dimension of animals' worlds. To behave adaptively, organisms must extract temporal structure from experience and construct temporally patterned behavior. Yet, the mechanisms for doing so are poorly understood. The striatum is the main input structure of the basal ganglia (BG) and has been implicated in several time-dependent functions such as reinforcement learning and timing behavior. Previously we manipulated and recorded the activity of striatal neurons while rats performed a duration categorization psychophysical task1. We found that simultaneously recorded neuronal ensembles could judge time as well as the animal and that striatal neurons were necessary for duration judgments, as muscimol infusions produced a clear decrease in temporal sensitivity. Lastly, we demonstrated that time was encoded in the speed of population dynamics: faster dynamics were correlated with longer duration judgments and vice-versa. To directly assess whether variability in the speed of striatal population dynamics causes variability in duration judgments, here we experimentally manipulated striatal temperature during task performance using a custom-made solid-state thermoelectric cooling (TEC) implant. Alterations in the temperature of neural tissue have been shown in a number of systems to systematically and selectively affect the speed of population dynamics2,3,4. Behavioral sessions were divided in fixed-time blocks: control blocks (in which the device was set to body temperature) were interspersed with manipulation blocks (in which the temperature change varied in sign and magnitude). Here we show that cooling or warming striatal tissue caused graded underestimation or overestimation of duration, respectively.

Critically, motor-related variables such as reaction and movement times were not affected. These data strongly support the hypothesis that dynamics in neural populations underlie duration estimation, and establish the dorsal striatum as a locus within neural circuitry where a decision variable related to elapsed time and independent of motor function may be manipulated.

III-38. Rapid embedding of low-dimensional dynamics in a recurrent neural network

Eli Pollock EPOLLOCK@MIT.EDU
Mehrdad Jazayeri MJAZ@MIT.EDU

Massachusetts Institute of Technology

In recent years, recurrent neural networks (RNNs) have become a popular tool for modeling the brain, and analysis of their dynamics has provided insight into how the brain might accomplish tasks. Using this framework, it is possible to relate neural computations to the evolution of network activity over a manifold. Recent theoretical work has investigated how low-rank structure in network connectivity influences network dynamics. However, if we wish to study the kind of connectivity that allows for movement over a particular manifold, it is useful to solve the inverse problem of creating RNNs whose activity is constrained to that manifold. In this abstract, we address that inverse problem and demonstrate a technique for rapidly creating RNNs whose intrinsic dynamics are fully defined in advance. We first demonstrate this technique with an RNN that moves over a ring attractor according to a desired drift function. We also show the generalization of this technique to three dimensions with a spherical attractor. Finally, we demonstrate how RNNs are limited in the kinds of dynamical landscapes they can traverse, which presents a puzzle for future work.

III-39. Hippocampal sequences and model-based planning in the rat

Sarah Jo Venditto[1] VENDITTO@PRINCETON.EDU
Carlos Brody[1] BRODY@PRINCETON.EDU
Kevin Miller[2] KJMD10@GMAIL.COM
Matthew Botvinick[3] BOTVINICK@GOOGLE.COM
Nathaniel Daw[1] NDAW@PRINCETON.EDU

[1]Princeton University
[2]University College London
[3]DeepMind

Humans and animals construct internal models of the world around them and use these models to guide behavior. Such model-based cognition is often referred to as “planning”, and its neural mechanisms remain poorly understood. Planning has been proposed to depend on hippocampal sequences, in which place cells “sweep out” trajectories through space while an animal is at rest (Foster & Knierim, 2012; Mattar & Daw, 2018). Research into the role of sequences in planning, and into the neural mechanisms of planning in general, has been hampered by a lack of tasks for animals which both demonstrably elicit planning behavior and are suitable for neural recordings. Recent work has lifted this limitation; advances from work with humans (Daw et al., 2011) has been adapted into a multi-step decision task for rats, showing that rats adopt a planning strategy which depends on neural activity in the dorsal hippocampus (Miller, Botvinick, & Brody, 2017). Here, we report the results of electrophysiological recordings made in dorsal hippocampus during planning behavior. We find that individual cells encode the states of the task, and that hippocampal sequences take place during sharp wave ripple events at the conclusion of rewarded trials but not omission trials. The content of these sequences reflects knowledge of the structure of the task, consistent with a role in model-based planning. We used a traditional Bayesian decoding approach to identify replay events during sharp wave ripples and characterize their relationship to the task on a trial-by-trial level. We

seek to understand whether these replay events support planning by encoding action-outcome representations that evolve with changing reward contingencies.

III-40. Power-law efficient codes provide general link between perceptual bias and discriminability

Michael Morais MJMORAIS@PRINCETON.EDU
Jonathan Pillow PILLOW@PRINCETON.EDU
Princeton University

Recent work in theoretical neuroscience has shown that efficient neural codes, which work to maximize the mutual information between stimuli and neural responses, give rise to a lawful relationship between perceptual bias and discriminability in psychophysical experiments (Wei & Stocker 2017). Here we generalize these results to show that the same law arises under a much larger family of optimal neural codes, which we call power-law efficient codes. These codes provide a unifying framework for understanding the relationship between perceptual bias and discriminability, and show how it depends on the allocation of neural resources. Specifically, we show that the same lawful relationship between bias and discriminability arises whenever Fisher information is allocated proportional to any power of the prior distribution. This family includes neural codes that are optimal for minimizing Lp error for any p, indicating that the lawful relationship observed in human psychophysical data does not require information-theoretically optimal neural codes. Furthermore, we derive the exact constant of proportionality governing the relationship between bias and discriminability for different choices of power law exponent q, which includes information-theoretic (q=2) as well as "discrimax" (q=1/2) neural codes, and different choices of decoder. As a bonus, our framework provides new insights into "anti-Bayesian" perceptual biases, in which percepts are biased away from the center of mass of the prior. We suggest that the bias-discriminability relation would be a signature of a generalized efficient coding principle, reinterpreted as a broader family of transformations that optimize the noise structure of neural codes.

III-41. Mechanisms of auditory evoked response generation revealed with bispectral pattern recovery

Christopher Kovach[1] CHRISTOPHER-KOVACH@UIOWA.EDU
Kirill Nourski[2] NOURSKI@GMAIL.COM
Phillip Gander[2] PHILLIP-GANDER@UIOWA.EDU
Hiroyuki Oya[2] HIROYUKI-OYA@UIOWA.EDU
Hiroto Kawasaki[2] HIROTO-KAWASAKI@UIOWA.EDU
Matthew Howard[2] MATTHEW-HOWARD@UIOWA.EDU
[1]University of Iowa
[2]University of Iowa Hospitals and Clinics

We present a new technique for isolating recurring transient waveforms of arbitrary shape in electrophysiological recordings and demonstrate its ability to blindly recover evoked responses in auditory cortex without prior information about stimulus timing. The technique exploits time-invariant information in the bispectrum to obtain filters matched to one or more independently recurring patterns within a signal under noise. We applied this approach to human intracranial recordings of responses during stimulation with acoustic click trains. We found averaged evoked responses in auditory cortex which arose not from the transient event-related emission (ERE) of a particular waveform, but through the transient modulation of a background of emitted waveforms. In particular, we note a role for transient event-related suppression, resulting in a negative image of the suppressed waveforms within ensemble averages (illustrated in Fig. 1). We thus identify a novel mechanism for averaged response generation, which departs from prior models that consider only ERE or oscillatory phase resetting. In addition, we observed

the predominant mechanism within auditory cortex to vary with location along Heschl's gyrus (Fig. 2). In the response to simple acoustic transients, ERE predominated within the posteromedial portion of the gyrus, whereas ERS characterizes the anterolateral portion. The latter finding was supported by negative correlations between the shape of the (suppressed) generating waveform and that of the averaged response and, separately, by the relative suppression of overall signal power at the time of the peak averaged response. We consider implications of these findings for the interpretation of ensemble averages and for the relationship between averaged evoked responses and background activity.

III-42. Unambiguous representations of ambiguous contexts in the hippocampus

Mark Plitt MARKPLITT@GMAIL.COM
Lisa Giocomo GIOCOMO@GMAIL.COM

Stanford University

The hippocampus is a crucial structure for contextual and spatial learning, and principal cells in this region show striking correlations with behavior. Place cells, for example, fire specifically when an animal occupies confined locations in space and change their firing properties with context. However, we know very little about how hippocampal populations behave when contextual cues are ambiguous. To address this question, we trained mice to distinguish two similar hallways in virtual reality while imaging calcium activity of large CA1 populations. Once the animals performed the task well, they were occasionally shown randomly chosen morphed environments which were a blend of the two hallways. Mice rapidly acquired this task (1 week), mitigating effects of overtraining and allowing imaging during initial stages of learning. During the first exposures to the hallways, low dimensional population dynamics were not informative about task variables; however, these dynamics rapidly reorganized to encode position and context on subsequent sessions. This coding strengthened as animals progressed to more difficult stages of the task—an effect likely driven by the dramatic increase in the proportion of place cells observed over the same period. Despite the continuous nature of the environmental cues, individual place cells show largely binary remapping on single trials in the morphed environments, firing in only one or none of their candidate place fields. The number of cells that remapped increased across the morphs. This results in a gradually decaying population vector correlation across morphs. In addition, a population decoding analysis found evidence for a pattern completion-separation mechanism where context was represented in a binary fashion even at first exposure to the morphed environments. Together these results suggest that rapidly organizing attractor dynamics at both the single cell and population level result in all-or-none representations of ambiguous contexts in the hippocampus.

III-43. Learning low-dimensional inputs for brain-machine interface control

Jorge Aurelio Menendez UCBPJAM@UCL.AC.UK
Peter Latham PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, University College London

A critical, but poorly understood, feature of the motor system is that it can learn new motor behaviors without forgetting old ones. A remarkable demonstration of this ability is given by brain-machine interfaces (BMI), which show that primates are capable of learning experimenter-imposed mappings from motor cortical activity to movement. Previous models of this phenomenon (Legenstein et al., 2010; Waernberg & Kumar, 2017) invoke synaptic plasticity in the local network. However, several observations are at odds with this hypothesis. In some cases, the correlation structure in the local population is largely conserved between BMI and manual control (Hwang et al., 2013) and between different BMI decoders (Golub et al., 2018), suggesting that the local circuitry has not undergone any changes. Furthermore, primates are able to remember and rapidly switch between previously

used decoders (Ganguly & Carmena, 2009) – something that local plasticity rules have considerable difficulty achieving. Motivated by these observations, here we propose an alternative hypothesis: learning is isolated to upstream commands driving the local motor cortical circuit. Under this hypothesis, efficient learning can be achieved by restricting the space of possible upstream commands to those used during natural movement, which has been previously suggested to underlie BMI control. We formalize this hypothesis as a control problem in which upstream commands drive a recurrent network model of motor cortex. The formalism provides a quantitative relationship between the experimenter-determined BMI decoder and the difficulty of learning the task; we show that this relationship can account for previous observations on learning different BMI decoders (Sadtler et al., 2014). Moreover, we argue that solving the control problem in this way makes efficient learning viable without computing gradients - a critical feature of BMI learning.

III-44. Interactions between decision bounds and metacognitive evidence accumulation

Tarryn Balsdon[1] TARRYN.BALSDON@GMAIL.COM
Pascal Mamassian[2] PASCAL.MAMASSIAN@ENS.FR
Valentin Wyart[1] VALENTIN.WYART@ENS.FR

[1]Ecole Normale Superieure
[2]PSL University

In order to resolve complex and ambiguous sensory input, the observer must accumulate multiple samples of sensory evidence. Accumulation-to-bound models have provided an adequate description of perceptual decisions in terms of behavior and neural correlates. However, these decisions are also accompanied by a feeling of confidence, and a consensus has not yet been reached on how best to computationally describe observers' metacognitive confidence judgments. In a two-session psychophysical experiment, human participants (N = 20) were asked to categorize the source of up to 40 noisy visual stimuli presented sequentially. In the first session, participants could end the presentation of stimuli at any point in the sequence to match target levels of decision accuracy. Bayesian accumulation-to-bound models were fitted to estimate participants' ability to adjust their decision bounds according to target levels of accuracy. In the second session, the same partici-pants first provided their decisions whenever they felt ready, which could be fitted in terms of a 'default' decision bound. They were then cued to provide decisions and confidence ratings at different points relative to this fitted bound. We found that participants' ability to adjust their decisions bounds in the first session predicted the sensitivity of their confidence ratings to the accuracy of their decisions in the second session. Despite this interrelation between decision-making and confidence, we further found that the accumulation of evidence for perceptual decisions and confidence could be dissociated at different points relative to the observers 'default' decision bound. Together, these findings offer a reconciliation of seemingly antagonistic effects through a confidence-regulated accumulation-to-bound process that controls decision-making even in the absence of any speed-accuracy trade-off.

III-45. A neuron type with variable receptive field properties is required for robust motion processing

Luis Giordano Ramos Traslosheros Lopez[1,2] L.RAMOS@ENI-G.DE
Marion Silies[3] M.SILIES@ENI-G.DE

[1]University of Goettingen
[2]European Neuroscience Institute
[3]European Neuroscience Institute, University Medical Center Gottingen

Moving objects are part of the world we experience, and the perception of visual motion is crucial for many organisms. The neural substrate for motion computation has been extensively studied in Drosophila. The fruit fly uses

two pathways to detect moving ON and OFF signals. These pathways split downstream of photoreceptors and connect through interneurons to the first direction-selective cells, T4 (ON) and T5 (OFF). Models for elementary motion detection rely on local correlations of two spatially restricted points.

Core circuits that process local motion cues have recently been worked out (1, 2). Interestingly, we identified a neuron with spatially wide receptive fields to be required for local motion detection. Blocking activity in this cell type (Tm9) abolished behavioral responses to OFF motion (3). Another study reported spatially narrow receptive fields. We set out to understand how a single cell type can display distinct receptive field properties, and how this contributes to its role in motion computation. We measured Tm9 receptive fields using a range of visual stimuli and in vivo two-photon calcium imaging. Our experiments reconcile the contradicting results and show that Tm9 wide receptive fields result from ON inhibition. At the same time, receptive fields are variable within this cell type. This variability is unique to Tm9 and does not occur in parallel OFF pathway neurons. We further dissected this pathway using optogenetics for functional connectivity mapping. We show that Tm9 receives inputs from numerous neuron types, including previously uncharacterized wide-field neurons. We are currently investigating how recruitment from a range of potential inputs relates to distinct physiological properties. We aim to understand how a variable synaptic partner choice ensures robust neuronal function. Our data suggest that the implementation of a robust computation requires the integration of a variable circuit element.

III-46. Joint neural-behavioral models of perceptual decision making

Matthew D. Golub[1] MGOLUB@STANFORD.EDU
Chandramouli Chandrasekaran[1] MOULI@STANFORD.EDU
William T Newsome[1] BNEWSOME@STANFORD.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU
David Sussillo[2,1] SUSSILLO@GOOGLE.COM

[1]Stanford University
[2]Google AI

An emerging hypothesis in systems neuroscience is that population-level dynamics implement the computations that map sensory cues to appropriate behaviors. If computation is indeed performed through neural dynamics, understanding the dynamical motifs expressed by a neural population could provide a rich description of how that population solves a task, flexibly generalizes across tasks, and learns to solve new tasks.

Existing approaches to identifying dynamical motifs can be largely categorized as either “task-modeling” or “neural data-modeling.” Task-modeling approaches have trained recurrent neural networks (RNNs) to solve a task, and then sought post-hoc correspondences between the RNN’s activity and neural activity recorded in animals performing that task. However, there may be many qualitatively different dynamical motifs that can solve a task, and the particular motif identified can depend dramatically on the model design and choice of hyperparameters. Neural data-modeling approaches have sought low-dimensional dynamical latent variables from which recorded neural activity can be reliably reconstructed. While these approaches provide powerful single-trial estimates of neural state, they lack an explicit link to the task, behavior and underlying computation.

Here we propose Goal-Oriented Learning of neural Dynamics (GOLD), a joint neural-behavioral modeling approach that overcomes the limitations of the aforementioned approaches while still enjoying the benefits of each. GOLD learns latent dynamics to generate single-trial behavior (i.e. task-modeling), constrained to be consistent with single-trial population recordings (i.e. neural data-modeling). Using behavior and population recordings from a perceptual decision-making task, we found that GOLD identifies latent dynamics that more closely resemble held-out single-trial neural activity compared to dynamics identified from task-modeling approaches. Furthermore, GOLD provides interrogable access to the set of dynamical motifs expressed in the data, with explicit links to those motifs’ correspondence to computation and behavior.

III-47. Universal statistical properties of the hippocampal code

John Briguglio BRIGUGLIOJ@JANELIA.HHMI.ORG
Albert Lee LEEA@JANELIA.HHMI.ORG
Jaesung Lee LEEJ13@JANELIA.HHMI.ORG
Sandro Romani ROMANIS@JANELIA.HHMI.ORG

HHMI Janelia Research Campus

Hippocampus is known to be important for spatial memory and navigation. Place fields are the neural correlates of spatial information observed during these tasks. Previous work (Rich et al, 2014) has shown large heterogeneities in the number of place fields observed across hippocampal CA1 cells, the nature of which becomes clear only in large environments. However, it remains unclear how these coding properties generalize across multiple conditions, such as different environments, extended periods of time, and the presence of salient environmental cues. To examine these properties, we use two-photon imaging to record from hundreds of CA1 cells as mice explore a variety of large, 1D VR environments with random- and fixed-reward locations over a timescale of several months. We find that CA1 cells approximately preserved their place-field number across environments and across long periods of time. Additionally, we find significant correlations between the place-field number during random-reward exploration and the number of reward-specific fields exhibited when reward location is fixed. To test whether these observations can arise from a similar underlying propensity for cells to respond to any of these conditions, we model the field formation as a stochastic process (spatially Poisson) with highly preserved underlying propensities for each cell. We find that this class of models explains many features of the data, including the joint distribution of place- and reward-responses. These findings suggest that cell-excitability is an intrinsic property of a cell that largely determines its activity pattern across different environments and over month-long timescales. Moreover, reward-specific responses can be accounted for by an approximately uniform increase across cells of their propensity for place responses. This statistical model of the hippocampal code opens the door for the study of storage, retrieval, and decoding of spatial information in more mechanistic descriptions of the hippocampal circuits.

III-48. Cortical dynamics drive the simultaneous growth of smooth maps and salt-and-pepper organization

Caleb Holt[1] CHOLT@UOREGON.EDU
Kenneth Miller[2] KEN@NEUROTHEORY.COLUMBIA.EDU
Yashar Ahmadian[1] YASHAR@UOREGON.EDU

[1]University of Oregon
[2]Columbia University

In many mammalian species, the primary visual cortex (V1) develops smooth maps for several receptive field features, such as preferred orientation or spatial frequency. Such maps may be beneficial in minimizing wiring lengths between neurons selective to similar features. Nevertheless, even in visual cortices with smooth maps, the receptive fields of nearby neurons show considerable heterogeneity. Correspondingly, some receptive field features are largely uncorrelated between nearby cells, and average signal correlations between nearby cells are near zero. Such random, "salt-and-pepper" organizations may in turn be advantageous in reducing the response redundancy of local V1 populations and increasing their information content. Thus a combination of smooth maps for some features, and salt-and-pepper organizations for others, may provide both the benefits of wiring length minimization and informational efficiency. Previous theoretical models have described the development of V1 feature selectivity and maps based on activity-dependent Hebbian plasticity. However, these models inevitably predict that V1 either develops smooth maps for all features (when long-range recurrent cortical excitation is strong) or salt-and-pepper organizations for all features (if cortical recurrent excitation is weak or short-range); they fail to account for the biologically observed co-presence. We propose that this failure is in part due to these models neglecting the intrinsic temporal dynamics of V1, and show that if properly considered, cortical interactions

at slow and fast time scales will couple to the slow and fast features of inputs to V1, respectively. This can lead to the development of smooth maps for slow input features and salt-and-pepper organizations for fast input features. In particular, by simulating one- and two-dimensional topographic models of development of plastic feedforward thalamocortical connections, we show that our framework can sustain both smooth maps and salt-and-pepper organizations, providing a more biologically plausible mechanism for receptive field feature development in V1.

III-49. Corticostriatal plasticity underlying learning and reversal of auditory – motor association in mice

Sanchari Ghosh SGHOSH@CSHL.EDU
Anthony Zador ZADOR@CSHL.EDU

Cold Spring Harbor Laboratory

Animals use complex sensory cues from their environment to make a variety of decisions which require sensory discrimination, decision-making and appropriate motor actions. The brain circuits underlying such behavior, and how they evolve during learning remain largely unexplained. Previous studies in rodents trained in an auditory discrimination task (‘tonecloud’ task) have established a critical role for the connections between auditory cortex and auditory striatum in performing such a behavior. It is also known that learning this task induces strong synaptic plasticity in this circuit in a manner that the memory trace of the learned association can be read out from acute slices containing these synapses. These findings suggest that even parts of the striatum receiving predominantly sensory inputs, might be involved in promoting contralateral movements. We tested this model in a reversal paradigm using the tonecloud task, where an animal initially trained to form a specific auditory-motor association was then forced to reverse its association to obtain reward successfully. We find that mice can indeed learn to reverse these associations, taking comparable training times and reaching similar performances in both training epochs. We then investigated the pattern of plasticity in this circuit following the reversal and found that reversing the association does not result in a simple reversal of the memory trace. In fact, we observe a strong persistence of the synaptic plasticity pattern that only reflects the initial association. Our results conform to the general belief that forming new memories don’t erase previous ones thereby allowing animals to form flexible associations between sensory inputs and motor outputs. However, these results suggest that sensory striatum might not be simply transducing sensory information into a contralateral motor output, and raises the question of how task-related context information is integrated in this circuit.

III-50. Identifying behavioral states and their sensorimotor transformations

Adam Calhoun ADAMJC@PRINCETON.EDU
Jonathan Pillow PILLOW@PRINCETON.EDU
Mala Murthy MMURTHY@PRINCETON.EDU

Princeton University

Animals need to be able to respond to the same stimuli when they are in different behavioral contexts or have differing internal needs. This source of behavioral variability is often ignored in behavioral studies due to our inability to quantitatively identify the internal state or behavioral context of an animal. To fill this need, we develop a novel unsupervised method to identify the internal state of an animal and apply it to the acoustic communication behaviors of Drosophila melanogaster. This behavior consists of discrete patterned sounds which they produce during their mating ritual and are dynamically shaped by feedback cues between the male and the female. Our unsupervised model combines hidden Markov models and generalized linear models to identify the hidden (latent) states of the animal and the timescales of interactions between the animals, enabling it to predict 94% of the moment-to-moment variation in song patterning decisions in Drosophila males (a 43% improvement over previous models). We identify three previously-unrecognized behavioral states that correspond to different sensorimotor

strategies. These three states are characterized by different relationships between motion and feedback cues and the outputs of the model (three different modes of song plus quiet). We find that activation of neurons previously identified to drive song production in males can in fact bias switching between behavioral states and hence distinct transformations of the same sensory input. Our results reveal how animals compose behavior from previously invisible internal states, a necessary step for quantitative descriptions of animal behavior that links environmental cues, internal needs, neuronal activity, and motor outputs.

III-51. Extracting universal algorithmic principles from large populations of recurrent networks

Niru Maheswaranathan[1] NIRUM@GOOGLE.COM
Alexander Williams[2] ALEX.H.WILLIA@GMAIL.COM
Matt Golub[2] GOLUB.MATTHEW@GMAIL.COM
Surya Ganguli[2] SGANGULI@STANFORD.EDU
David Sussillo[3,2] SUSSILLO@GOOGLE.COM

[1]Google Brain
[2]Stanford University
[3]Google AI

Many recent studies suggest that artificial neural networks learn to represent task variables in a similar fashion to neurobiological networks that solve analogous problems. These results are surprising because artificial and biological networks differ dramatically in obvious respects. An intriguing hypothesis is that similarities in network representations reflect universal algorithmic principles of network-based computation and learning. Evidence that such principles emerge despite low-level implementational differences would provide strong motivation for neuroscientists to study artificial networks alongside their biological counterparts. However, it is also possible that artificial and biological systems utilize similar task representations, but nonetheless implement qualitatively different algorithmic solutions. Systematically studying these questions in biological systems would be challenging given experimental constraints. Therefore, we trained thousands of artificial recurrent neural networks (RNNs) on a series of canonical tasks (cued sensory integration, discrete memorization, dynamics generation, and prediction), and used this diverse in silico population to search for universal features of computation. Interestingly, minor changes in RNN implementation induced large differences in network representations, as quantified by canonical correlation analysis (CCA), a common method for comparing artificial and biological network activity. However, using tools from graph theory and dynamical systems theory to both define and extract an algorithmic description of network computation, we found that networks with different implementation details and representational properties can exploit similar algorithmic principles to solve a task. Thus, while network representations can vary with implementation, algorithms may be more universal.

III-52. Acetaminophen reduces the sensitivity of macaque anterior cingulate gyrus neurons to the valence of decision outcomes

Km Sharika CATCHTHESHARK@GMAIL.COM
Michael Platt MPLATT@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Human imaging studies have implicated areas associated with the affective component of pain, namely, dorsal anterior cingulate cortex (dACC) and anterior insula (Rainville et al., 1997) in empathy for pain in others (Singer et al., 2004; Lamm et al., 2015). More recently, painkillers like acetaminophen - active ingredient in tylenol/ paracetamol – as well as placebo analgesia have been shown to reduce BOLD activity in ACC, ease self-reported pain and also diminish empathy for others in distress (Rutgens et al., 2015; DeWall et al., 2016; Mischkowski et al., 2016) suggesting that analgesics may interfere with the normal processing of affect associated with self and other outcomes. We recorded the electrophysiological activity of neurons in macaque ACC gyrus (ACCg) – an area that has been shown to be sensitive to others' positive outcomes (Chang et al., 2013) - to test 1) if these neurons also process negatively valenced information for self and other and if yes, 2) whether this activity is sensitive to oral acetaminophen. Two monkeys were trained to make choices between two targets associated with varying magnitudes of fluid. The fluid on offer was cued to be either pleasant (fruit juice) or unpleasant (diluted quinine) for either self or a recipient monkey sitting across the room. Monkeys rapidly learnt the cue associations and showed an overall preference for good tasting fluid for both self and other ($p<0.05$, n=2 pairs). Preliminary analysis suggests that acetaminophen not only reduced the monkey's behavioral sensitivity to negative outcomes in general, but also diminished the sensitivity of ACCg neurons to discriminate between positive and negative outcomes for both self and other. These results aim to expand our understanding of the neural circuitry underlying the processing of affective outcomes for self and others.

III-53. Differential encoding of information by coordinated neuronal ensembles in primary auditory cortex

Jermyn See JERMYN.SEE@UCSF.EDU
Craig Atencio CRAIGATENCIO@GMAIL.COM
Natsumi Homma NATSUMI.HOMMA@UCSF.EDU
Vikaas Sohal VIKAAS.SOHAL@UCSF.EDU
Christoph Schreiner CHRIS@PHY.UCSF.EDU

University of California, San Francisco

How individual neurons work together to encode sensory information and influence behavior remains one of the fundamental questions in sensory neuroscience. However, most studies of information processing in the primary auditory cortex (AI) involve either single-unit spectrotemporal receptive field (STRF) estimation or paired neuronal correlation analyses, assuming that AI neurons filter auditory information either as individual entities or pairs. Determining how AI encodes information will require an integrated approach that combines receptive field and multi-neuronal analyses.

We previously showed that coordinated neuronal ensembles (cNEs, i.e. groups of neurons with reliable synchronous activity) are stable functional constructs that may represent the basic unit of information processing in AI (See et al., 2018). To further assess the functional properties of cNEs in AI, we performed dense extracellular recordings in rat AI while presenting dynamic broadband stimuli, identified cNEs using dimensionality reduction techniques, and isolated cNE events, using them to assess spectrotemporal information processing. We found that there were different types of cNEs – some improved spectrotemporal information processing over that of member neurons, while others seem agnostic to the presented stimuli and may represent the convergence of top-down inputs. We also identified single neurons that participated in multiple cNEs and found that their spikes associated with one cNE had receptive field properties that were significantly different from those of their spikes associated with other cNEs.

These findings challenge the classical idea that AI neurons produce a homogeneous set of spikes that may be equally weighted to estimate a single STRF. Instead, AI neurons can have several receptive fields based on associations with different cNEs, with each cNE representing the convergence of thalamocortical, intracortical and top-down inputs into AI. Therefore, by considering the stimulus preferences associated with each cNE, we may gain a more comprehensive evaluation of information processing in AI.

III-54. Accumulated evidence inferred from neural activity accurately predicts behavioral choice

Brian DePasquale DEPASQUALE@PRINCETON.EDU
Jonathan Pillow PILLOW@PRINCETON.EDU
Carlos Brody BRODY@PRINCETON.EDU

Princeton University

Accumulating evidence in service of decision making is a core cognitive function. Methods exist for inferring accumulated evidence models from subject choices but few studies quantify how well these models describe the neural processes that underlie those choices. We present a maximum likelihood based method for inferring models of accumulated evidence from neural activity. Our method postulates a one-dimensional latent variable model that describes the full distribution of accumulated evidence and relates this value to observed neural activity. We compute the temporal evolution of the distribution of accumulated evidence using the Fokker-Planck equation, which allows us to compute the likelihood with a single forward pass over the data. We apply our method to activity from three brain regions which reflect accumulation—posterior parietal cortex (PPC), frontal orienting fields (FOF) and striatum (STR)—in rats performing a pulsed-based accumulation task. Many neurons were well described by the model, supporting its use for describing behavioral choice. Consistent with accumulator models fit to choice data, neural-inferred models describe an accumulation process with a high decision commitment bound and in which accumulation noise is caused by uncertainty in the sensory stimulus, not diffusive noise. Whereas choice-inferred accumulation models exhibit near perfect accumulation or mild information loss due to leak, we find that accumulator models extracted from neural data exhibit positive feedback, consistent with neurons favoring early evidence. Despite these differences, our model predicts choices nearly as well as choice-inferred accumulator models. Our work introduces an efficient method for learning evidence accumulation models from neural activity and illustrates that this model class, historically used to describe behavioral choices, accurately describes the neural activity that subserves those choices. Additionally, our results hint at additional processing that must exist downstream of the neurally accumulated evidence so that choices are consistent with choice-inferred accumulator models.

III-55. Divergent strategies for learning in males and females

Sijin Chen CHEN5388@UMN.EDU
Becket Ebitz REBITZ@GMAIL.COM
Sylvia Bindas BINDA008@UMN.EDU
Benjamin Hayden BENHAYDEN@GMAIL.COM
Nicola Grissom NGRISSOM@UMN.EDU

University of Minnesota

While gender and sex differences in most behavioral outcomes are small, there is evidence to suggest more substantial divergence in the cognitive strategies preferentially used by males and females. Unobserved computational differences due to sex or gender could cloud any attempt to understand interindividual variability. To address this omission, we examined strategy selection in a large sample of both male and female mice performing a classic decision-making task: the two-armed bandit. In this task, animals adopt a variety of strategies, which evolve as they learn. This means that identical final levels of performance can be achieved through widely divergent strategic paths. Here, we quantified these strategic paths. We found that one of the major axes of interindividual variability in strategy was the sex of the animals. While males and females ended at the same performance level, females learned more rapidly than their male counterparts because the sexes differed by the strategy applied during learning. Female mice as a group adopted a unified, systematic approach which reduced the dimensionality of the decision-space early in learning. Conversely, males engaged in ever-changing strategies not only between males but within an individual male over multiple iterations of the task. These results suggest that similar levels of performance can be achieved through widely divergent approaches, within and between

subjects, and that sex is a significant factor governing strategy selection in decision making and learning. These results highlight the need to consider sex and gender influences on cognitive strategies in decision making and reinforcement learning.

III-56. Integration of visual and motor information in retrosplenial cortex

Lukas Fischer[1] LFF@MIT.EDU
Raul Mojica Soto-Albors[2] RMOJICA@MIT.EDU
Mark Harnett[1] HARNETT@MIT.EDU

[1]Massachusetts Institute of Technology
[2]University of Puerto Rico at Mayaguez

Environmental landmarks represent advantageous reference points for guiding navigation. Salient visuo-spatial cues anchor the mammalian head direction system and the hippocampal cognitive map. Yet the process by which visual information is incorporated into the brain's spatial framework to represent landmarks is poorly understood. Converging evidence points to the retrosplenial cortex (RSC) as an important locus for these computations.

To investigate how visual and spatial information are integrated in RSC, we developed a behavioral task in which head-fixed mice learned the spatial relationship between visual cues that act as landmarks and unmarked reward locations along a virtual corridor. Optogenetic inactivation of RSC on a subset of randomly interleaved trials produced a marked decrease in task performance. Two-photon GCaMP6f imaging revealed a large fraction of RSC neurons were active during behavior (60%). Spatially-tuned neurons tiled the virtual environment, however the majority of neurons were anchored to visual landmarks (30%), as opposed to trial onset (9%) or reward (18%). Landmarks were represented conjunctively by visual, motor, and task contingency components: presenting the same visual stimuli at a constant flow speed decoupled from the treadmill resulted in attenuated responses, which were modulated by animal motion.

To address if receptive field tunings were the result of local computations or were inherited from visual inputs we imaged GCaMP6f-expressing axonal boutons from primary visual cortex (V1) projections in RSC. We observed surprising similarity in the receptive fields of V1 boutons and RSC neurons. However, preliminary analyses indicate that V1 axons are less conjunctive and more responsive to purely visual stimuli. Ongoing work is aimed at identifying the relative contributions of visual and motor components to tuning of V1 inputs. Our experiments provide insight into the process by which visual inputs are integrated with motor and spatial information to produce landmark-like neural representations during goal-directed behavior in RSC.

III-57. Stable working memory representations in the presence of code morphing in the lateral prefrontal cortex

Camilo Libedinsky CAMILO@NUS.EDU.SG
Aishwarya Parthasarathy AISHU.PARTH@GMAIL.COM
Cheng Tang CHENGTANG827@GMAIL.COM
Roger Herikstad ROGER.HERIKSTAD@NUS.EDU.SG
Fah Cheong Loong ELECLF@NUS.EDU.SG
Shih-Cheng Yen SHIHCHENG@NUS.EDU.SG

National University of Singapore

Maintenance of working memory is driven by endogenous processes, which presumably involve networks of neurons with strong recurrent connections. We have previously shown that during the maintenance of memories, distractors evoke large changes in the population dynamics of the lateral prefrontal cortex, such that the population activity morphs from one stable state to another after the distractor. However, it is unclear how stable memory information can be extracted from these highly dynamic population responses. Here we show that we can extract stable memory information from the dynamic activity of the lateral prefrontal cortex. We used an optimization algorithm that reduced the distance between the Delay 1 and Delay 2 population activity while maintaining memory information. This optimization allowed us to find a mnemonic subspace where the pattern of activity remains invariant throughout the course of a trial. This invariance extended to periods of the trial that were not used to identify the mnemonic subspace, and it was absent in error trials. The existence of the mnemonic subspace appeared to be dependent on the parallel movement of trajectories associated with different memories in state space. Compared to other recurrent neural network models, the bump attractor model was the best at fitting the data, including the existence of code-morphing and the existence of a mnemonic subspace. The model also provided predictions that could be verified in the data. We conclude that the mnemonic subspace in prefrontal cortex could be used by downstream regions to read out stable memory information in the presence of code morphing.

III-58. A motion detector learned on transformed images replicates features of the fly neural circuit

Yanis Bahroun[1] Y.BAHROUN@LBORO.AC.UK
Dmitri Chklovskii[2] DCHKLOVSKII@FLATIRONINSTITUTE.ORG
Anirvan Sengupta[3,4] ANIRVANS.PHYSICS@GMAIL.COM

[1]Loughborough University
[2]Flatiron Institute and NYU Langone Medical Center
[3]Flatiron Institute
[4]Rutgers University

Classic models, such as Hassenstein-Reichardt correlator [5], detect motion by comparing light intensity in two adjacent points in space (transduced by adjacent photoreceptors). Yet, according to recent connectomics and physiology experiments in the fruit fly [7][1][2][6], elementary motion detectors (EMDs) require differently pre-processed signals from at least three adjacent photoreceptors. To address this discrepancy, we adopt a normative approach based on transformation learning and derive a global motion detector that integrates outputs of EMDs each receiving signals from three consecutive pixels. Moreover, we demonstrate that such a motion detector can be learned in an unsupervised setting as was suggested originally by Rao and Ruderman [4]. The novelty of our approach is to learn the detector in a biologically plausible network using recently developed similarity matching framework. Unlike previous work, we apply similarity not to individual images but to pairs of consecutive images. The normative nature of our approach makes it applicable to the general problem of learning content-invariant transformations which is crucial for successful pattern recognition. We map our algorithm onto a neural network (Fig.1) that recapitulates several salient anatomical, physiological and behavioral features of the fly. 1) EMD neurons receive signals from at least three adjacent locations in the visual field each pre-processed differently and in agreement with physiological measurements [1]. 2) Outputs of EMDs from different parts of the visual field are combined together in downstream units similar to the lobula plate tangential cells (LPTCs) in Drosophila. 3) The output of downstream units has the properties consistent with the physiology of the LPTCs: the output is proportional to contrast squared, modulated at the frequency of the stimulus, optimal velocity scales with stimulus wavelength. 4) Individual EMDs exhibit both phi and reverse-phi illusions [6].

III-59. Localized semi nonnegative matrix factorization of widefield imaging data

Shreya Saxena[1] SS5513@COLUMBIA.EDU
Simon Musall[2] SIMON.MUSALL@GMAIL.COM
Sharon Kim[1] SHK2146@COLUMBIA.EDU
Mohammed Shaik[1] MOHAMMED.ALTAF.SHAIK@GMAIL.COM
David Thibodeaux[1] DNT2111@COLUMBIA.EDU
Ying Ma[1] YM2415@COLUMBIA.EDU
Hanzhi Zhao[1] HZ2356@COLUMBIA.EDU
Elizabeth Hillman[1] EH2245@COLUMBIA.EDU
Anne Churchland[2] CHURCHLAND@CSHL.EDU
John Cunningham[1] JPC2181@COLUMBIA.EDU
Liam Paninski[1] LMP2107@COLUMBIA.EDU

[1]Columbia University
[2]Cold Spring Harbor Laboratory

Widefield calcium imaging is becoming more prevalent in mice as a high-throughput and non-invasive alternative to traditional approaches, such as two-photon imaging. Widefield recordings provide a global view of superficial dorsal cortex, with spatial resolution of around $20\mu m/pixel$, and temporal resolution of >30Hz. While high-throughput acquisition of these videos is becoming more routine, tools to process and understand the widefield data are lagging behind. Our goal is to compress, denoise, and demix the videos into meaningful spatial and temporal components that can be compared across mice and experimental conditions. One way to obtain an interpretable, stable decomposition is to map the inferred components onto well-defined brain regions. Although the Allen Mouse Common Coordinate Framework (CCF) defines a detailed map of the cortical structures for an average mouse, no known tools satisfactorily extract the activity of the different brain regions in individual mice in a data-driven manner, while taking into account mouse-specific and surgery-specific differences. Here, we introduce a variant of Non-Negative Matrix Factorization (NMF), termed localized semi-NMF (L-sNMF), that efficiently decomposes widefield video data and allows us to directly compare activity across multiple mice by outputting mouse-specific localized functional regions. L-sNMF uses a fast low-rank version of Hierarchical Alternating Least Squares (HALS), and outputs components that are significantly more interpretable than more traditional NMF or SVD-based techniques. Moreover, it provides a natural subspace to directly compare the correlation maps across different mice, and enables downstream identification of common signatures of activity and the modeling of dynamics across mice.

III-60. Determining biomarkers for anxiety by studying the neural basis of foraging behavior

Arjun Ramakrishnan ARJUNRA@UPENN.EDU
Dilara Berkay DBERKAY@SAS.UPENN.EDU
Michael Platt MPLATT@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Foraging for food efficiently is crucial for survival. Although animals, including humans, are on average optimal foragers, individuals vary considerably in their foraging strategies. Variation in the fundamental computations underlying foraging may provide a powerful opportunity to identify biomarkers of neural circuit dysfunctions associated with neuropsychiatric conditions. A neural circuit connecting posterior and anterior cingulate cortices (ACC), modulated by locus coerulus norepinephrine, has been implicated in variations in foraging behavior. This circuitry is also associated with anxiety disorders, suggesting a possible biological connection between the two. To test this idea, we measured the behavior of human participants (N=84) performing a patch foraging task where they made stay/leave decisions in environments varying in richness, while we monitored pupil size, reflective of

norepinephrine tone, and EEG source-localized to the ACC. Anxiety was assessed using standard self-report measures. We found that anxious individuals stayed at a patch for less time. Inducing anxiety led to the same effect, validating the result. A supervised machine learning algorithm trained on patch leaving time, pupil diameter and EEG, could predict anxiety levels (high/medium/low) with 70% accuracy in a subclinical cohort. To understand the neural basis of the ACC-EEG signal, we trained non-human primates (N=4) on the same task and monitored dACC (areas 24 a,b,c) spiking activity and local field potentials (LFPs). While dACC spiking activity reflected both the available reward and travel time costs, dACC LFPs reflected only the available reward. Model building to relate the ACC spiking activity and LFPs with ACC-EEG is currently underway. Overall, our results show that both momentary and long-term anxiety can be decoded from foraging behavior and associated physiological measures. These findings may provide neural circuit-based biomarkers for early detection of the neuropsychiatric disorder, monitoring response to treatment, and relapse prevention, particularly in individuals, such as children, for whom self-report is less reliable.

III-61. Using brain-score to evaluate and build neural networks for brain-like object recognition

Martin Schrimpf[1] MSCH@MIT.EDU
Kohitij Kar[1] KOHITIJ@MIT.EDU
Pouya Bashivan[1] BASHIVAN@MIT.EDU
Aran Nayebi[2] ANAYEBI@STANFORD.EDU
James DiCarlo[1] DICARLO@MIT.EDU
Jonas Kubilius[1,3] QBILIUS@MIT.EDU
Ha Hong[4] HAHONG@MIT.EDU
Najib J Majaj[5] NAJIB.MAJAJ@NYU.EDU
Rishi Rajalingham[1] RISHIR@MIT.EDU
Elias B Issa[6] ELIAS.ISSA@COLUMBIA.EDU
Daniel Bear[2] DBEAR@STANFORD.EDU
Jonathan Prescott-Roy[1] JJPR@MIT.EDU
Kailyn Schmidt[1] KAILYN@MIT.EDU
Daniel LK Yamins[2] YAMINS@STANFORD.EDU

[1]Massachusetts Institute of Technology
[2]Stanford University
[3]KU Leuven
[4]Bay Labs Inc.
[5]New York University
[6]Columbia University

The internal representations of early deep artificial neural networks (ANNs) were found to be remarkably similar to the internal neural representations measured experimentally in the primate brain. Here we ask, as ANNs have continued to evolve, are they becoming more or less brain-like? We therefore developed Brain-Score: a composite of multiple neural and behavioral benchmarks that score any ANN on how similar it is to the brain's mechanisms for core object recognition. Evaluating state-of-the-art ANNs, we found that ResNet and DenseNet families of models are the closest models from the Machine Learning community to primate ventral visual stream. Curiously, best current ImageNet models, such as PNASNet, were not the top-performing models on Brain-Score.

Despite high scores, these deep models are often hard to map onto the brain due to their vast number of layers and missing biologically-important connections, such as recurrence. We thus built CORnet-S: a neural network developed by using Brain-Score as a guide with the anatomical constraints of compactness and recurrence. Although a shallow model with four anatomically mapped areas with recurrent connectivity, CORnet-S is a top model on Brain-Score and outperforms similarly compact models on ImageNet.

Finally, to further validate our claims, we are including new behavioral measurements that models have not yet been tested on. The scores on this novel data inform us about how well model rankings and specifically

CORnet-S performance generalize to new benchmarks. Altogether, we propose that evaluating and tracking model-benchmark correspondences through a Brain-Score can be used to guide machine network evolution, and machine networks as mechanistic hypotheses of the brain's network can drive next experiments. To facilitate both of these, we release Brain-Score.org: a platform that hosts the neural and behavioral benchmarks, where ANNs can be submitted to receive a Brain-Score, and where new experimental data can be naturally incorporated.

III-62. Gradient-based analysis of wide-field calcium imagining data using Gaussian processes

Mikio Aoi[1] MAOI@PRINCETON.EDU
Benjamin Scott[2] BBS@BU.EDU
Anqi Wu[1] ANQIW@PRINCETON.EDU
Stephan Thiberge[1] THIBERGE@PRINCETON.EDU
Carlos Brody[1] BRODY@PRINCETON.EDU
David W Tank[1] DWTANK@PRINCETON.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Boston University

Wide-field calcium imaging (WFI) is a relatively recent modality for the analysis of neuronal activity. Data from WFI are arranged on a 3D grid; i.e. observations are on a regular lattice of two spatial dimensions and one time dimension. This grid structure provides opportunities for new analysis methods that examine the spatio-temporal evolution of neuronal activity across the cortex. Indeed, reports of WFI data have described activity organized as "waves", and "moving blobs." However, few studies have identified an effective way of describing these types of spatio-temporal dynamics by taking explicit advantage of the grid structure. In this study we develop a new modeling framework for WFI that allows experimenters to identify explicit dynamics across the cortical surface on a trial-by-trial basis using Gaussian processes (GP). The GP is a probabilistic generative distribution over functions. For WFI, the GP acts as a prior over a smooth instantiation of the calcium fluorescence and the data are a noisy observation of this smooth instantiation. The advantage of the GP framework is that, because the learned GP is smooth (i.e. differentiable), it allows us to do inference on partial derivatives of the GP of any order. We may therefore examine the dynamics of WFI data by studying the spatio-temporal gradient over time. To demonstrate our method, we apply this modeling approach to simulated WFI data as well as real data recorded in vivo from GCaMP transgenic rats using a head-mounted microscope while the rats performed a perceptual decision making task.

III-63. Sparse lateral inhibition mediates robust phase coding in the olfactory bulb

Marin Ozaki[1,2] OZAKIM@STUDENT.ETHZ.CH
Takeshi Imai[3] T-IMAI@MED.KYUSHU-U.AC.JP

[1]University of Zurich
[2]ETH Zurich
[3]Kyushu University

While temporal coding is considered a ubiquitous phenomenon in the brain, its circuit basis remains largely unknown. In the rodent olfactory bulb, odors are represented not only by the firing rate (rate coding) but also by the temporal patterns of oscillatory activity in mitral/tufted (MT) cells, known as phase coding. Oscillatory activity in MT cells is produced by the sniff-coupled mechanosensory inputs to olfactory sensory neurons (OSNs). While oscillation phases are homogeneous in OSNs, they are heterogeneous in MT cells, suggesting that phase coding

emerges in the olfactory bulb circuitry. Oscillation phases in MT cells are unaffected by sniff conditions (frequency and magnitude) but demonstrate reliable and concentration-invariant shifts upon odor stimulation. The rate code, in contrast, varies under different sniffing conditions and is affected by odor concentration. Furthermore, the phase code patterns remain consistent across multiple sniff cycles, whereas the rate code patterns are reorganized and decorrelate over time. Based on these observations, it has been proposed that phase coding is the basis for encoding odor identities in the olfactory bulb (Iwata et al., 2017). To explore the possible circuit mechanisms underlying phase coding in the olfactory system, we carried out computational simulations for a network model of the olfactory bulb. This model consists of OSNs, MT cells, and short axon (SA) cells, which mediate lateral inhibition among glomeruli. We found that sparse, rather than dense, lateral inhibition is essential for the concentration- and sniff mode-invariant phase coding of odors in MT cells. We also found that the co-existence of two different types of inhibitions (fast and slow) accounts for stable phase coding and labile rate coding in MT cells. Our study suggests a previously unappreciated role for the sparse lateral inhibition in sensory coding.

III-64. Spontaneous cortical waves modulate neural and perceptual sensitivity in awake, behaving marmosets

Zachary Davis[1] ZDAVIS@SALK.EDU
Lyle Muller[1] LMULLER@SALK.EDU
Terrence Sejnowski[1] TERRY@SALK.EDU
Julio Martinez-Trujillo[2] JULIO.MARTINEZ@ROBARTS.CA
John Reynolds[1] REYNOLDS@SALK.EDU

[1]Salk Institute for Biological Sciences
[2]University of Western Ontario

Our ability to perceive a subtle environmental stimulus is not fixed. Rather, perceptual thresholds fluctuate with changes in arousal, attention, and expectation. Similarly, individual neurons within the visual cortex exhibit variable responses across repeated presentations of the same stimulus, and there is evidence that these fluctuations contribute to perceptual variability. Injection of identical noisy currents evoke repeatable and highly precise spike trains, indicating that the spiking mechanism is reliable, and thus not a major source of this variability. Spiking variability instead likely reflects moment-by-moment changes in synaptic and neuromodulatory input from the cortical network. These network fluctuations are reflected in local field potentials (LFPs), which reflect the summed synaptic currents near the recording electrode. We used a Utah array to record neuronal responses and LFP fluctuations in extrastriate cortex of the awake, behaving marmoset (Callithrix jacchus) during a threshold detection task in which a low-contrast drifting Gabor appeared at a variable time following fixation. While LFP fluctuations are typically treated as narrowband oscillations, here we introduce an approach to analyze the generalized phase of wideband, non-frequency-resolved signals. We combine this approach with a statistical method to detect spontaneous traveling waves (STWs) with high temporal precision, on a moment-by-moment basis. We find that (1) wideband network fluctuations are STWs propagating across the cortex in the awake marmoset, (2) STWs modulate ongoing neuronal spiking activity as well as stimulus-evoked responses, and (3) the modulation of stimulus-evoked responses by STWs leads to a robust modulation of perceptual sensitivity. We explain these results using large-scale spiking models and effective population descriptions. Our results suggest that STWs are an important feature of ongoing activity in the awake cortex and that they actively shape sensory perception.

III-65. Learning successor representations in partially observed environments

Eszter Vertes VERTES.ESZTER@GMAIL.COM
Maneesh Sahani MANEESH@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, University College London

Animals need to devise strategies to maximize returns collected while interacting with their environment based on incoming noisy sensory observations. A crucial step in this process is to accurately estimate the value of each state from past experience. Approximating value functions is a non-trivial problem since rewards are typically spatially and temporally sparse. Moreover, the task-relevant states, such as the agent's location within an environment or presence of a predator, are often not directly observable but have to be inferred using available sensory information. A number of different approaches have been proposed in the reinforcement learning literature for computing values, most of which fall into the category of model based or model free approaches. Successor representations (SR, Dayan, 1993) can be thought of as a middle-ground between these strategies allowing for fast value computation reminiscent of model free methods, while being able to quickly adapt to changes in the reward function or goal locations. While their algorithmic properties have been well studied and recently consistent experimental findings have been presented from hippocampal place cells (Stachenfeld et al., 2017), the extension of SR-s to more challenging, partially observed environments has been lacking. Here, we introduce a framework that allows for efficient value function computation in partially observed environments via the successor representation. Our approach relies on Distributed Distributional Codes (DDC, Vertes and Sahani, 2018, Sahani and Dayan, 2003) as neural representation of uncertainty and is applicable to both discrete and continuous state spaces. We show that representing uncertainty over latent variables using DDC-s can be naturally integrated with representations of uncertainty about future states and therefore can generalize SR-s to more realistic environments with partial observability.

III-66. Topographic organization of distinct tuning maps for optimal tiling of sensory modules in visual cortex

Min Song NIGHT@KAIST.AC.KR
Jaeson Jang JAESONJANG@KAIST.AC.KR
Gwangsu Kim KGSPIANO@KAIST.AC.KR
Se-Bum Paik SBPAIK@KAIST.AC.KR

Korea Advanced Institute of Science and Technology (KAIST)

Visual information is composed of various features of objects, such as orientation of edges and spatial frequency of patterns, and thus requires multiple types of sensory tuning modules for encoding. For this, neurons in the primary visual cortex (V1) of higher mammals are tuned to each visual feature, and are organized into distinct tuning maps across the cortex. Moreover, the topography of these maps is spatially correlated in an orthogonal/parallel manner, resulting in an efficient tiling of sensory modules. However, it remains unclear how such a systematic organization could develop in V1. Here, we propose a developmental model in which the spatial organization of each tuning map is seeded altogether from the structure of retinal mosaics, and this common framework provides a blueprint of efficient tiling of multiple tuning maps. By extending the Paik-Ringach model (2011), we show that the spatial arrangement of neighbouring ON and OFF retinal ganglion cells (RGCs) can initiate diverse functional tunings in V1. First, the temporal delay between ON and OFF RGC responses induces direction tuning of a V1 neuron. Second, a periodic variation of ON-OFF RGC distances results in spatial variation of the contralateral cortical response, and this induces a periodic pattern of contra- vs. ipsi-lateral balance in the ocular dominance map. Next, phase difference between contra- and ipsi-lateral receptive fields in the binocular region induces narrower ON and OFF subregions, resulting in a preference for higher spatial frequency than that in monocular region. As a result, our model simulation successfully produced not only tuning maps of orientation, frequency, direction, and ocular dominance, but also topographic correlations among the maps. Our results suggest that the structure of the peripheral visual system can provide a common framework of diverse cortical tuning maps and a parsimonious mechanism of building an efficient tiling of distinct sensory modules.

III-67. Learning sensorimotor affordances by goal-directed action

Guy Isely GISELY@GMAIL.COM
Friedrich Sommer FSOMMER@BERKELEY.EDU

University of California, Berkeley

The sensorimotor affordances available to an organism determine the ease or difficulty with which it achieves its goals. In attempting to achieve a goal, new affordances can be learned that generalize across goals. We propose a model of learned sensorimotor affordances where sensory inputs are encoded in a quasimetric contextual state space, in which distances between contexts are optimized to match the experienced behavioral costs of navigating between those contexts. Previous work on goal-directed action in reinforcement learning has often used the structure of the reward function to represent different goals [2, 3]. In contrast, in our model goals are represented directly as contextual states. This allows direct estimation of the behavioral cost to a goal as a distance in state space. Reduction of this distance is a progress signal that can be used to optimize production of motor outputs toward a goal. Thus, experiences produced under goal-directed action provide self-supervised training signals for both the sensory encoding and motor production components of our model. For a simple simulated environment where we can directly estimate the minimum path length between points in the environment, we show that distances in learned contextual space are correlated with path length. The model predicts that the overlap of hippocampal place cell representations should vary with the behavioral costs of moving between states rather than just the spatial distance between the places they represent.

III-68. Noise or signal? Psychophysical evidence for the role of sensory variability

Adam Koblinger[1] KOBLINGER_ADAM@PHD.CEU.EDU
David Zoltowski[2] ZOLTOWSKI@PRINCETON.EDU
Jozsef Fiser[1] FISERJ@CEU.EDU
Mate Lengyel[3] M.LENGYEL@ENG.CAM.AC.UK

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

Stimulus-independent fluctuations in the responses of sensory neurons are traditionally considered as mere noise, and thus a source of perceptual ambiguity. In contrast, sampling-based models of perceptual inference suggest that the magnitude of this intrinsic variability acts as a signal: it conveys information about the uncertainty in low-level perceptual estimates. In both cases, to improve accuracy, downstream areas need to average sensory responses over time, as in classical models of evidence accumulation. However, due to the different roles that upstream sensory variability plays under the “noise” and “signal” hypotheses, the *uncertainty* about this average behaves in fundamentally different ways in them: it is respectively related to the standard error or the standard deviation of responses. In order to compare these hypotheses, we used a modified orientation estimation paradigm in which, on every trial, subjects simultaneously reported their best estimate of one of several briefly viewed, static line segments and their confidence about this estimate. We varied the difficulty of trials by changing the number of line segments, their contrast level, and the presentation time of the display. In general, we found that subjects’ confidence predicted their accuracy even when controlling for these experimentally manipulated stimulus parameters. This indicated that subjects had a well-calibrated trial-by-trial subjective measure of their uncertainty and did not only rely on extrinsic stimulus parameters to gauge the difficulty of a trial. Critically, while both models could account for changes in estimation performance with stimulus parameters, only the “signal” model predicted correctly the experimentally observed changes in confidence reports, and in the strength of correlation between confidence reports and actual accuracy. These results offer a new psychophysical window onto the role of sensory variability in perception and indicate that it conveys useful information about uncertainty.

III-69. Pattern completion by recurrent cortical circuits stabilizes odor representations across activity regimes

Kevin Bolding BOLDING@NEURO.DUKE.EDU
Shivathmihai Nagappan SHIVA.NC@DUKE.EDU
Bao-Xia Han BAO-XIA.HAN@DUKE.EDU
Fan Wang FAN.WANG@DUKE.EDU
Kevin Franks FRANKS@NEURO.DUKE.EDU

Duke University

Sensory representations must be stabilized across varying, noisy stimulus conditions and across behavioral states to maintain perceptual constancy. Stabilization can occur through pattern completion, by which a circuit generates stable output despite degraded input through the recruitment of interconnected cells within a recurrent network. By selectively eliminating recurrent connectivity in olfactory cortex, we provide direct mechanistic evidence that pattern completion stabilizes cortical odor representations. We simultaneously recorded odor-evoked spiking from populations of olfactory bulb (OB) and piriform cortex (PCx) neurons in head-fixed mice before and during anesthesia. Both odor responses and the ability to decode odor identity diminished under anesthesia in OB, but were preserved in PCx. Furthermore, a decoder trained on responses from one state and tested on responses from the other could accurately classify PCx responses, but not OB responses. To determine whether stabilization in PCx is indeed mediated by intracortical recurrent connections we first eliminated recurrent excitation, and recruitment of feedback inhibition, by selectively expressing tetanus toxin in PCx principal cells. Without recurrent collaterals, within-state PCx decoding was impaired and cross-state decoding was abolished. Second, we developed a mouse line to distinguish odor responses between the two major classes of PCx principal cells: semilunar (SL) cells, which lack recurrent inputs, and superficial pyramidal (SP) cells, which receive extensive recurrent inputs. Most SP cells were robust across states whereas SL cell responses were mostly state-specific, as in OB. Thus, recurrent collateral circuitry stabilizes cortical odor representations. Typically, only the earliest OB responses are robust across state. We show that recurrent excitatory connections first amplify the impact of these early inputs in PCx, and then recruit strong, global inhibition to discount later OB inputs, thus ensuring reliable odor coding from degraded OB input.

III-70. State space models for multiple interacting neural populations

Joshua Glaser[1] JIG2124@COLUMBIA.EDU
Scott Linderman[1] SCOTT.LINDERMAN@COLUMBIA.EDU
Matthew Whiteway[1] THEMATTINTHEHATT@GMAIL.COM
Matthew Perich[2] MPERICH@GMAIL.COM
Brian Dekleva[3] BRIAN.DEKLEVA@GMAIL.COM
Lee Miller[4] LM@NORTHWESTERN.EDU
Liam Paninski[1] LIAM@STAT.COLUMBIA.EDU
John Cunningham[1] JPC2181@COLUMBIA.EDU

[1]Columbia University
[2]University of Geneva
[3]University of Pittsburgh
[4]Northwestern University

As we move toward more complex recordings spanning multiple brain areas and cell types, existing data analysis methods may not provide a clear picture of the intra- and inter-population dynamics. Lacking any knowledge of subpopulation structure, factor analysis and its linear dynamical systems generalizations find intermixed representations that explain all populations with the same factors, leaving the practitioner to disentangle the result. Here, we leverage our knowledge of subpopulation structure to constrain state space (latent variable) models such that the state space is partitioned into separate subspaces for each subpopulation. The latent variables of these

subpopulations interact with each other across time through a linear dynamical system. This allows separating internal and external contributions to dynamics. In simulations, we demonstrate this approach applied to Poisson linear dynamical systems (PLDS) and switching linear dynamical systems (SLDS) models. In these simulations, the PLDS model can accurately recover important aspects of the interaction between populations of neurons. In particular, it can accurately recover which changes in neural activity are due to internal dynamics versus input from another population. Additionally, it can recover the dimensionality of interactions between populations. Moreover, in a simulation in which one population sporadically provides input to another population, the constrained SLDS model can accurately determine these times of interaction. We then provide a preliminary demonstration of this approach on real data, by fitting the constrained PLDS model to simultaneously recorded neurons in primary motor cortex (M1) and dorsal premotor cortex (PMd). The model suggests that most changes in activity are due to internal dynamics. Constrained state space models promise to be tools for better understanding the interactions between neural populations.

III-71. Inferring animal learning rules from behavior with nonlinear state space models

Farhan Damani FDAMANI@CS.PRINCETON.EDU
Nicholas Roy NROY@PRINCETON.EDU
Athena Akrami ATHENA.AKRAMI@UCL.AC.UK
Carlos Brody BRODY@PRINCETON.EDU
Ryan Adams RPA@PRINCETON.EDU
Jonathan Pillow PILLOW@PRINCETON.EDU

Princeton University

High-throughput behavioral neuroscience presents an exciting opportunity to investigate the dynamics of animal learning during decision making. We describe a probabilistic recurrent time-series model that captures a rat's underlying time-dependent decision making strategy while learning a forced choice discrimination task. Through priors that explicitly characterize learning strategies—including policy gradient and Q-learning—our approach presents the opportunity to infer reliable learning and exploration strategies despite observing only a single sequence of decisions for each rat. Importantly, our model is probabilistic, which means we can reason about a rat's competence, memory lapse, perception, and motivation levels using the marginal likelihood. To infer a rat's underlying knowledge state during a particular trial, we use a scalable black box variational inference approach, which allows for efficient learning and scientifically guided model selection. We apply our method to both simulated data and a large-scale rat training dataset (20,000 trials) (Akrami 2018) and discover plausible learning rules. Our framework allows for reasoning about how rats integrate information from its environment and from its history to make decisions. We also evaluate a rat's generalization capabilities when confronted with an unknown stimuli. Our method can also be combined with reinforcement learning to allow for a "teacher" to intelligently decide what sequence of stimuli to present to a rat in order to speed up task learning. Beyond learning efficiency, tailored instruction can also help push our understanding of what tasks we believe rats are cognitive capable of mastering. Together our composable machine learning framework can assist in discovering novel learning mechanisms in a data-driven manner and devise highly tailored instructions for effective learning.

III-72. Degenerate solution networks (DSNs) for theoretical neuroscience

Sean Bittner SEANRBITTNER@GMAIL.COM
John Cunningham JPC2181@COLUMBIA.EDU

Columbia University

Theoretical neuroscientists design and test mathematical models of neural activity, assessing a model's quality by its ability to replicate experimental results. A general challenge has been addressing the indeterminacy of model parameterizations; different parameterizations of these models accurately reflect established experiments, yet make incompatible predictions about other aspects of neural activity. We present a novel machine learning methodology called degenerate solution networks (DSNs), which learn the full (i.e. maximum entropy) distribution of generative model parameterizations that yield a behavior of interest. DSNs are a tool designed for theorists, that enables a new class of model analyses relying on the full distribution of generative parameters that result in some statistically specified activity. For example, consider the stomatogastric ganglion (STG) circuit in the crustacean, which generates tri-phasic rhythms and has a well-studied biophysical model. Rather than examining the parameter space of this model though extensive simulation as is common practice, we could directly learn the maximally random distribution of channel conductances and synaptic efficacies that yield tri-phasic rhythms with a DSN. Here, we demonstrated the utility of DSNs in three (of many possible) use cases. First, we learned the degenerate solution space of 2D linear systems producing a band of pure oscillations. Then, relying on input-driven dynamic mean field theory, we use DSNs to study solution spaces of recurrent neural networks (RNNs), one of the most widely used models in theoretical neuroscience. Finally, in contrast to RNNs, we discuss using DSNs to learn the parametric solution space of biophysically realistic models such as the STG. More generally, we speculate that access to such degenerate parameterizations can facilitate model comparison, inform model revision, and guide experimental design.

III-73. Towards an understanding of the algorithmic utility of recurrent connections

Brett Larsen[1] BWLARSEN@STANFORD.EDU
Shaul Druckmann[2,1] SHAULD@STANFORD.EDU

[1]Stanford University
[2]HHMI Janelia Research Campus

One of the brain's most striking anatomical features is the amount of feedback and lateral/recurrent connections in neural circuits. Despite the ubiquity of this circuit organization, most theoretical studies of recurrent neural networks have focused on analyzing the impact of recurrence on tasks that have a direct temporal nature, rather than considering what general computations would benefit from a recurrent (as opposed to feed-forward) architecture. In this work, we demonstrate two answers to this question: (1) Recurrent networks can much more efficiently perform tasks that require repeated, local computations to propagate information. And (2) recurrent models can be trained to perform such tasks more easily the more we know about the structure of this local computation, in essence the better prior we have about which neurons are considered local in the calculation.

In considering such computations, we build on the work of Minsky and Roelfsema studying a task (detection of edge-connected pixels) for which an efficient recurrent solution exists that propagates local information and implementation of the same solution in a single hidden layer feed-forward network is extremely inefficient. We extend this work by comparing the recurrent solution to a deep network trained on the task and demonstrate how computations of this form are found in a broad array of tasks in reinforcement learning and inference. This reflects the underlying understanding from physics that causal influence spreads locally and demonstrates how such tasks can be constructed in abstract connection spaces.

Finally, we empirically study how differing levels of knowledge of the computational structure affects a recurrent network's ability to learn such tasks. This is done by training RNNs to detect edge-connected pixels with increasingly restrictive parameter spaces, pruning the connections to represent various priors on locality of neurons and studying how such priors affect performance and sample complexity.

III-74. 4D analyses reveal molecular requirements for single neuron identity during circuit wiring

Samantha Ing-Esteves SAMANTHA.ESTEVES@MAIL.UTORONTO.CA
Julie Lefebvre JULIE.LEFEBVRE@SICKKIDS.CA

University of Toronto

Mature neural circuits require complex dendritic arbour patterns. During development, neurons produce excess branches and filopodia that are selectively pruned, but the mechanisms that regulate this selection and final arbour pattern are poorly understood. One mechanism of dendrite patterning is called dendrite self-avoidance, in which dendrites originating from the same cell (self-dendrites) minimize overlap and contacts with each other. We have shown that self-avoidance is regulated by the clustered protocadherins (Pcdhs), a family of 60 recognition molecules able to confer neurons with unique molecular identities. Past studies showed that Pcdh isoforms form multimers that interact homophilically. We hypothesize that self-avoidance is a dynamic process mediated by Pcdh-dependent homophilic contact between self-dendrites and that the Pcdh intracellular region is required to signal self-avoidance.

To test this, we acquired time-lapsed 3D volumes of a developing retinal interneuron followed by automatic neuron reconstruction. In these 4-dimensional videos (3D+time) we observe transient self-contacts that form ‘closed-loop structures’ in the arbour. SWC files, the common output format for open-sourced tracing algorithms, do not allow for the encoding of closed-loop structures. Here we are developing an analysis pipeline to correct for closed-looped structures and extract geometrically defined dendritic features from SWC traces. The resulting corrected reconstructions containing closed-loop structures will be essential in modernizing current simulations of functional circuits as cell morphologies can now be accurately represented. Additionally, we are interrogating the mechanism of dendrite patterning at the molecular level. To identify the Pcdh domain required for self-avoidance we expressed a transgene lacking the Pcdh intracellular region in a previously established Pcdh-knockout mouse. Assaying self-avoidance defects in these animals will reveal the essential Pcdh intracellular domain needed for self-avoidance signaling. These studies will provide an open-sourced pipeline for 4D feature extraction and reveal the molecular and cellular mechanism of dendrite patterning.

III-75. A large-scale physiological survey reveals higher order coding throughout the mouse visual cortex

Saskia de Vries SASKIAD@ALLENINSTITUTE.ORG
Michael A Buice MICHAELBU@ALLENINSTITUTE.ORG
Jerome Lecoq JEROMEL@ALLENINSTITUTE.ORG
Gabriel Ocker GABEO@ALLENINSTITUTE.ORG
Michael Oliver MICHAELO@ALLENINSTITUTE.ORG
Daniel Millman DANIELM@ALLENINSTITUTE.ORG
Peter Ledochowitsch PETERL@ALLENINSTITUTE.ORG
Peter Groblewski PETERG@ALLENINSTITUTE.ORG
Nicholas Cain NICHOLASC@ALLENINSTITUTE.ORG
Clay Reid CLAYR@ALLENINSTITUTE.ORG
Christof Koch CHRISTOFK@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

An important open question in visual neuroscience is how visual information is represented in cortex. Important results characterized neural coding by assessing the responses to artificial stimuli, with the assumption that responses to gratings, for example, capture the key features of neural responses, and deviations, such as extra-classical effects, are relatively minor. The failure of these responses to have strong predictive power has renewed these questions. It has been suggested that this characterization of visual responses has been strongly influenced by the biases inherent in recording methods and the limited stimuli used in experiments. In creating the Allen Brain Observatory, we sought to reduce these biases by recording large populations of neurons using a broad array of stimuli, both artificial and natural, surveying over 63,000 neurons in the mouse visual cortex. We find that visual responses throughout the mouse cortex are highly variable. Few cell's tuning curves are predictive on a trial-by-trial basis. This variability cannot be accounted for by running speed or eye position. While for a few cells the temporal response within trials can be explained reasonably well by linear-nonlinear models, for the majority it cannot. Comparing stimuli, we find that whether a neuron is responsive to one stimulus is independent of whether it responds to another, i.e. whether a neuron responds to natural scenes provides no information about whether it responds to natural movies or gratings. This surprising phenomenon cannot be explained by standard local filter-based models, but we show that this is consistent with multi-layer computation, as found in deeper layers of standard convolutional networks. These results suggest that visual responses in the mouse cortex are more selective and context dependent than previously thought, and demonstrate that the traditional view of visual cortical computation is inadequate for describing the function for the majority of cortical cells.

III-76. Stimulus-driven bursting dominates neural activity measured by calcium imaging

Peter Ledochowitsch PETERL@ALLENINSTITUTE.ORG
Nicholas Cain NICHOLASC@ALLENINSTITUTE.ORG
Joshua Siegle JOSHS@ALLENINSTITUTE.ORG
Xiaoxuan Jia XIAOXUANJ@ALLENINSTITUTE.ORG
Ulf Knoblich KNOBLICH@GMAIL.COM
Lawrence Huang LAWRENCEH@ALLENINSTITUTE.ORG
Lu Li LUL@ALLENINSTITUTE.ORG
Gabriel Ocker GABEO@ALLENINSTITUTE.ORG
Daniel Millman DANIELM@ALLENINSTITUTE.ORG
Michael Oliver MICHAELO@ALLENINSTITUTE.ORG
Hongkui Zeng HONGKUIZ@ALLENINSTITUTE.ORG
Shawn Olsen SHAWNO@ALLENINSTITUTE.ORG
Clay Reid CLAYR@ALLENINSTITUTE.ORG
Stefan Mihalas STEFANM@ALLENINSTITUTE.ORG
Saskia de Vries SASKIAD@ALLENINSTITUTE.ORG
Michael A Buice MICHAELBU@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

For the purpose of analysis and modeling of visually driven neural responses, stimulus-evoked events extracted from calcium imaging experiments via 'spike-deconvolution' algorithms are commonly treated like electrophysiologically recorded action potentials. However, such calcium events may not correspond to single action potentials; a fact often acknowledged, yet rarely quantified or discussed in detail. In this work, we used 'ground truth' data – captured by two-photon microscopy and resampled to match the noise profile of the Allen Brain Observatory [1], combined with simultaneous, juxtacellular electrophysiology in vivo – to characterize, which spike train features elicited extractable calcium responses. We found that detectable changes in fluorescence were mostly associated with short bursts, containing brief (intra-burst) inter-spike intervals (ISI), and rarely with isolated spikes ($<$10% detection probability). Next, we investigated how such bias in favor of specific types of bursting activity affected the analysis of visually evoked neural responses. We recorded extracellular spike trains from five regions of mouse (awake, head-fixed) visual cortex using Neuropixels probes [2] while presenting images (natural scenes). All spike trains were decomposed into components: bursting [3, 4], doublet, and isolated spikes. We observed that for some neurons the change in firing rate evoked by their preferred image was in fact a selective increase in the rate of just one component, consistently across trials. Moreover, for some cells, distinct components preferred different stimuli. Using the model MLSpike [5] trained on the 'ground truth' data set, we generated synthetic fluorescence data for each recorded spike train, and obtained event trains by L0-regularized detection [6]. A comparative visual response analysis performed on the spike train components, and on the synthetic calcium events, respectively, confirmed that stimulus preferences and response metrics computed from synthetic calcium data were more consistent with those computed from bursts and doublets, than with results obtained from isolated spikes.

III-77. Do neurons learn how to learn?

Benjamin Lansdell LANSDELL@SEAS.UPENN.EDU
Prashanth Prakash PRPRAK@SEAS.UPENN.EDU
Konrad Kording KOERDING@GMAIL.COM

University of Pennsylvania

Backpropagation learning is driving today's artificial neural networks (ANNs). However, despite extensive research, it remains unclear if the brain implements this algorithm, in particular in recurrent neural networks (RNNs). Reinforcement learning (RL) is often seen as a realistic alternative: neurons can randomly introduce change, and use unspecific feedback signals to observe their effect on the cost and thus approximate their gradient. However, such learning scales poorly to large problems. Here we propose a hybrid learning approach. Each neuron uses an RL-type strategy to learn how to approximate the gradients that backpropagation would provide – it learns to learn (L2L). Such an approach may naturally align with proposals that cortical pyramidal neurons have a teaching circuit implemented in their apical dendrites. In such a strategy, learning involves two components, a learner that updates a neuron's weights and a meta-learner that modifies the learner. In a feed-forward network, we show that L2L learns to approximate the gradient provided by backpropagation, and matches the performance of gradient-based learning. This performance is expected since even fixed random feedback weights, known as feedback alignment, are known to converge for small deep learning problems. In an RNN, the system closely matches the performance of backpropagation, whereas feedback alignment fails to converge. Learning to learn promises to allow neurons to obtain good performance without the need for precise pre-specified learning rules.

III-78. Inclusive normativity: control costs and motor delays shape optimal perceptual decision making

Juan Castineiras[1,2] JUAN.CASTINEIRAS@NEURO.FCHAMPALIMAUD.ORG
Alfonso Renart[3] ALFONSO.RENART@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Foundation
[2]Champalimaud Research
[3]Champalimaud Centre for the Unknown

Adaptive behavior in a perceptual decision making (PDM) task involves a trade-off between the performance deficit of responding too soon and the cost of time associated to gathering evidence. The optimal solution to this problem has been understood using the partially observable markov decision process framework (POMDP); however, these kinds of policies ignore constraints that real agents have to confront. Here, we generalize ideas from optimal control theory to derive optimal policies that solve a categorical PDM task in the presence of two types of generic costs. One is a "cost of control", which penalizes policies that deviate from the agent's default actions. In particular, we consider the effect of impulsivity, a spontaneous tendency to respond (consistent with the concept of exploration). The other is the cost derived from "motor delays", i.e., the unavoidable delay between making a choice and the registration of the corresponding action that stops the stimulus. We show that the optimal policy in this setting involves a separation between a commitment to act, and the specification of the action – because the agent can use information arriving during the "motor delay" to improve action selection. We provide semi-analytical solutions for the probability that the agent will choose an option at a given time with a given belief, deriving predictions on measurable observables: choice, reaction time (RT) and decision confidence. We show that when the cost of control or the motor delay are significant, the behavior of decision confidence departs

from the POMDP solution and resembles predictions of signal detection theory. Overall, our results clarify the link between the observed phenomenology of decision confidence and different notions of optimality, and, more generally, clear the way to a more inclusive notion of normative behavior that includes both task contingencies as well as unavoidable costs faced by real organisms.

III-79. The perceptual logic of a spatio-temporal sensory code

Edmund Chong[1] EDMUND.W.CHONG@GMAIL.COM
Monica Moroni[2,3] MONICA.MORONI@IIT.IT
Christoper Wilson[1] CHRISTOPHER.WILSON@NYUMC.ORG
Stefano Panzeri[2] STEFANO.PANZERI@IIT.IT
Dmitry Rinberg[1] RINBERG@NYU.EDU

[1]New York University
[2]Istituto Italiano di Tecnologia
[3]Center for Mind and Brain Sciences (CIMeC)

Brains commonly encode sensory information within spatio-temporal activity patterns. However, little is known of how these patterns are read out to inform behavior and perception. In the olfactory bulb (OB), odors are represented in a spatially-distributed, temporally-precise code. To measure how features in the code map to behavior, we trained mice to recognize optogenetically-driven OB activity, or artificial 'odors', manipulating features in an independent, precise and parametric manner. We developed a novel experimental framework to measure how perceptual distances vary with changes along various spatial and temporal dimensions within activity patterns. We examined changes along each feature dimension, compared the importance of multiple dimensions within a common framework, and uncovered a model of how various spatial and temporal features are jointly processed to produce olfactory perception. Replacing spatial components in activity patterns produce graded perceptual effects that sum linearly, which is not predicted by previously-suggested coding schemes that maximize representational capacity (small changes representing unrelated odors). Changing temporal latencies relative to sniff affects animals' responses, consistent with previous electrophysiological experiments. More surprisingly, changing relative latencies within each pattern, as opposed to sniff-aligned latencies, has a far greater effect on animals' perceptual decisions. Both spatial and temporal features are modulated by a preferential weighting of early activity, such that replacing or temporally shifting the earlier-activated components in the pattern has stronger effects, supporting a previously-proposed 'primacy' code in olfaction. A process model performing linear summation of spatial components defined in sniff and relative time coordinates, and modulated by primacy, explains animals' responses, revealing the perceptual logic of the spatio-temporal olfactory code. To our knowledge, this is the first study to measure behavioral responses across a broad space of precisely and systematically-manipulated neural activity patterns.

III-80. Arousal enhances reliability of population coding in primary auditory cortex

Charles Heller HELLERC@OHSU.EDU
Daniela Saderi SADERID@OHSU.EDU
Zachary Schwartz SCHWARZA@OHSU.EDU
Stephen David DAVIDS@OHSU.EDU

Oregon Health and Science University

Arousal increases the gain of sensory cortical responses and reduces the amount of shared variability, or noise correlations. It remains unknown what physiological processes underlie these changes and how they impact sensory encoding. We investigated this question by monitoring the arousal level of awake ferrets via pupillometry

while using a 64-channel silicon probe to record spiking activity from primary auditory cortex (A1) and presenting natural sound stimuli. We clustered units based on their spike waveform shape to distinguish between putative fast-spiking (FS) interneurons and regular-spiking (RS) pyramidal neurons. For both classes of units, we found that over half exhibited significant arousal-dependent modulation in firing rates and that sensory responses were more reliable during high arousal. To determine to what extent these arousal-dependent modulations influenced population coding, we measured noise correlations before and after regressing out the effects of pupil-indexed arousal on the activity of individual units. Noise correlations were reduced for pairs of units showing significant firing rate modulation, demonstrating that pupil-indexed arousal acts as a common modulator in A1. However, even after this correction, noise correlations continued to depend on arousal. FS/FS pairs significantly decreased their correlations during high arousal while RS/RS pairs did not. This suggests that pupil-indexed arousal reflects multiple, cell-type specific processes in A1. Finally, we asked if these changes affected coding at the population level by measuring the degree to which trial-trial residual activity in the population diverged from the mean stimulus response. The magnitude of this divergence was significantly and positively correlated with pupil-indexed arousal, suggesting that arousal enhances sensory encoding by decreasing the interference between correlated neural activity and sound evoked responses.

III-81. Stable memories despite large spontaneous synaptic fluctuations

Samuel Eckmann ECKMANN@FIAS.UNI-FRANKFURT.DE
Suneet Singh Jhutty JHUTTY@FIAS.UNI-FRANKFURT.DE
Jochen Triesch TRIESCH@FIAS.UNI-FRANKFURT.DE

Frankfurt Institute for Advanced Studies

Memories are thought to be stored as patterns of synaptic efficacies. However, individual synapses exhibit strong efficacy fluctuations even in the absence of neural activity. Such fluctuations seem to be at odds with the long-term storage of information in synaptic weight patterns. Here we present a model that tries to resolve this apparent paradox. The model is based on the stochastic trafficking and redistribution of synaptic resources within a dendritic segment. We model this behavior as a stochastic process where a single synaptic resource variable is constantly redistributed among dendritic spines on the same dendritic segment, while the total amount of the resource is controlled in a homeostatic fashion. The constant redistribution gives rise to strong spine size fluctuations as observed experimentally. However, the time-averaged efficacy of individual synapses can still be systematically controlled through the rate at which a synapse captures the synaptic resource. Therefore, the post-synaptic expression of long-term plasticity in the model relies on modifying the rate at which synaptic resources are attracted into the synapse. The resulting model is able to learn stable associations despite fluctuating synaptic efficacies. Furthermore, it explains a wealth of data and makes some testable predictions. First, the model reproduces experimental data on spine size fluctuations and the distribution of synaptic efficacies. Second, the model captures the local heterosynaptic normalization of synaptic efficacies in response to long-term plasticity induction. Third, the model reproduces data on the global homeostatic scaling of spine sizes in response to altered activity levels. Finally, we show that inevitable variations in tuning properties due to the stochastic fluctuations of synaptic efficacies are greatly reduced by multi-synaptic contacts between pre- and postsynaptic neurons. Overall, the model explains how the formation of stable memories can be reconciled with the strong fluctuations of individual synaptic efficacies observed experimentally.

III-82. A new role for sparse expansion in neural networks

Madhu Advani[1] M.ADVANI81@GMAIL.COM
Julia Steinberg[1] JSTEINBERG@G.HARVARD.EDU
Haim Sompolinsky[2] HAIM@FIZ.HUJI.AC.IL

[1]Harvard University
[2]Hebrew University of Jerusalem

Multiple sensory pathways in the brain rely on sparsely active populations of neurons downstream from the input stimuli. The biological reason for the occurrence of expanded structure in the brain is unclear, but may be because expansion can increase the expressive power of a neural network: the universal approximation theorem proves a two layer network with a wide hidden layer can approximate any function. Here we consider whether expanding the hidden layer of a network can lead to improved learning even when this is not a factor, meaning the expansion does not permit the network to better express the underlying rule the data is generated from.

To study this setting we use a teacher-student framework where a perceptron teacher network generates labels which are corrupted with small amounts of noise. We then train a student network that is structurally matched to the teacher and can achieve optimal accuracy if given the teacher's synaptic weights. In this setting, adding hidden units with random connectivity to the input does not improve the capacity of the student network to match the teacher. Moreover, adding such units will increase the expressive capacity of the network leading to a risk of over-fitting the noise.

However, we show that by making the hidden units sparsely active, their addition can substantially improve learning in the student. We also demonstrate that ablation of the expanded layer after training the student enhances this effect, and provide theoretical justification as to why. This finding has implications for the design of artificial neural networks and for the understanding the role of neurogenesis and short-lived synapses for learning in biological neural networks.

III-83. A thalamocortical model solves sequencing tasks more robustly than backpropagation-trained RNNs

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

Columbia University

Training networks to produce long temporal sequences is notoriously difficult because of the issue of gradient instability plaguing the most popular learning algorithm: back propagation through time (BPTT). For sequences that can be decomposed into a series of shorter time-varying "motifs", one solution to this problem is to train a recurrent network to produce each motif in response to a corresponding static input. Then, (i) any number of motifs can be learned assuming the parameter updates necessary for training new motifs do not disrupt previously learned ones; and (ii) arbitrarily long sequences can be generated by switching between inputs, assuming that the network activity at the end of a motif can be used to start any subsequent one.

We identify networks architectures fulfilling the first requirement, and attempt to address the second by learning to generate individual motifs when initializing the network activity from an isotropic Gaussian. However, end-of-motif neural activities are not fully captured by this distribution, and we observe transition failures between motifs during sequence production. If all transition pairs are presented during training, sequence production succeeds, but the quadratic scaling with the number of motifs is prohibitive.

Gaining insight from the anatomy and physiology of the mammalian motor system, we rescue these issues by constructing a network model in which a motif-dependent thalamocortical loop drives the dynamics for each motif and motif-independent transition loops quickly prepare the network for any next motif. We analytically show that we can identify loop weights for the production of any motif and that transition loop weights can be tuned independently from the motifs involved. During sequence production, the performance is similar to that of BPTT-trained networks trained on all transition pairs while keeping the number of parameters fixed.

Our results emphasize that neuroscience insights can help designing more robust functional networks.

III-84. A formal neural synchrony model for unsupervised image grouping

Matt Ricci MATTHEW_RICCI_1@BROWN.EDU
Charles Windolf CHARLES_WINDOLF@ALUMNI.BROWN.EDU
Thomas Serre THOMAS_SERRE@BROWN.EDU

Brown University

The accuracy of modern deep convolutional neural networks (CNNs) is now approaching that of humans on certain visual recognition tasks while providing the best quantitative fit to visual cortex data. CNNs have thus rapidly become the de facto “standard model" of vision. Despite these successes, the most impressive CNN results have been achieved in a processing mode akin to rapid, pre-attentive, recognition, and a growing body of literature suggests that current CNNs remain outmatched by the brain. Recent work has shown that CNNs are limited in their ability to learn visual relations and neuroscience evidence indicates that feedback mechanisms play a critical role for solving visual reasoning tasks beyond image categorization. Yet, current models of feedback processes are typically more concerned with biological realism than utility in computer vision. How do we link these levels of analysis? Here, we propose a graphical model for unsupervised perceptual grouping based on an abstraction of neuronal synchrony using complex-valued units. In a series of experiments, we show how this generative neural network with complex-valued neurons representing both firing rate and phase can flexibly group simple scenes into a maximum of around four assemblies, close to the number reported in visual working memory studies. We additionally confirm the model is able to group novel objects as long as they obey certain gestalt rules, like good continuity. What's more, our network gives rise to a mathematical link between neuronal synchrony and statistical inference. We argue that the network's balance between biological plausibility and statistical rigor allows it both to model computations in cortex and scale to hard computer vision problems. To that end, we also propose a reinforcement learning extension of our architecture for visual reasoning.

III-85. Subcortical connectivity correlates selectively with choice bias in a visuospatial attention task

Varsha Sreenivasan VARSHAS@IISC.AC.IN
Devarajan Sridharan SRIDHAR@IISC.AC.IN

Indian Institute of Science

Attention enhances behavioral performance by either enhancing sensory processing of attended information (enhanced sensitivity) or by prioritized gating of attended information for guiding decisions (enhanced bias) [1]. While recent studies have investigated cortical correlates of these attention components [2-3] comparatively little is known about the involvement of subcortical regions in controlling sensitivity and bias [4]. Here, we combine model-based analysis of behavior with computational neuroimaging to examine the role of the superior colliculus (SC) – a subcortical midbrain structure – whose specific role in controlling attention remains a topic of active debate [5-9]

Human participants performed a multialternative visual change detection task with endogenous spatial cueing (Fig. 1A). We decoupled sensitivity and bias effects of cueing by modeling observers' decisions in a multidimensional decision space, with a recently developed signal detection model (Fig. 1B; [5,10]). We observed striking, complementary patterns of hemifield asymmetries in sensitivity and bias: sensitivity for detecting changes was consistently higher for the left hemifield (Fig. 1C, $p<0.01$), whereas bias for reporting changes was higher on the right hemifield (Fig. 1C, $p<0.05$). We sought to identify putative neural substrates for these behavioral asymmetries. We mapped white-matter fibers connecting the SC to multiple cortical regions using diffusion MRI and tractography [11] in two datasets comprising 82 subjects. We observed the converse pattern of asymmetries – significantly stronger connectivity of the SC with left, as compared to right, hemispheric cortical regions, consistently across both datasets (Fig. 2B). Finally, we found that asymmetry in SC connectivity correlated strongly with asymmetry in choice bias (Fig. 3A, $p<0.05$), but not sensitivity ($p>0.05$). Moreover, average SC connectivity correlated strongly with cueing-induced modulation of bias but not of sensitivity (Fig. 3B). Our results indicate a

distinct subcortical neural substrate for controlling choice bias during visuospatial attention.

III-86. Selective filtering of accumbal inputs by neuromodulators

Paul Hoerbelt HOERBELT@STANFORD.EDU
Daniel Christoffel DANC2@STANFORD.EDU
Jessica Walsh JWALSH2@STANFORD.EDU
Boris Heifets BHEIFETS@STANFORD.EDU
Karl Deisseroth DEISSERO@STANFORD.EDU
Charu Ramakrishnan CHARUR@STANFORD.EDU
Lili King-Adas LKINGADAS@SCU.EDU
Robert Malenka MALENKA@STANFORD.EDU

Stanford University

The neuromodulators dopamine (DA) and serotonin (5-HT) influence a wide range of motivated behaviors but the detailed mechanisms by which they modify critical neural circuits remain largely unknown. The nucleus accumbens (NAc) is a key node for reward processing and receives dense DA and 5-HT innervation. To determine whether DA and 5-HT selectively modulate synaptic transmission in the NAc in an input-specific manner, we examined how these neuromodulators influence excitatory postsynaptic currents (EPSCs) generated by four of the major glutamatergic inputs to the NAc. In mice in which D1-MSNs are labelled by tdTomato, we injected two different adeno-associated viruses each encoding one of two optogenetic constructs (blue light-activated Chronos or red light-activated ChrimsonR) into pairs of the major NAc input areas: the ventral hippocampus, the medial prefrontal cortex, the basolateral amygdala, or the periventricular thalamus. After 5-6 weeks of incubation, we made whole cell voltage clamp recordings from NAc MSNs in acute slices prepared from these animals while optogenetically activating two independent sets of inputs. EPSCs generated by these different inputs were selectively depressed by DA or 5-HT in an input-specific manner. Specifically, DA and 5-HT differentially regulated EPSCs generated by inputs from medial prefrontal cortex and periventricular thalamus compared to other inputs. These findings are recapitulated during drug-induced release of endogenous DA or 5-HT. Preliminary results suggest that the input-specific neuromodulator-evoked regulation of excitatory inputs occurs in either D1- or D2-MSN subpopulations. Moreover, consistent with previous work, measurement of EPSC paired pulse ratios suggests a presynaptic mechanism of neuromodulator action. Taken together, our findings suggest that DA and 5-HT selectively filter incoming excitatory input to the NAc. This input-specific modulation of excitatory synaptic transmission in the NAc may contribute to the discrete behavioral consequences of DA and 5-HT release in the NAc.

III-87. Inferring implicit inference

Rajkumar Vasudeva Raju[1] RV12@RICE.EDU
Zhe Li[2] ZHE.LI@BCM.EDU
Scott Linderman[3] SCOTT.LINDERMAN@COLUMBIA.EDU
Xaq Pitkow[2] XAQ@RICE.EDU

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

Repeating patterns of structure and function in the cerebral cortex motivate the idea that cortical circuits perform canonical computations. When neural representations are distributed, it remains an open challenge how to define canonical computations because the relevant operations are only indirectly related to single-neuron transformations. We present a new mathematical framework for inferring such canonical computations in distributed representations.

We focus on the hypothesis that the brain performs probabilistic inference via message-passing implemented by overlapping probabilistic population codes. Our general analysis method simultaneously finds (i) the neural representation of relevant latent variables, (ii) mapping from input to latent variable space, (iii) interactions between these decoded variables that define the brain's internal probabilistic graphical model of the world, and (iv) canonical message-functions that specify the implicit inference algorithm. These properties are statistically distinguishable due to the symmetries inherent in any canonical computation.

We applied our method to artificial neural recordings generated by a model brain that implicitly implements advanced mean-field inference. Given external inputs and noisy neural activity from the model brain, we successfully estimate the latent dynamics and canonical parameters that explain the simulated neural measurements. In this first example application, we used a simple polynomial basis to characterize canonical transformations. While this construction matched our model, it is unlikely to capture the brain's nonlinearities efficiently. In a companion submission, we apply a flexible variant of our method based on Graph Neural Networks to infer approximate inferences over circular variables (e.g. orientations) with known neural embeddings. We also apply the method to recordings from mouse visual cortex using naturalistic stimuli designed to highlight lateral connectivity, and identify population-level nearest-neighbor interactions. Finally, analysis of our models reveals certain features of experimental design needed to successfully extract canonical computations from neural data.

III-88. Dynamic multiplexing of modules in presynaptic neurons underlies feature encoding in cortical neurons

Arjun Bharioke[1] ARJUN.BHARIOKE@FMI.CH
Stuart Trenholm[2] STUART.TRENHOLM@MCGILL.CA
Adrian Wertz[1] ADRIAN.WERTZ@WEB.DE
Emilie Mace[1] EMILIE.MACE@FMI.CH
Daniel Hillier[1] DANIEL.HILLIER@FMI.CH
Josephine Juettner[1] JOSEPHINE.JUETTNER@FMI.CH
Zoltan Raics[1] ZOLTAN@RAICS.HU
Gergely Szalay[3] GSZALAY@FEMTONICS.EU
Balazs Rozsa[3] ROZSABAL@KOKI.HU
Georg Keller[1] GEORG.KELLER@FMI.CH
Karl Klaus Conzelmann[4] CONZELMANN@GENZENTRUM.LMU.DE
Keisuke Yonehara[5] KEISUKE.YONEHARA@DANDRITE.AU.DK
Rava Azeredo da Silveira[6] RAVA@ENS.FR
Botond Roska[7] BOTOND.ROSKA@IOB.CH

[1]Friedrich Miescher Institute for Biomedical Research
[2]McGill University
[3]Hungarian Academy of Sciences
[4]Ludwig-Maximilians-University
[5]The Danish Research Institute of Translational Neuroscience
[6]Ecole Normale Superieure
[7]Institute of Molecular and Clinical Ophthalmology Basel

As the seat of higher-order sensory processing, multi-sensory integration, decision-making and other complex computations, the neocortex has been extensively studied. Anatomically, cortical circuits are characterized by a high degree of convergence onto individual neurons, suggesting that each input neuron can only weakly affect output activity. Hence, analyses have focused on the computational role of populations of neurons. However, a detailed understanding of the specificity by which individual cortical input neurons collect into modules, that contribute to an output neuron's activity, is still unknown.

To understand how specific modules of input neurons are organized, and to differentiate such modules from the overall population of cortical neurons, we used single-cell-initiated monosynaptically-restricted rabies to express gCaMP6s transynaptically, in a single cortical neuron and its population of input neurons. We then used a rapid

random-access, acousto-optic-deflection 2-photon microscope to simultaneously record from both the output cell and its input population in vivo – sampling the response dynamics throughout the input network of a single cortical neuron, across different visual stimuli.

Repeating this analysis for multiple circuits, across different mice, we characterized both the stimulus response dynamics of individual neurons and their spontaneous activity. We demonstrated that similarities in presynaptic neurons' responses are not correlated across different stimulus domains. This generalizes to spontaneous responses, with neurons participating in different modules during different windows of activity. We observed that neurons within the presynaptic network show higher pairwise correlations, and greater synchrony of spontaneous responses, than randomly sampled cortical neurons. Modules of neurons also show greater correlation with the output neurons' firing than any individual cell alone, or all cells combined. Our results suggest that modules of presynaptic neurons provide dynamically varying inputs to the output neuron, with individual neurons participating in multiple different modules, providing a novel understanding of the functional connectivity of neurons within cortex.

III-89. Hierarchy of cortical population characteristic timescales inferred from field potentials

Richard Gao R.DG.GAO@GMAIL.COM
Bradley Voytek BVOYTEK@UCSD.EDU

University of California, San Diego

Neuronal populations across the macaque cortex exhibit intrinsic characteristic timescales in their spiking fluctuations. Specifically, sensory areas have shorter timescales of spiking autocorrelation (i.e., faster decay), while frontal areas have longer timescales [1]. Importantly, these regional differences predict working memory function, suggesting possible physiological mechanisms for how populations temporally maintain information [2]. Extending measurements of population timescales across many brain regions would considerably broaden our understanding of neurodynamics and how they support behavior. However, it is prohibitively expensive and technically challenging to record single units across the whole macaque cortex under different tasks, and impossible in the human brain due to the scarcity of single-unit recordings, even in rare clinical settings. Here we infer neuronal population timescales from field potentials recorded via electrocorticography (ECoG) grids across the macaque cortex. In frequency domain, the power spectral density (PSD) of field potentials often have a 'knee' [3] connecting two distinct power law-like regimes. The frequency at which this knee occurs corresponds mathematically to the time constant of autocorrelation decay in time domain. Applying a spectral parameterization tool we developed on 8 whole-cortex ECoG recordings from 2 monkeys [4,5], we find that ECoG timescales are tightly correlated with the reported spiking fluctuation timescales in select cortical areas from [1] (rho=0.884, p=0.008). We extend the analysis across cortex and find a U-shaped trend along the anterior-posterior axis, with the shortest timescales occurring in sensorimotor regions. Finally, we find that timescales decrease by half from eyes open to eyes closed resting. We validate the method with simulation of varying timescales, and discuss results pertaining to the scale-free nature of neuronal activity. In summary, this work offers a method for characterizing neuronal population timescales in ECoG, which can be extended to human datasets, as well as possibly M/EEG and field potential data in other model organisms.

III-90. Differences in the heading encoding circuit of insect species

Ioannis Pisokas[1] I.PISOKAS@SMS.ED.AC.UK
Stanley Heinze[2] STANLEY.HEINZE@BIOL.LU.SE
Barbara Webb[1] BWEBB@INF.ED.AC.UK

[1]University of Edinburgh
[2]Lund University

The ability of insects to keep track of their heading with respect to their environment is crucial for many behavioural tasks. A brain structure called the central complex plays an essential role in this heading encoding. Recent studies in Drosophila have shown that a group of neurons interconnecting two parts of the central complex, the ellipsoid body (EB) and the protocerebral bridge (PB), form a ring-attractor that maintains an activity "bump" that tracks the animal's heading. In this ring-attractor the excitation is mediated via a recurrent circuit formed by the E-PG, P-EN and P-EG neurons, while the $\Delta 7$ neurons mediate the inhibition. Similar, but not identical, circuitry is found in other insects, raising the question of how the connectivity details in this structure relate to function. We explore this question in computational models that accurately replicate the anatomical evidence from different species. In Drosophila, the inhibitory $\Delta 7$ neurons make connections uniformly in all columns of the PB. In other insects, e.g. the locust, the $\Delta 7$ neurons have synapses only in some columns of the PB. Additionally, the number of columns differs between species and the Drosophila EB forms a closed torus, while in the locust the EB is a linear structure prohibiting local connections between its two ends. Our hypothesis is that circuit differences match the ecological needs of the respective species. Our models show that the ring-attractor of Drosophila follows heading changes twice as fast as the locust circuit, matching Drosophila's ability to make fast turns. We suggest that the circuit identified in Drosophila is a specialisation of flies, possibly derived from a simpler circuit present in other insects. Our models predict that not all $\Delta 7$ neurons have the same activity level but the difference is smaller in Drosophila than in locust, and that the P-EG neurons are essential to the stability of the circuit.

III-91. A framework for linking computation and rhythm-based timing patterns in neural firing

E. Paxon Frady EPAXON@BERKELEY.EDU
Pentti Kanerva PKANERVA@CSLI.STANFORD.EDU
Friedrich Sommer FSOMMER@BERKELEY.EDU

University of California, Berkeley

A major challenge in neuroscience is to develop models that can relate observed neural firing patterns with computational functions. We present a two-level theory framework based on attractor networks in complex state space for understanding computations with rhythmic spike trains. This theory goes beyond traditional rate coding, and in particular takes into account spike timing codes, temporal sequences, oscillations, and synaptic delays. The first level describes computation by operations on high-dimensional complex representations. We describe a novel model of Hopfield-like attractor networks that can store and retrieve sparse patterns of complex phasors, referred to as *threshold phasor associative memory* (TPAM) networks. These networks store analog information in the phase of a unit magnitude complex number, and a threshold mechanism allows only neurons with high-confidence of their phase information to fire. We derived the energy function and show empirically how threshold phasor patterns can increase the capacity of an associative memory. The second level describes mechanisms how such operations can be mapped onto spiking neural networks. We describe how TPAM networks can be implemented with spiking integrate-and-fire neurons using a model of synaptic integration that implements algebraic operations on complex numbers in a biophysically plausible manner. Here, the timing of a spike relative to an internal oscillation represents the phase of a complex number. We show how the stable fixed-point attractor states of stored memories in the TPAM correspond to self-generated oscillations and stable limit-cycles of spiking activity in the spiking neural network. Further, we apply this modeling approach as a theory for hippocampus, and capture observed experimental phenomena, such as place-fields and phase-precession.

III-92. Dynamic updating of sensory history shapes perceptual judgments

Iacopo Hachen IHACHEN@SISSA.IT
Sebastian Reinartz REINARTZ.SEBASTIAN@GMAIL.COM
Romain Brasselet ROMAIN.BRASSELET@SISSA.IT
Mathew E Diamond DIAMOND@SISSA.IT

International School for Advanced Studies (SISSA)

The role of trial history on perceptual decision making is an increasingly debated topic in systems neuroscience. What are the exact dynamics by which past events affect ongoing perceptual processes? To gain insights into this issue, we study the role of stimulus history in shaping the current judgment, on time scales that range from days (the previous session) to seconds (the time span between trials). We designed a task in which rats had to judge a single stimulus as belonging to one of two categories, ‘high’ or ‘low’. Rats initiated each trial by placing their whiskers in contact with a plate which delivered a stochastic vibration. Stimuli differed in the average speed of the plate during vibration (9 possible intensities). After stimulus termination, an auditory cue instructed rats to select one of two spouts, which delivered fluid only after ‘correct’ responses. The central stimulus in the set corresponded to the category boundary, with choices rewarded randomly. To correctly classify stimuli, rats needed to create an ‘internal boundary’ based on previous experience of task contingencies. We found a systematic shift in animals’ psychometric functions in the direction of past stimulus values. That is, after a ‘low’ speed stimulus rats were more likely to report the next stimulus as ‘high’, and vice versa, independently of the previous choice, and proportionally to the past speed value. Further analysis revealed that this influence lasted for sequences of trials, decaying exponentially throughout the session. Importantly, we studied the effect’s dynamics at different timescales. Based on our data, we provide simple statistical models (easily testable, behaviorally and physiologically) that can account for the effect of both sensory features and timing of past stimuli. Overall, this work supports the view of perception as an active and continuous integration process of information about the environment, highlighting its functionality and limitations.

III-93. Learning a reward distribution with reward prediction errors in a circuit model of the Drosophila mushroom body

James Bennett JAMES.BENNETT@SUSSEX.AC.UK
Thomas Nowotny T.NOWOTNY@SUSSEX.AC.UK

University of Sussex

Drosophila exhibit matching behaviour, whereby their relative preferences for different foods determine the frequencies with which they visit the respective food sources (Beshel & Zhong, 2013). Key to this behavior is the ability to learn and accurately evaluate the valence of different options, and to associate those valences with appropriate sensory cues. An important site of learning in Drosophila is the mushroom body (MB) (Congnigni et al., 2018). There, sensory cues are encoded by Kenyon cells (KCs), which in turn drive MB output neurons (MBONs) that bias either approach or avoidance behaviors (Aso et at., 2014). MBON activity can thus be interpreted as encoding the learned valence, or a reward prediction. Valence memories are learned via reward modulated synaptic plasticity, and are stored in KC-MBON synapses, at which rewards are signalled by dopaminergic neurons (DANs). DANs that signal positive rewards tend to modulate MBONs that bias avoidance, whereas punishment signalling DANs modulate approach biasing MBONs. Here, we investigate the hypothesis that DANs receive a copy of the MBON reward prediction (Aso et al., 2014; Felsenberg et al., 2017, 2018) in addition to the actual reward signal, and thus compute reward prediction errors (RPEs). We present a MB circuit model that captures these data and, assuming that the objective of the MB is to minimise RPEs, derive a learning rule that yields a neuronal implementation of the Rescorla-Wagner model (Rescorla & Wagner, 1972). However, under this objective, the valence-specific circuity of the MB is unable to learn. We therefore propose a modified MB circuit and corresponding plasticity rule that is not only able to learn, but captures additional features of MB physiology. Finally, we propose a role for the B’1 MBONs in learning the reward variability, and suggest how this information may be utilised to incorporate risk into choice evaluations.

III-94. Sparse coding protects against adversarial attacks

Dylan Paiton DPAITON@BERKELEY.EDU
Joel Bowen JOEL_BOWEN@BERKELEY.EDU
Jasmine Collins JAZZIE@BERKELEY.EDU
Charles Frye CFRYE@BERKELEY.EDU
Alex Terekhov AVTEREKHOV@GMAIL.COM
Bruno Olshausen BAOLSHAUSEN@BERKELEY.EDU

University of California, Berkeley

Lateral and feedback connections abound in biological neural networks [1], but are absent from the feedforward, convolutional architectures popular in artificial neural networks. Artificial neural networks also exhibit adversarial examples [2]: small perturbations to the input can produce large changes to the output. Recent work [3] has identified adversarial examples that transfer to human observers, but only when those observers are time-limited. This suggests that the slower lateral and feedback computations help protect biological neural networks from adversarial examples. In this work, we demonstrate that a classifier with lateral connectivity in its hidden layer is more robust to adversarial examples than a feedforward classifier, requiring larger input perturbations to achieve the same change in output. The hidden layer uses the Locally-Competitive Algorithm (LCA) [4], a leaky integrator network implementation of sparse coding [5]. We hypothesize that the curved geometry of neuron responses with non-feedforward connections [6] improves selectivity and provides this protection. We analyze this geometry in the context of adversarial examples.

III-95. Multimodal integration in L5 pyramidal neurons of secondary visual cortex

Alessandro Galloni ALESSANDRO.GALLONI@CRICK.AC.UK
Aeron Laffere AERON.LAFFERE@CRICK.AC.UK
Ede Rancz EDE.RANCZ@CRICK.AC.UK

The Francis Crick Institute

Sensory processing in the brain involves integration of both feedforward sensory information and multimodal feedback signals. The function of cortical feedback is still poorly understood and may convey information about attention, context, and sensory predictions. Thick-tufted layer 5 (ttL5) pyramidal neurons are particularly well-placed for integrating these diverse inputs. Their extensive dendrites receive thousands of local and long-range inputs and show nonlinearities like calcium plateau potentials, endowing them with unique computational properties. Here, we explore the mechanisms of top-down and bottom-up integration by describing a genetically targeted population of ttL5 neurons in the mouse medial secondary visual cortex (V2m). Using monosynaptic retrograde tracing, we find that this population receives strong long-range input from L5 of primary visual cortex (V1) and retrosplenial cortex (RSC). To map the precise dendritic targets of these connections, we use subcellular optogenetic circuit mapping (sCRACM) and find that feedforward inputs from V1 surprisingly target mainly apical and tuft dendrites in V2m, whereas inputs from RSC, considered as feedback, target mainly basal dendrites. To determine how these inputs functionally interact, we evoke EPSPs from each pathway and show that they sum linearly across a wide range of time intervals. We also directly stimulate the soma and synaptic input to layer 1, which normally evokes calcium plateau-mediated spike bursts when stimulated simultaneously. Surprisingly, we find that unlike previously described ttL5 neurons in V1 and somatosensory cortex, ttL5 neurons in V2m generally do not display such bursts or other hallmarks of plateau potentials. We explore the conditions required for linear and non-linear summation further in a biophysically accurate ttL5 neuron model. These results argue against the classical notion that sensory inputs target the basal compartment and top-down inputs target the apical tuft, and demonstrate that ttL5 properties differ depending on brain region.

III-96. Functional interactions between Vip and Sst interneurons are stimulus and state dependent

Daniel Millman DANIELM@ALLENINSTITUTE.ORG
Gabriel Ocker GABEO@ALLENINSTITUTE.ORG
Nicholas Cain NICHOLASC@ALLENINSTITUTE.ORG
Peter Ledochowitsch PETERL@ALLENINSTITUTE.ORG
Rylan Larsen RYLANL@ALLENINSTITUTE.ORG
Kevin Takasaki KEVINT@ALLENINSTITUTE.ORG
Michael Oliver MICHAELO@ALLENINSTITUTE.ORG
Clay Reid CLAYR@ALLENINSTITUTE.ORG
Michael A Buice MICHAELBU@ALLENINSTITUTE.ORG
Saskia de Vries SASKIAD@ALLENINSTITUTE.ORG

Allen Institute for Brain Science

In cortex, vasoactive intestinal peptide-expressing (Vip) interneurons make strong inhibitory synapses onto somatostatin-expressing (Sst) interneurons and, reciprocally, Sst cells make inhibitory synapses onto Vip cells. This reciprocal connectivity establishes a circuit motif of competitive inhibition with poorly understood effects on cortical processing. Here, we investigate the visual stimulus tuning and behavioral state-dependent modulation of these two major classes of cortical interneurons in mouse visual cortex. First, analysis of data from the Allen Brain Observatory reveals signatures of functional interactions between Vip and Sst cells that are highly dependent on stimulus statistics (i.e. natural versus artificial stimuli) and differ across visual areas. In V1 and LM, but not in PM, full-field gratings drive strong and reliable responses in Sst cells and strongly suppress activity of Vip cells. This is consistent with direct inhibition from Sst onto Vip cells. Furthermore, we confirm previous reports that Vip cells in V1 are far more predictive of running behavior than cells labelled by other Cre-lines, but do not find this effect in higher visual areas (LM and PM). Surprisingly, the prevalent model of mutual inhibition appears to be stimulus specific as Vip and Sst cells are both highly responsive to natural stimuli and their activity is not anti-correlated. To directly observe the functional interactions between Vip and Sst cells, we established an experimental approach to express GCamp in both Vip and Sst cells in the same animal and simultaneously image populations of 50-100 of these sparsely-labeled neurons at high temporal resolution (30 Hz) using a Bessel scope. We find ensembles of coactive cells with distinct stimulus tuning and running modulation. Overall, our results suggest that the functional interaction between these cell types is both stimulus and state-dependent, and does not necessarily generalize across areas of visual cortex.

III-97. A mechanism for generating modular activity in the neocortex

Bettina Hein[1] HEIN@FIAS.UNI-FRANKFURT.DE
Dmitry Bibichkov[1] BIBICHKOV@FIAS.UNI-FRANKFURT.DE
Gordon B Smith[2] GBSMITH@UMN.EDU
David E Whitney[3] DAVID.WHITNEY@MPFI.ORG
David Fitzpatrick[3] DAVID.FITZPATRICK@MPFI.ORG
Matthias Kaschube[1] KASCHUBE@FIAS.UNI-FRANKFURT.DE

[1]Frankfurt Institute for Advanced Studies
[2]University of Minnesota
[3]Max Planck Florida Institute for Neuroscience

Cortical activity patterns in primates and carnivores typically comprise domains of coactive neurons with an approximately regular spacing on the order of 1mm. Such modular organization is seen in spontaneous activity early in the developing cortex and it persists even after deactivation of LGN suggesting a cortical origin [Smith et al., Nat Neurosci 2018]. However, a biologically plausible and robust mechanism for the intra-cortical generation of modular activity is still missing. Previous network models generating such activity patterns either assume short-range excitation and longer-ranging inhibition (Mexican hat connectivity) [Ermentrout, ROPP 1998] or rely on fast inhibition [Pinto & Ermentrout, SIAM 2001; Kang et al., PNAS 2003], but these assumptions appear in conflict with current experimental evidence. Here we show that spatially localized self-inhibition, consistent with the observed prevalent autapses of fast-spiking interneurons in cat visual cortex [Tamas et al., J Neurosci 1997], relaxes the constraints in the above models, such that biologically more plausible network motifs with shorter ranging inhibition than excitation robustly generate modular activity. Considering a network of excitatory and inhibitory rate units we use linear stability analysis to determine the boundaries of the parameter regime for the formation of modular activity and show that the size of the regime consistent with experimental evidence increases with the strength of local inhibition. A critical prediction of our model is the decrease in spacing of active domains when the total amount of inhibition increases. Indeed, consistent with an increase of inhibition in the developing visual cortex [White & Fitzpatrick, Neuron 2007], we observe that the spacing of spontaneously active domains decreases with age in early ferret visual cortex. We conclude that cortical recurrent networks involving local self-inhibition can robustly generate modular activity patterns.

III-98. Sparse-coding variational auto-encoders

Gregory Barello[1] GBARELLO@UOREGON.EDU
Adam Charles[2] ADAMSC@PRINCETON.EDU
Jonathan Pillow[2] PILLOW@PRINCETON.EDU

[1]University of Oregon
[2]Princeton University

The sparse coding model posits that the visual system has evolved to efficiently encode natural stimuli using a sparse set of features from an overcomplete dictionary. The classic sparse coding model suffers from two key limitations, however: (1) computing the neural response to an image patch requires minimizing a nonlinear objective function, which was initially not easily mapped onto a neurally plausible feedforward mechanism and (2) fitting the model to data relied on an approximate inference method that ignores uncertainty. Here we address these two shortcomings by formulating a variational inference method for the sparse coding model inspired by the variational auto-encoder (VAE) framework. The sparse-coding variational auto-encoder (SVAE) augments the classic sparse coding model with a probabilistic recognition model, parametrized by a deep neural network. This recognition model provides a neurally plausible implementation for the mapping from image patches to neural activities, and enables a principled method for fitting the sparse coding model to data via maximization of the evidence lower bound (ELBO). The SVAE differs from the traditional VAE in three important ways: the generative model is the sparse coding model instead of a deep network; the latent representation is overcomplete, with more latent dimensions than image pixels; and the prior over latent variables is a sparse or heavy-tailed instead of Gaussian. We fit the SVAE to natural image data under different assumed prior distributions, and show that it obtains higher test performance than previous fitting methods. Finally, we examine the response properties of the recognition network and show that it captures important nonlinear properties of neurons in the early visual pathway.

III-99. Detecting and correcting false transients in calcium imaging

Jeff Gauthier JEFF.L.GAUTHIER@GMAIL.COM
Sue Ann Koay SUEANNKOAY@GMAIL.COM
Edward Nieh ENIEH@PRINCETON.EDU
David W Tank DWTANK@PRINCETON.EDU
Jonathan Pillow PILLOW@PRINCETON.EDU
Adam Charles ADAMSHCH@GMAIL.COM

Princeton University

Population recordings of calcium activity are a major source of insight into neural function. Large dataset sizes often require automated methods, but these can introduce errors that are difficult to detect and have the potential to distort scientific findings. Here we show that automatic time course estimation can sometimes lead to significant misattribution errors, in which fluorescence is ascribed to the wrong cell. Misattribution arises when the shapes of overlapping cells are imperfectly defined, or when entire cells or processes are not identified. Importantly, misattribution can even be produced by methods specifically designed to handle overlap. Current evaluation metrics are not able to detect or remove these cases of false activity. We developed a novel procedure to identify misattribution at the level of individual transients, and found substantial contamination, particularly in a tissue with densely packed and overlapping sources. To filter out this erroneous activity, we designed a robust estimator that explicitly accounts for contaminating signals in a generative model. Using manual human classification of activity as ground truth, we show that our method achieves near-perfect performance at correcting misattribution. Our methods are post hoc, meaning they can be combined with essentially any cell finding technique, thus empowering users to diagnose and correct at large scale a problem that has the potential to significantly alter scientific conclusions.

III-100. Avoidance, extinction and dopamine-guided learning in Drosophila melanogaster

Bayar Menzat BAYAR.MENZAT@CNCB.OX.AC.UK
Johannes Felsenberg JOHANNES.FELSENBERG@DPAG.OX.AC.UK
Richard Lonsdale RICHARD.LONSDALE3@BTINTERNET.COM
Scott Waddell SCOTT.WADDELL@CNCB.OX.AC.UK
Tim P Vogels TIM.VOGELS@CNCB.OX.AC.UK

University of Oxford

Drosophila Melanagoster can form sparse representations of various odours and learn odour-specific associations such as the history of the positive or negative experiences, in a brain area called the mushroom body. Here, we introduce a four-layer spiking model of the fruit fly's olfactory system to investigate the mechanisms behind sparse odour representation and associative learning. We tune our model to elicit realistic firing activity and show that the interplay between multiple neural mechanisms (lateral inhibition in the antennal lobe and feedback inhibition in the mushroom body) produces temporal and spatial sparseness of the odour representation. We test the robustness of our model with a data set of the first layer responses to 110 different odours and show that our model network encodes sparse and decorrelated odour representations from the presented data set. To study the plasticity mechanism behind aversive learning and extinction (odour w/o expected shock) as seen in recent experimental results (Felsenberg et al., Cell, 2018), we introduce a dopamine-modulated learning rule in which high levels of dopaminergic activity paired with an odour trigger depression of KC-MBON synapses. We show that a depression-only learning rule can train all of the odours in our dataset and classify them as either naive, aversive or extinguished, according to context. We then add a mechanism of strengthening KC-MBON synapses at low levels of dopaminergic activity to enhance the fly's ability to recover the naiive state and re-learn odours after extinction. Our model predicts that silencing appetitive DANs required for extinction learning during odour exposure would create an aversive odour association. Our spiking model of the fly's brain provides a mechanistic explanation of how aversive odour associations can be acquired, extinguished and re-learned.

III-101. Tracking attention fluctuations in the human brain with an EEG-based cognitive brain-machine interface

Abhijit Chinchani[1] ABHIJITCHINCHANI@GMAIL.COM
Suhas Ganesh[2] SUHASG10495@GMAIL.COM
Byron Yu[2] BYRONYU@CMU.EDU
Devarajan Sridharan[1] SRIDHAR@IISC.AC.IN

[1]Indian Institute of Science
[2]Carnegie Mellon University

Visual attention improves perceptual accuracy and reaction times for attended stimuli, and also enhances the activity (gain) of visual neurons. Previous studies have shown that momentary fluctuations in attention can be measured by tracking neural activity in primate visual areas (e.g. V4) (1, 2, 3); such fluctuations correlate with animals' perceptual accuracy on a trial-by-trial basis. Yet, the link between such momentary neural fluctuations and attention state remains to be shown in the human brain. We investigate this link using a real-time cognitive brain machine interface (Fig. 1A), based on steady state visually evoked potentials (SSVEPs): occipital EEG potentials evoked by periodically flickering stimuli, whose power systematically modulates with attention (4). Human subjects (n=12) viewed two flickering pedestals, each flickering at a distinct frequency (Fig. 1B), and were cued to attend to one side (target). With dimensionality reduction (5, 6) we identified two latent dimensions, each quantifying SSVEP power evoked by one of the two pedestals (Fig. 1C). On a trial-by-trial basis, the presentation of behaviorally-relevant stimuli (oriented gratings) was triggered when power in the target or distractor dimension reached predetermined threshold values (Fig. 1D). Subjects detected and reported the orientation of the grating on the cued side. Orientation discrimination accuracy (d') was significantly higher on trials in which the grating presentation occurred when power was momentarily high in the target dimension (p=0.008, ANOVA); in contrast, no differences occurred in other metrics like reaction times or decision criteria (Fig. 2A). Moreover, when grating presentation was triggered based on power in the distractor dimension (n=7) no systematic behavioral effects were observed (Fig. 2B). Our results demonstrate a direct link between attentional effects on perceptual accuracy and neural gain in the human brain. Additionally, attention's effects on behavioral accuracies and reaction times may engage distinct neural mechanisms.

III-102. VIP disinhibition: interneuron contributions to state-dependent cortical computation

Katie Ferguson KATIE.FERGUSON@YALE.EDU
Calvin Fang CALVIN.FANG@YALE.EDU
Jessica Cardin JESS.CARDIN@YALE.EDU

Yale University

Visual information from the environment is rapidly and dynamically processed by a complex network of cortical neurons. Cortical activity patterns reflect not only changing sensory inputs, but also behavioral state (e.g., quiet wakefulness vs. active locomotion). GABAergic interneurons expressing vasoactive intestinal peptide (VIP-INs) are strongly activated by neuromodulatory afferents during distinct behavioral states, and are thus uniquely situated to be key contributors to flexible cortical function. Research suggests that VIP-INs predominantly inhibit somatostatin expressing interneurons (SST-INs), leading to a disinhibition of pyramidal neurons (PNs) and an increase in overall network activity??. However this disinhibitory model remains incomplete, as recent research proposes a role for VIP-INs in direct inhibition of local PNs. The relative impact of these highly non-linear cortical interactions may vary with behavioral or environmental context??, and thus VIP-INs may have a complex and largely unexplored role in the regulation of cortical activity.

To assess the contribution of VIP-INs to arousal-mediated changes in cortical processing, we expressed the genetic calcium indicator gCaMP6 in VIP-INs and, using intersectional genetic tools, in their postsynaptic targets

SST-INs and PNs. We used large-scale two-photon dual population imaging and high-throughput analyses to assess simultaneous patterns of VIP-IN and SST-IN (or PN) activity in the primary visual cortex of awake mice across behavioral states and cortical layers. We then ablated VIP-INs by Cre-dependent expression of a genetically targeted toxin, Caspase-3, and recorded from their postsynaptic targets, SST-INs and PNs, using gCaMP6.

Although SST-INs in V1 showed robust responses to visual stimuli in healthy mice, these responses were diminished in mice following VIP-IN ablation. Instead, loss of VIP-INs revealed locomotion-dependent activity in SST-INs and reduced activity in PNs. These results uncover underlying state-dependent excitatory inputs to SST-INs and suggest a complex role for VIP-INs in regulating large-scale changes to sensory processing across behavioral states.

III-103. Divisive normalization in a population of coupled GLM neurons

Lukasz Kusmierz[1] NALEWKOZ@GMAIL.COM
Jonathan Pillow[2] PILLOW@PRINCETON.EDU

[1]RIKEN
[2]Princeton University

Divisive normalization is a ubiquitous phenomenon in neural data and is widely regarded as a canonical cortical computation. However, the question of how divisive normalization can be implemented in a network of spiking neurons remains an open problem. Here we analyze the ability of networks of coupled Poisson generalized linear model (GLM) neurons to carry out divisive normalization. We show that a network of Poisson GLM neurons with inhibitory coupling between them can perform a nonlinear form of divisive normalization. Moreover, this normalization can be made linear if the inhibition is mediated by a single auxiliary neuron. Importantly, neurons do not need to have access to unnormalized activities and do not have to have specialized normalizing inputs. Using synthetic data, we confirm that GLMs can be fitted to spike train data exhibiting divisive normalization, providing an efficient tool that can help to uncover brain circuits implementing the divisive normalization. This adds an important population-level behavior the extensive repertoire of single-neuron dynamical behaviors the GLM has been shown to be able to express.

III-104. The translation invariant bispectrum for feature analysis in complex cells

Sophia Sanborn SANBORN@BERKELEY.EDU
Ryan Zarcone ZARCONE@BERKELEY.EDU
Bruno Olshausen BAOLSHAUSEN@BERKELEY.EDU
Christopher Hillar CHILLAR@MSRI.ORG

University of California, Berkeley

We present a method for recovering the shape selectivity of complex cells (or other invariant neurons) via reverse correlation in the bispectral domain. Invariant representations of stimuli, such as the Fourier power spectrum, have been of great importance in the analysis of natural signals. However, the power spectrum has fundamental limitations, which we explicate here. We emphasize the need for a complete set of invariants, for which all points in a transformation orbit map to the same invariants, and distinct orbits map to distinct points. We advocate determining a complete sets of invariants with the machinery of the bispectrum, which can be effectively computed and inverted using the concrete representation theory of the underlying transformation group. As a proof-of-concept for the approach, we present a method for inferring invariant stimulus features for the special case of two-dimensional translation, demonstrating the potential of the 2D bispectrum for understanding the response properties of complex cells in visual cortex.

III-105. Identity domains in complex behavior: Toward a biology of personality

Oren Forkosh[1] OREN.FORKOSH@MAIL.HUJI.AC.IL
Stoyo Karamihalev[2] STOYO_KARAMIHALEV@PSYCH.MPG.DE
Alon Chen[3,4] ALON_CHEN@PSYCH.MPG.DE

[1]Hebrew University of Jerusalem
[2]Max Planck Institute of Psychiatry
[3]Max Planck Institute of psychiatry
[4]Weizmann Institute of Science

Personality traits offer considerable insight into the biological bases of individual differences. However, existing approaches toward understanding personality across species rely on subjective criteria and limited sets of behavioral readouts, resulting in noisy and often inconsistent outcomes. Here, we introduce a mathematical framework for studying individual differences along traits with maximum consistency and discriminative power. We validate this framework in mice, using data from a system for high-throughput longitudinal monitoring of group-housed mice that yields a variety of readouts from all across an individual's behavioral repertoire. We found three distinct behavioral strategies that emerge from the data, and most mice possess a mixture of these strategies. We describe a set of stable characteristics that capture variability in behavior and gene expression in the brain, allowing for better informed mechanistic investigations into the biology of individual differences.

III-106. Cell assembly model for retinal ganglion cell population

Chris Warner CWARNER@BERKELEY.EDU
Friedrich Sommer FSOMMER@BERKELEY.EDU

University of California, Berkeley

We use a Noisy-Or model (Heckerman, 1990) to learn latent structure in observed spiking patterns from a population of rat retinal ganglion cells in response to white noise and natural movie stimulus. The probabalistic generative model, fit with the expectation maximization algorithm, captures approximately repeated groups of neurons often coactive together without requiring exact repeats. These cell assemblies capture the fine-temporal correlations in population spike trains that differ across stimulus conditions but are not necessarily stimulus locked. We discuss the spatial structure of cells in cell assemblies and what features of the stimulus is represented by a cell assembly above and beyond the union of the representations of its member cells.

III-107. Confirmation bias in active learning

Yul Kang[1] YUL.KANG.ON@GMAIL.COM
Daniel Wolpert[2] WOLPERT@COLUMBIA.EDU
Mate Lengyel[1] M.LENGYEL@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Columbia University

When gathering information, di?fferent sources typically have distinct levels of informativeness. Thus, we need to actively select the source of information from which to learn (i.e., perform "active learning"). There is conflicting evidence as to whether humans use e?fficient active learning strategies or instead use the simple heuristic of seeking information that only confirms their beliefs ("positive testing" or "confirmation bias"). Deciding which strategy dominates in human active learning requires four elements: 1. an experimental paradigm in which both strategies could in principle be employed, 2. principled statistical methods to estimate subjects' beliefs trial-by-trial (which form the basis of both optimal active learning and positive testing strategies), 3. mathematical models

to formalize each strategy, and 4. a data-driven characterization of the influence of each strategy on behavior. We combined these elements in a perceptual decision-making task. On each trial, subjects categorized a target animal (stick drawing) into one of three species, given prior experience with the distribution of animals in each species. Critically, on some trials, they had the opportunity to sample another example of one of the species of either their own choosing or chosen for them randomly. This allowed us both to assess the informativeness of each new sample as a function of the target stimulus (contrasting the random condition with baseline), and to test subjects' strategy in selecting samples (in the active choice condition). By fi?tting ideal observer models to categorization behavior, we estimated subjects' beliefs and predicted their sampling choices given either optimal active learning, positive testing, or a combination of both. We found positive testing strongly dominated in all subjects, and was only weakly modulated by an optimal active learning strategy. Thus, in contrast to eye movements where near-optimal active learning has been demonstrated, overt choice behavior is severely suboptimal.

III-108. The roots of the stabilized supralinear network in the asynchronous irregular regime of cortex

Takafumi Arakaki TARAKAKI@UOREGON.EDU
Yashar Ahmadian YASHAR@UOREGON.EDU

University of Oregon

Understanding the fundamental operations of the cerebral cortex is a central goal of neuroscience. The Stabilized Supralinear Network (SSN) is a simple rate model of cortical circuitry that accounts for a range of nonlinear cortical computations. These include sub-additive multi-input integration, known as normalization, and the saturating or supersaturating contrast-response function of visual cortical neurons. Importantly, the SSN further predicts a transition from sub-additive to super-additive multi-stimulus integration when stimuli become sufficiently weak. A key component of the SSN is the power-law input-output neural transfer function, which is known to result from fluctuation-driven spiking that emerges in the asynchronous irregular (AI) regime of cortical activity. A quantitative linking of SSN to AI dynamics in spiking networks of excitatory and inhibitory neurons which self-consistently accounts for the dynamically emerging power-law transfer function is nevertheless lacking. Previous theoretical mean-field studies of the AI regime in randomly connected spiking networks focused on the limit in which the number, K, of presynaptic neurons tends to infinity. This approximation yields the so-called balanced solution which predicts a linear dependence of population responses on inputs, missing the observed cortical nonlinearities. Here, we show that spiking neural networks with stimulus selectivity, statistically structured connectivity, and a K in the biological range, reproduce all the basic nonlinearities of SSN. Specifically, we show that structured spiking networks in the AI regime also exhibit supersaturation of excitatory responses and a transition from super-additive to sub-additive integration with growing stimulus strength. Finally, to bridge the theoretical gap between the SSN and mean-field theories of the AI state, we introduce an extension of the original SSN, which we call sigma-SSN. This model self-consistently accounts for the strength of dynamically emerging input fluctuations and provides an approximation that parallels the traditional mean-field theory when the network operates in the biologically relevant regime.

III-109. Tree-structured locally linear dynamics model to uproot Bayesian neural data analysis

Josue Nassar[1] JOSUE.NASSAR@STONYBROOK.EDU
Scott Linderman[2] SCOTT.LINDERMAN@COLUMBIA.EDU
Memming Park[1] MEMMING.PARK@STONYBROOK.EDU
Monica Bugallo[1] MONICA.BUGALLO@STONYBROOK.EDU

[1]Stony Brook University
[2]Columbia University

Statistical neuroscience has been pushing the boundary of what can be inferred bottom-up from neural data. As higher-dimensional population recordings become available, the question arises of whether we can learn the underlying dynamical law of population dynamics solely based on the recorded neural activities. While modern machine learning techniques have been successful in learning latent dynamical models of neural populations, these techniques often trade interpretability for predictive power. These methods essentially become “black boxes,” making it difficult for neuroscientists to decipher their inner workings and understand the computational principles implemented by neural systems. In contrast, linear dynamical systems (LDS) are very interpretable, but sadly, the limited expressive power of LDS can only capture trivial neural dynamics. To increase the flexibility of the LDS, we partition the underlying latent space into multiple regions, each with linear dynamics. This combination of locally linear dynamics can represent a much richer class of dynamical systems while retaining interpretability. We call our method the *tree-structured recurrent switching linear dynamical system* (TrSLDS) because it partitions the latent space using a novel tree-structure and leverages a hierarchical LDS prior to enforce structural smoothness of the dynamics. We propose a computationally efficient, fully Bayesian inference algorithm to infer the model parameters. The learned model is a binary tree where each node represents a spatially constrained LDS, providing a multi-scale view of the full dynamics with increasing complexity for deeper levels of the tree. We apply the TrSLDS to a spiking neural network with winner-take-all dynamics. Using a small subset as observed neurons, our method recovers an effective 2-dimensional dynamics that recapitulate the theoretically derived reduction of the high-dimensional spiking neural network model. This demonstrates how our method can be used to bring theoretical understanding of population dynamics in a purely data-driven manner.

III-110. High-dimensional filtering supports context-dependent neural integration

Jonathan Gill[1] JVG219@NYU.EDU
Kelsey Allen[2] KRALLEN@MIT.EDU
Alexander Williams[3] ALEX.H.WILLIA@GMAIL.COM
Mark Goldman[4] MSGOLDMAN@UCDAVIS.EDU

[1]New York University
[2]Massachusetts Institute of Technology
[3]Stanford University
[4]University of California, Davis

Humans and other animals are capable of using sensory cues to guide flexible, context-dependent behavior. A key question is how context-relevant cues may be used to guide behavior while context-irrelevant cues are ignored. Here we consider this question using a decision-making task in which an animal must temporally accumulate (i.e. integrate) a relevant sensory signal while ignoring an irrelevant stimulus. Previous theories proposed that neural responses to irrelevant inputs are actively suppressed upstream, preventing transmission to decision-making circuits. However, this idea is inconsistent with neural recordings demonstrating that both relevant and irrelevant signals are represented at the level of decision-making circuits, such as the prefrontal cortex (PFC). Here we show that, in large, high dimensional networks, irrelevant signals are 1) expected to be present in neural recordings, but 2) generically expected to have little effect on network computation. This seemingly paradoxical result derives from the fact that, in high dimensions, any two randomly selected vectors are orthogonal. Therefore, during accumulation of evidence, sensory inputs are by default orthogonal to the direction of integration (integrating eigenvector) and therefore enter the network and influence activity, but are not integrated. We demonstrate this principle of “high-dimensional filtering” in a network of mixed-selectivity sensory neurons subject to contextual gain modulation, that project onto an integrating network. We show that, with appropriately tuned connection strengths from the sensory to the decision-making (integrator) network, the relevant signals are selected in each context by rotating the sensory network output slightly towards the integrating mode. As the network grows, it becomes more robust to irrelevant inputs, requiring smaller rotations of the relevant input for integration. These results reveal a general principle for how large biological networks may filter out irrelevant inputs, while using

contextual signals to subtly rotate neural responses for sensory selection.

III-111. Geometric properties of face representationsin the macaque face patch system and deep neural networks

Uri Cohen[1] URI.CO@MAIL.HUJI.AC.IL
SueYeon Chung[2] SUEYEON@MIT.EDU
Daniel Lee DDL46@CORNELL.EDU
Winrich Freiwald WFREIWALD@MAIL.ROCKEFELLER.EDU
Doris Tsao[3] DORTSAO@CALTECH.EDU
Haim Sompolinsky[1] HAIM@FIZ.HUJI.AC.IL

[1]Hebrew University of Jerusalem
[2]Massachusetts Institute of Technology
[3]California Institute of Technology

The response of a population of sensory neurons to an object may vary significantly as it is presented under different conditions (such as distance, pose, background, etc.). When multiple objects are represented by a neural population, each of them is associated with an “object manifold”, defined in the neural response space. The ability of a neural population to support object classification is quantified through “classification capacity” and a rich theoretical framework connects this capacity with the geometry of the object manifolds. The theory suggests per-manifold measures of radius and dimension, as well as between-manifolds correlations of either centers or axes of variation. Here we apply the theory to analyze the geometry of “face manifolds” created at different stages of the macaque face patches system, in response to faces at different orientations; those are compared to the manifolds created when the same stimuli is processes by a deep convolutional neural network trained to perform face recognition tasks. Our analysis suggests that the face patches representations form a hierarchy from ML/MF through AL to AM, along which the face manifold radii and dimensions decrease and the face classification capacity increases. Propagating the same face stimuli through the deep network shows such trend in the last three stages, consisting of rectified fully-connected linear transformations. Additionally, in both cases the ability to estimate the stimulus’s pose is degraded along those stages, suggesting the manifold extent is reduced by suppressing pose-related variability. A notable difference between the cortical face patches representations and those of the network layers is the high degree of alignment of the axes-of-variation between manifolds of different faces which is evident in the biological manifolds and absent in the artificial ones. Both the mechanistic origin and computational implication of those axes-axes correlations are currently explored.

III-112. A rep of time that minimizes TDRL’s valuation error explains observed temporal decision making

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

Johns Hopkins University

By what neural mechanisms do animals represent the passage of time, learn temporal patterns, and decide how to spend time? Is temporal decision making consistent with algorithms hypothesized from other domains? How do these representations relate to evolutionary pressure to maximize reward accumulation? We investigated how evaluations assigned by temporal difference reinforcement learning (TDRL) [1] relate to decisions about how to spend time. We derive the general solution for how to optimize reward accumulation and prove that memoryless TDRL evaluations (infinite sums of exponentially discounted future rewards) systematically fail to achieve this optimization. However, this failure can be best mitigated by representing time using a time-dilating state space, wherein the amount of time spent in subsequent state increases by a precise proportion. We show

that this representation has been previously neurally evidenced [2]. TDRL applied to a time-dilating state space explains the diverse suboptimalities observed over decades of investigating how animals decide to spend time [3],[4],[6]. Patterns of timing behavior [7] are also consistent with a time-dilating state space. In particular, this compromise between memoryless exponential discounting and reward rate optimization preserves optimal forgo behavior [4], creates a suboptimal bias toward sooner-smaller rewards in mutually exclusive choices [4], and leads to a suboptimal unwillingness to abandon pursuits midway (sunk cost) [6]. Thus, TDRL applied to a precisely time-dilating state space provides 1) the first general mechanistically descriptive explanation of temporal decision making, 2) a normative rationalization for the neural representation of time, and 3) support for the TDRL decision-making framework in the time domain. Temporal decision making can consequently be understood as a locally-biased misestimation of reward rate and opportunity cost, a representation of the infinite future within a finite horizon time, and the representation of the time spent outside an upcoming option as a smaller subjective time inside this horizon time.

III-113. Neural substrates for visual orientation

Sabine Renninger S.RENNI@GMAIL.COM
Michael Orger MICHAEL.ORGER@NEURO.FCHAMPALIMAUD.ORG
Jens Bierfeld JENS.BIERFELD@NEURO.FCHAMPALIMAUD.ORG
Ruth Diez del Corral RUTH.DIEZDELCORRAL@NEURO.FCHAMPALIMAUD.ORG
Rita Esteves RITA.ESTEVES@NEURO.FCHAMPALIMAUD.ORG

Champalimaud Centre for the Unknown

When animals navigate in their environment they use local and global motion cues and luminance information to orient themselves and adaptively interact with objects. Brain-wide circuits integrate such signals across visual space and time to compute appropriate behavioural responses. Understanding the underlying circuit mechanisms requires knowledge of the response properties of large population of neurons and their genetic identity. Using two-photon and light-sheet microscopy we systematically mapped brain-wide activity dynamics with subcellular resolution or high volume rates to a diverse repertoire of behaviourally relevant stimuli, including optic flow pattern and object motion. We dissected neuronal responses by recording from genetically identified cell types, which provided information also on neurotransmitter specificity. To identify conserved functional units, we segmented brain-wide activity data into 3-dimensional regions of interest and applied clustering techniques to relevant dimensions of the data. Our analysis revealed clusters with distinct temporal response dynamics and stimulus selectivity, with luminance being processed largely independently from whole-field motion. We further registered our data to a common brain reference and correlated functional units with anatomically identified brain structures. This allowed us to identify new circuit components, specifically processing global and local motion and to create a brain-wide functional framework for visual processing. One newly identified group is a population of thalamic neurons with stereotypically organized responses to optic flow pattern. To understand their computational relevance for optomotor behaviour we compared their responses with neurons in other brain areas and performed functional perturbations. The majority of thalamic neurons selectively responded to motion when directions in both visual hemispheres matched the neurons preferred direction. Functional-guided laser ablation of these cells tested their role in gating optomotor swimming in freely moving animals.

III-114. A dual integrator model determines the when and what in expectation-guided perceptual decisions

Lluis Hernandez-Navarro[1] LLHERNANDEZ@CLINIC.CAT
Ainhoa Hermoso-Mendizabal[1] AHERMOSM@CLINIC.CAT
Jaime de la Rocha[1] JROCHAV@CLINIC.CAT
Alexandre Hyafil[2] ALEXANDRE.HYAFIL@GMAIL.COM

[1]IDIBAPS
[2]Universitat Pompeu Fabra

In perceptual categorization tasks, reaction times (RTs) not only depend on current stimulus information, but also on urgency and prior expectations. We trained rats in a two-alternative forced choice auditory discrimination task in a free-response paradigm, where correlations in the stimulus sequence induced trial-dependent expectations to repeat/alternate the previous response. Rats showed high accuracy levels that increased with RT for fast responses (RTs$<$80ms), and then plateaued for longer RTs ($>$80ms). Paradoxically, 'express' RTs ($<$80ms, $\approx$ 35% of trials) were independent of stimulus strength, while longer RTs were modulated by stimuli. Since these results are inconsistent with classical Drift Diffusion Models (DDMs), we propose a novel Dual DDM (D3M) in which rats independently integrate time and stimulus to trigger the response. D3M is implemented by (1) a single-threshold DDM that integrates time starting during fixation prior to stimulus onset; (2) a two-threshold DDM that integrates sensory information with a sensory delay following the stimulus onset. Although RTs are triggered by the first among the two integrators reaching threshold, choices are always set by the sign of the stimulus integrator at the time of the decision. Express RTs occur soon after stimulus onset, when only the time integrator, that starts ramping during the fixation period, can reach the threshold. Longer RTs arise from a mixture of time-triggered and stimulus-triggered responses. Rats responded faster when the stimulus was expected, and more slowly when the stimulus was unexpected. Evidence suggests that expectations change the initial offset of the stimulus integrator but also modulate accumulation in the time-integrator. Additionally, rats became slower after an error trial, which seems to stem from a lower initial offset of the temporal integrator. Overall, the novel D3M suggests that the when and the what of a perceptual choice can be more dissociated than predicted by the classic DDM

Author Index