

COSYNE


WORKSHOPS

Breckenridge, CO:

2-3 mar

2020

cosyne.org



 GATSBY

 DeepMind

SIMONS
FOUNDATION

BURROUGHS
WELLCOME
FUND 

COSYNE 2020 WORKSHOPS

**Breckenridge, CO
March 2–3, 2020**

**Organizers:
Catherine Hartley
Blake Richards**

COSYNE 2020 Workshops

March 2–3, 2020, Breckenridge, CO

Monday, March 2, 2020	Organizer(s)	Location
1.1 Interpretable computational neuroscience: What are we modeling and what does it have to do with the brain? – Day 1	M Whiteway J Glaser	<i>Breckenridge Peak 14</i>
1.2 Scrutinizing models of brain function: From in-silico stimulus synthesis to direct brain perturbation – Day 1	K Kar A Afraz J Feather P Bashivan	<i>Breckenridge Peak 17</i>
1.3 Comparative decision making workshop – Day 1	P Reinagel T Hanks R Kiani A Huk	<i>Summit Peaks 6, 7, 8</i>
1.4 Fifteen years of grid cells – Day 1	Y Burak D Derdikman	<i>Colorado Peak 3</i>
1.5 Non-canonical neural responses: Where do they come from and what do they do?	M Insanally J Zylberberg	<i>Colorado Peak 1</i>
1.6 Neural networks that are neuronal networks: Considerations of neuron morphology in circuit computations	S Ing-Esteves R Farhoodi	<i>Breckenridge Peaks 15, 16</i>
1.7 Representational drift	T O'Leary A Fink	<i>Colorado Peak 2</i>
1.8 What's a behavior: Systems neuroscience meets behavioral ecology	A El Hady J Davidson	<i>Summit Peaks 11, 12</i>
1.9 Structure learning: Graphs, manifolds, and geometries	K Stachenfeld I Momennejad	<i>Imperial Ballroom</i>
1.10 The neurobiology of control	C Buckley E Chiappe	<i>Summit Peaks 9, 10</i>

Workshop Co-Chairs

Blake Richards, McGill University
Catherine Hartley, New York University

Email

blake.richards@mila.quebec
cate@nyu.edu

Map of Breckenridge workshop locations is on page 5 of this booklet.

COSYNE 2020 Workshops

March 2–3, 2020, Breckenridge, CO

Tuesday, March 3, 2020	Organizer(s)	Location
2.1 Interpretable computational neuroscience: What are we modeling and what does it have to do with the brain? – Day 1	M Whiteway J Glaser	<i>Breckenridge Peak 14</i>
2.2 Scrutinizing models of brain function: From in-silico stimulus synthesis to direct brain perturbation – Day 1	K Kar A Afraz J Feather P Bashivan	<i>Breckenridge Peak 17</i>
2.3 Comparative decision making workshop – Day 2	P Reinagel T Hanks R Kiani A Huk	<i>Summit Peaks 6, 7, 8</i>
2.4 Fifteen years of grid cells – Day 2	Y Burak D Derdikman	<i>Colorado Peak 3</i>
2.5 Modules in the brain: Compartmentalized and distributed computation	L Driscoll L Duncker	<i>Breckenridge Peaks 15, 16</i>
2.6 Learning to learn. Novel approaches to studying synaptic plasticity.	T Vogels K Kording B Confavreux	<i>Imperial Ballroom</i>
2.7 Closing the gap between neural networks and the brain: A collaborative effort for bridging computational models and experimental data for visual processing	T Marques M Schrimpf	<i>Summit Peaks 11, 12</i>
2.8 Learning in piriform cortex	C Poo C Schoonover	<i>Colorado Peak 1</i>
2.9 Memory, modularity, and attention: Efficient information dispatching in neural computations	R Ke A Goyal B Lansdell G Lajoie	<i>Summit Peaks 9, 10</i>

Workshop Co-Chairs

Email

Blake Richards, McGill University

blake.richards@mila.quebec

Catherine Hartley, New York University

cate@nyu.edu

Map of Breckenridge workshop locations is on page 5 of this booklet.

Schedule

Each workshop group will meet in two sessions from ~8–11am and from ~4.30–7.30pm.

Workshop summaries and schedules are available starting on page 6 of this booklet.

Transportation

For all travel discounts available to Cosyne attendees, please visit www.cosyne.org.

Hilton Denver City Center to Breckenridge: Free shuttle provided for registered attendees (first shuttle leaves @ 4pm, last @ 5pm on Sunday, 01 March 2020).

Breckenridge to Denver Airport: Shuttle can be arranged at Breckenridge, or online at: www.coloradomountainexpress.com

Further information about transportation to/from Breckenridge is available at: www.breckenridge.com.

For further information on transportation or other logistics please contact Leslie Weekes (leslie.weekes@cosyne.org).

Discounted workshop rates

For more details on ski rental and lift tickets discounts, please visit www.cosyne.org.

Ski rental discount for Cosyne attendees: 25 % off ski rentals at www.rentskis.com/breckvip

Group rates for lift tickets. Use your lift ticket coupon for the following walk up window rates:

Off Peak 01 March Coupon Adult 1-Day/2-Day/3-Day/4-Day \$154/\$302/\$426/\$532

Peak >=02 March Coupon Adult 1-Day/2-Day/3-Day/4-Day \$173/\$346/\$501/\$532

Meals included with registration

Breakfast (Day 1 and Day 2) — Colorado ballroom, Peaks 4 & 5

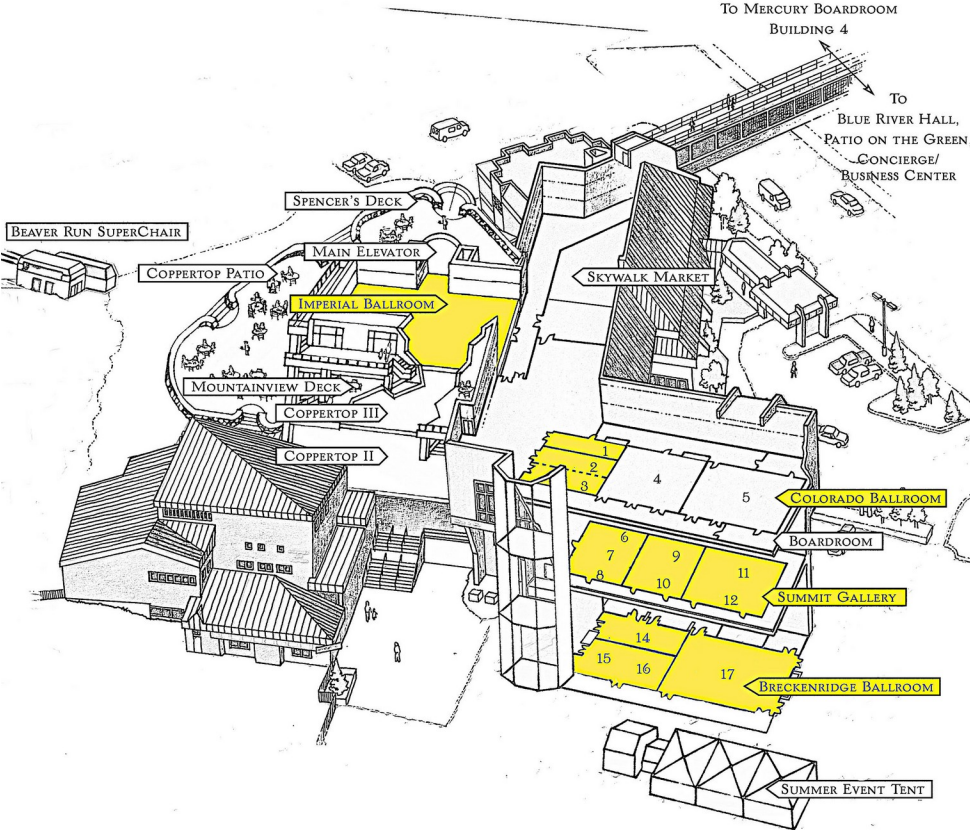
Dinner (Day 2) — Colorado ballroom

Coffee breaks during morning and afternoon sessions

Workshops wi-fi password

Password: *cosyne2020*

Breckenridge workshop locations



1.1 Interpretable computational neuroscience: What are we modeling and what does it have to do with the brain? – Day 1

Monday, March 2, 2020

Organizers: Matt Whiteway, Josh Glaser

Abstract: A principal goal of computational neuroscience is to create mathematical models that lead to biological insight. While the recent development of powerful machine learning techniques has proven effective at discovering predictive structure from large amounts of data, these techniques lack the critical interpretability needed in neuroscience. This is a fundamental problem in a growing and ever-dominant segment of computational neuroscience. We therefore present a two day workshop that aims to both highlight recent advances in interpretable statistical models, and stimulate a discussion on what the “neural state” we are after even is. That is, how does the latent structure extracted mathematically correlate with the true nature of the brain? We aim to bring together computationalists, theorists, and experimentalists to define and meet this central challenge in modern neuroscience.

Day 1: Interpretable statistical models

The first day of this workshop will survey current approaches for creating interpretable statistical models of neural activity. Participants will discuss current challenges to including domain-specific experimental knowledge and theories into statistical models. This includes broader questions such as what constitutes model “interpretability”, and which structural or functional components should be incorporated in our models given current experimental paradigms and recording technologies.

1.1 Interpretable computational neuroscience: What are we modeling and what does it have to do with the brain? – Day 1

Breckenridge ballroom, Peak 14

Morning session

09:00–09:10	<i>Matt Whiteway and Josh Glaser</i> , Introduction
09:10–09:40	<i>Max Burg</i> , Learning divisive normalization in primary visual cortex
09:40–10:10	<i>Daniel Butts</i> , The topology of computation across sensory neuron populations
10:10–10:30	Coffee break
10:30–11:00	<i>Anqi Wu</i> , Extracting structure from high-dimensional neural recordings with Bayesian latent variable models
11:00–11:30	<i>Terry Sejnowski</i> , TBA
11:30–12:00	<i>Kristin Branson</i> , Searching for the rules of animal behavior using deep imitation learning

Afternoon session

16:30–17:00	<i>Grace Lindsay</i> , Merging neural circuit models with deep learning
17:00–17:30	<i>Memming Park</i> , Can dynamical systems be interpreted as cognitive algorithms?
17:30–18:00	Coffee break
18:00–18:30	<i>Lea Duncker</i> , Adding biological constraints to state space models can explain robustness to optogenetic stimulation in motor cortex
18:30–19:00	<i>Scott Linderman</i> , Unifying and generalizing models of neural dynamics during decision-making
19:00–19:30	Panel Discussion/wrap-up

1.2 Scrutinizing models of brain function: From in-silico stimulus synthesis to direct brain perturbation – Day 1

Monday, March 2, 2020

Organizer: Kohitij Kar, Arash Afraz, Jenelle Feather, Pouya Bashivan

We are currently at a stage where neuroscience is coming out of its adolescence, moving beyond under-powered, first-generation investigations of the structure and function of the brain. Recently, much of this progress has been driven by the emergence of high-performing computational models that can predict neural activity and behavior. Nevertheless, these models are far from being perfect, and human intuitions are often insufficient to guide model improvements. Given that the space of all possible models and testable stimuli is very large, we need targeted model-centric techniques to address this issue. This workshop will explore two such methods -- model-based stimulus synthesis and in-vivo causal perturbations. It will bring together researchers who are at the forefront of developing these techniques, with the objective of establishing shared goals and common concerns that will guide the next phase of progress.

Day 1 will focus on in-vivo causal perturbation experiments. The debate over correlation vs causation with respect to brain and behavior is often settled by causal perturbation experiments (e.g. pharmacological, optogenetic, chemogenetic, and electrical interventions). However, given the limited and arbitrary levels of control that current causal tools offer, inferences drawn from such experiments often only confirm prior correlative intuitions, while stronger claims about causal circuit motifs remain challenged. Specific mapping between the brain tissue and the computational models allow us to employ causal perturbation techniques to serve as a tool for screening amongst models. We can now directly test whether in-silico perturbations can predict outcomes of in-vivo perturbations. The goal of Day 1 is to establish a synergy between computational modeling and causal experimental frameworks.

1.2 Scrutinizing models of brain function: From in-silico stimulus synthesis to direct brain perturbation – Day 1

Breckenridge ballroom, Peak 17

Morning session

- 09:00–09:10 *Organizers (Kohitij Kar), Introduction*
Theme 1: Optogenetics in non human primates
- 09:10–09:35 *Arash Afraz, Studying the effect of optogenetic stimulation of inferior temporal cortex on "perception"; a new approach*
- 09:35–10:00 *Adriana Galvan, Advances and challenges using optogenetics in system neurosciences in non-human primates*
- 10:00–10:10 *Organizers and Speakers, Q & A Session*
- 10:10–10:25 *Coffee break*
- Theme 2: Pharmacological and Chemogenetic Approaches*
- 10:25–10:50 *Alexander Huk, Convergent evidence in interpreting perturbation experiments*
- 10:50–11:15 *David Freedman, Assessing the causality of parietal correlates of flexible categorical and perceptual decisions*
- 11:15–11:40 *Kohitij Kar, Chemogenetic and pharmacological perturbations to test circuit motifs critical for core object recognition*
- 11:40–12:00 *Organizers and Speakers, Q & A Session*

Afternoon session

- Theme 3: Causal Perturbation across species*
- 16:30–16:55 *Nancy Kanwisher, Causal evidence for functional and anatomical specificity of face and color processing in the human brain*
- 16:55–17:20 *Lauri Nurminen, Optogenetic inactivation of corticocortical feedback connections in non-human primates*
- 17:20–17:45 *Lindsey Glickfeld, Diverse contributions of mouse visual areas to perceptual decision making*
- 17:45–18:00 *Organizers and Speakers, Q & A session*
- 18:00–18:20 *Coffee break*
- 18:20–19:30 *Audience, Organizers and Speakers, Panel Discussion/wrap-up*

1.3 Comparative decision making workshop – Day 1

Monday, March 2, 2020

Organizers: Pamela Reinagel, Tim Hanks, Roozbeh Kiani, Alex Huk

There is a large body of experimental and theoretical literature on perceptual decision-making, much of which comes from studying the random dot motion visual discrimination task in primates, recording in parietal cortex, and modeling with bounded drift diffusion. Recent years have marked a period of rapid divergence in the study of perceptual decision-making, with empirical studies in many different animal species and sensory modalities, exploring evidence integration over longer timescales, and a plethora of modified or alternative theoretical models. The field is now poised to start synthesizing insights across these dimensions. This workshop will engage comparisons explicitly, with two over-arching themes:

- (1) Can the observed differences in decision-making behavior, neural mechanisms, or computational theories be understood as different manifestations of the same principle at some higher level of abstraction?
- (2) Can differences be understood as individually optimal for the unique constraints of the respective systems?

The workshop is organized around four comparative themes which weave through all the talks: comparing timescales, circuits, species, and theories. Behavioral, neurobiological, and computational approaches will be equally emphasized, and communication that is intelligible across disciplines will be stressed. In the literature each model is often tested on a different data set, and each data set is analyzed or modeled with different code. A goal of this workshop is to promote direct, controlled comparisons. We will promote this agenda in several ways. Many of the selected speakers have made direct comparisons across one or more dimension within the lab. The workshop speakers have agreed to share data and code with one another prior to the workshop, in order to collaborate and present direct comparisons bridging among the talks. A deliverable of the workshop is to create a perceptual decision-making data and code portal, collating links to publicly available data sets, analysis code and model code.

1.3 Comparative decision making workshop – Day 1

Summit gallery, Peaks 6, 7, 8

Morning session	<i>Species Comparisons</i>
09:00–09:15	<i>Pam Reinagel, Tim Hanks</i> , Introduction
09:15–09:45	<i>Tess Oram</i> , Feature Integration Drives Probabilistic Behavior in the <i>Drosophila</i> Escape Response
09:45–10:15	<i>Florian Engert</i> , Neural basis of evidence accumulation and decision making in larval zebra fish
10:15–10:45	<i>Devarajan Sridharan</i> , The role of the midbrain in stimulus selection in birds and humans: A comparative analysis with multidimensional, signal detection theory
10:45–11:00	Coffee break
11:00–11:30	<i>Tim Hanks</i> , Connecting human and rodent research on flexibility of timescales of evidence evaluation
11:30–12:00	<i>Pam Reinagel</i> , Rat versus human visual decision making: Explaining species differences with variable or state-dependent DDMs
Afternoon session	<i>Model Comparisons</i>
16:30–17:00	<i>Zoe Ashwood</i> , Discrete latent states underlie sensory decision making behavior
17:00–17:30	<i>Bo Shen</i> , A dynamic divisive normalization model for choice and working memory
17:30–18:00	Coffee break
18:00–18:15	<i>Gaia Tavoni</i> , The complexity dividend: when sophisticated inference matters: Subjects modulate adaptivity and working-memory load as predicted by the theory
18:15–18:45	<i>Ruben Moreno Bote</i> , Reliability-dependent hierarchical integration of contextual information
18:45–19:30	Panel Discussion, Collaboration and Team Science Opportunities; Highlighting inter-speaker collaborative comparisons

1.4 Fifteen years of grid cells – Day 1

Monday, March 2, 2020

Organizers: Yoram Burak, Dori Derdikman

The discovery of grid cell activity in the entorhinal cortex has provided fundamental new insights on the cortical coding of a high-level cognitive variable (an animal's position in its environment). Fifteen years later, the precise role of grid cells in spatial processing is not yet clear, but much has been learned about grid cell coding – going beyond neural representations of position in simple 2d arenas, about the anatomy of the grid cell system, and on the relationship between grid cells to other spatially modulated cells in the entorhinal cortex and related brain areas. Throughout this period experimental techniques in systems neuroscience have undergone a dramatic revolution, holding a promise to identify mechanisms of network computation in the grid cell system, and to more tightly relate the experimental work to theoretical models of network dynamics, coding, and plasticity. The goal of our proposed workshop is to bring together experimentalists and theoreticians in order to map the recent developments, discuss the key open questions, and propose ways for the field to take advantage of new experimental and theoretical tools to advance our understanding of neural computation in the high-end cortices.

The target audience of the workshop includes a wide range of experimentalists and theorists interested in computational and systems neuroscience, as well as researchers working on the entorhinal cortex and the hippocampus.

1.4 Fifteen years of grid cells – Day 1

Colorado ballroom, Peak 3

Morning session

09:00–09:10

Yoram Burak & Dori Derdikman, Introduction

Introductory talk

09:10–10:00

Edvard Moser, Grid cells and the medial entorhinal cortex: From single cells to networks

Grid cells and the hippocampus

10:00–10:25

Dori Derdikman, The Chicken and egg problem of grid cells and place cells

10:25–10:45

Coffee break

10:45–11:10

Cliff Kentros, Distinct functional roles for distinct entorhinal inputs to the hippocampus

11:10–11:35

Yoram Burak, Functional and emergent consequences of synaptic coupling between grid cells and place cells

11:35–12:00

Neil Burgess, Place and grid cells: encoding states and transitions?

Afternoon session

16:30–16:55

Julia Krupic, Real-time formation of the hippocampal cognitive map

Grid cell formation

16:55–17:20

Vijay Balasubramanian, Dynamical self-organization and reorganization in the grid system

17:20–17:45

Henning Sprekeler, Getting grid cells by failing to learn an excitation-inhibition balance

17:45–18:15

Coffee break

18:15–18:40

Group discussion on Grid formation and function

18:40–19:05

Marianne Fyhn, Perturbing the spatio-temporal organization of the grid cell network

19:05–19:30

Surya Ganguli, A computation origin for grid cells

1.5 Non-canonical neural responses: Where do they come from and what do they do?

Monday, March 2, 2020

Organizers: Michele Insanally, Joel Zylberberg

A shocking number of neurons in the brain do not seem to respond to sensory stimulation, and are otherwise not obviously modulated during behavior. These cells- including cells in primary sensory cortex- typically comprise the majority of datasets yet are often neglected from analysis. What are these cells doing? In recent years, using a range of sophisticated techniques the experimental and computational neuroscience community have begun to look beyond 'classically-responsive' neurons and investigate those with more complex response profiles such as 'mixed selectivity' cells in primate prefrontal cortex, 'category-free' neurons in rodent parietal cortex, non-classically responsive cells in rodent auditory cortex, 'untuned' neurons in visual cortex, and non-canonical hippocampal place cells. These new studies have produced a wealth of new insights and offer an unprecedented opportunity to understand neural circuit function in all of its inherent complexity.

The overall goal of this workshop is to understand the factors that produce non-canonical neural responses and the putative functional role of these responses. To achieve that goal, this workshop will bring together experimentalists monitoring neural activity in behaving animals, theorists studying population dynamics in neural networks, and machine learning researchers studying artificial neural networks. The aim of this workshop is to determine whether there are general principles that govern non-canonical neural activity and explore their putative function in neural circuits. There are three basic topics to be addressed:

- 1) The relation between non-canonical and canonical neural responses
- 2) The plasticity mechanisms that produce non-canonical activity
- 3) The functional role of diverse response types

1.5 Non-canonical neural responses: Where do they come from and what do they do?

Colorado ballroom, Peak 1

Morning session

- 09:00–09:15 *Michele Insanally and Joel Zylberberg*, Introduction
- 09:15–09:45 *Michele Insanally*, Neuronal dynamics of non-classically responsive cortical neurons
- 09:45–10:15 *Stefano Fusi*, Abstract representations: neural responses and their geometry
- 10:15–10:45 *Kanaka Rajan*, Recurrent neural network models of evidence accumulation and decision making
- 10:45–11:00 Break, Non-canonical coffee and discussion
- 11:00–11:30 *Kishore Kuchibhotla*, Behavioral and neural correlates for latent knowledge during learning
- 11:30–12:00 *Konrad Kording*, How non-canonical are neural responses?

Afternoon session

- 16:30–17:00 *Lisa Giocomo*, Mixed selectivity in neural codes for navigation
- 17:00–17:30 *Ari Morcos*, Selectivity considered harmful: evaluating the causal impact of class selectivity in DNNs
- 17:30–18:00 *Joni Wallis*, Prefrontal dynamics underlying decision-making
- 18:00–18:15 Break, Non-canonical coffee and discussion
- 18:15–18:45 *Anitha Pasupathy*, Diverse response properties of neurons in mid-level visual cortex
- 18:45–19:15 *Jane Wang*, What can artificial agents tell us about neural responses?
- 19:15–19:30 Panel Discussion/wrap-up

1.6 Neural networks that are neuronal networks: Considerations of neuron morphology in circuit computations

Monday, March 2, 2020

Organizers: Samantha Ing-Esteves, Roozbeh Farhoodi

Neural network models the brain by point-neurons and connections between them. This simple but powerful model is far from a realistic model of the brain. One of the important shortcomings is morphology of neurons, i.e. geometrical shape of neurons. Imaging technologies have enabled us to measure the morphologies in unprecedented details. Experimentally investigating the impact of neuron morphology on function is obscured, in part, due to the large volume of data produced and lack of standardized data processing techniques. Despite these hurdles, it is emerging that spatial segregation of computations between dendritic compartments and the soma occurs and adds an additional degree of freedom in network computations. The goal of this workshop is to share ideas on fundamental principles of neuron morphologies in terms of theory and experiment. We will have plenty of discussion on how these ideas can potentially build a more realistic model of the brain.

1.6 Neural networks that are neuronal networks: Considerations of neuron morphology in circuit computations

Breckenridge ballroom, Peaks 15, 16

Morning session

- | | |
|-------------|---|
| 09:00–09:10 | <i>Organizers</i> , Introduction |
| 09:10–09:40 | <i>Lida Kanari</i> , A topological understanding of neuronal morphologies; from single cells to detailed neuronal networks |
| 09:40–10:10 | <i>Hermann Cuntz</i> , Quantitative analysis and modeling of dendrite growth |
| 10:10–10:40 | <i>Shreejoy Tripathy</i> , Using single-cell transcriptomics to understand the origins of neuronal morphological and electrophysiological diversity |
| 10:40–10:55 | Coffee break |
| 10:55–11:30 | <i>Giorgio Ascoli</i> , From neuronal classification to network circuitry, and back |
| 11:30–12:00 | <i>Eva Dyer</i> , Barycenters in the brain: An optimal transport approach for modeling connectivity |

Afternoon session

- | | |
|-------------|--|
| 16:30–17:05 | <i>Michael Hausser</i> , Dendritic computations in vivo - linking experiments and theory |
| 17:05–17:40 | <i>Aaron Milstein</i> , Dendritic mechanisms of rapid associative learning |
| 17:40–17:55 | Coffee break |
| 17:55–18:30 | <i>Richard Naud</i> , A role for dendritic interactions in the coordination of learning |
| 18:30–19:05 | <i>Peter Li</i> , Deep Learning of Neuronal Morphologies |
| 19:05–19:30 | Panel Discussion |

1.7 Representational drift

Monday, March 2, 2020

Organizers: Timothy O'Leary, Andrew Fink

Recent work has uncovered marked reorganization over time of neuronal representations of sensory stimuli, cognitive variables, and actions. This raises the question of how processes thought to correlate directly with the responses of individual neurons, such as sensory perception, memory, motor behavior, or task performance, are faithfully maintained.

For example:

- Spatial representations in hippocampus become sufficiently uncorrelated that spatial decoding accuracy falls to chance in just two weeks (Lee, Ziv, Leutgeb).
- Sensory representations in primary auditory as well as olfactory cortices change markedly, merely as a function of time (tones, Rumpel; odors, Fink, Schoonover & Axel).
- Task representation in posterior parietal cortex exhibits wholesale reconfiguration after several weeks, despite consistent behavioral performance (Driscoll & Harvey).

Together these findings present an emerging view that the brain is constantly in flux. Yet not all neuronal representations drift: in primary visual cortex, for example, retinotopy, orientation tuning and ocular dominance are stable (Bonhoeffer, Kuhlman); in somatosensory cortex, although single cells are noisy, population representations remain correlated across time (Helmchen); and in HVC (Long), basal ganglia (Ölveczky), and motor cortex (Chestek) representations remain stable across extended periods.

What preserves stability over time in some systems and what induces representational drift in others? How does an unstable system support stable perception and behavior? Are there coding schemes that permit faithful readout of changing neuronal representations (e.g. an invariant, low dimensional manifold, Solla)? Can downstream systems learn to follow changing representations (O'Leary)? Can a system made of unstable components nonetheless exhibit stability at larger scales (Barak, Loewenstein)? More generally, is there any utility to instability or just a good tolerance for "slop"?

This workshop brings together experimentalists using diverse methods to monitor neuronal representations over time in various brain areas with theorists actively crafting a framework to handle this fundamental problem.

1.7 Representational drift

Colorado ballroom, Peak 2

Morning session

- 09:00–09:15 Introduction (*T. O'Leary/A. Fink*),
- 09:15–09:45 *Albert Lee*, Stability and instability in hippocampal representations
- 09:45–10:15 *Alon Rubin*, Deconstructing the long-term stability and dynamics of hippocampal neuronal code
- 10:15–10:45 *Andrew Fink / Carl Schoonover*, Rapid representational drift in primary olfactory cortex
- 10:45–11:00 Coffee break
- 11:00–11:30 *Maria Neimark Geffen*, Dynamics and stability of sound representation in the auditory cortex
- 11:30–12:00 *Simon Rumpel*, Long-term dynamics of auditory representations during basal conditions and learning

Afternoon session

- 16:30–17:00 *Sara Solla*, Neural manifolds for the stable control of movement
- 17:00–17:30 *Samuel Nason*, Temporal stability of finger-related neural activity in M1 for low-power brain-machine interfaces
- 17:30–18:00 *Yonatan Loewenstein*, From synapses to behavior: on the long-term dynamics of choice bias
- 18:00–18:15 Coffee break
- 18:15–18:45 *Laura Driscoll*, Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex
- 18:45–19:15 *Timothy O'Leary*, Decoding despite drift
- 19:15–19:30 Discussion

1.8 What's a behavior: Systems neuroscience meets behavioral ecology

Monday, March 2, 2020

Organizers: Ahmed El Hady, Jacob Davidson

Recent technological advances in systems neuroscience have given experimenters the ability to train animals on complicated behavioral tasks in order to study cognitive phenomena such as decision making and working memory. Tools to probe the neurophysiological aspects in freely moving animals have been developed in order to unravel the details of the neural mechanisms underlying such behavior.

Although the use of trained laboratory tasks has helped to elucidate the neurobiological mechanisms of decision-making, the behavior analyzed can be remote from the natural behavior of the animal. This makes it difficult to connect to the structure of the natural environment in which the behavior is performed, and evolutionary considerations in terms of the forces shaping the adaptivity of the behavior. These aspects are usually studied in ecology and more specifically behavioral ecology. It is well known in behavioral ecology that a full understanding of behavior must address both proximate (mechanistic) and ultimate (evolutionary) causes. For example, in decision neuroscience, it is becoming clear that 'sub-optimal' decisions can only be understood in the appropriate environmental context.

The workshop will bring together experimentalists and theorists from systems neuroscience and behavioral ecology working to bridge the gap between the study of trained and natural behavior. This is an exciting emerging interdisciplinary area that we think would be of interest for many conference attendees.

The workshop will address the following questions: How one can one use tools from systems neuroscience to understand behavior in a natural context? How can one design laboratory experiments that mimic the natural environments in which the animal behave? How can the concept of optimality, or 'sub-optimality', be used to understand context-dependent behavior? How one can apply the concept of evolutionary homology to behavior, opening up the door for comparative systems neuroscience?

1.8 What's a behavior: Systems neuroscience meets behavioral ecology

Summit gallery, Peaks 11, 12

Morning session

- 09:00–09:15 *Organizers*, Why should behavioral ecology and systems neuroscience converge?
- 09:15–10:00 *Nachum Ulanovsky*, Neural codes for natural behaviors in flying bats
- 10:00–10:30 *Paloma Gonzalez-Bellido*, Neural adaptations and behavioral strategies for interception in two lineages of aerial predatory insects
- 10:30–10:45 Coffee break
- 10:45- 11:15 *Jan Drugowitsch*, The limits of diffusion models for decision-making
- 11:15-11:45 *Song Qi*, Decision making under threat - an ecological framework

Afternoon session

- 16:30–17:00 *Pepe Alcami*, Birdsong in natural conditions. Insights from canaries
- 17:00- 17:30 *Jimena Berni*, A motor rivalry model predicts the environment and genetic interplay during exploratory behaviour in *D. melanogaster*
- 17:30–18:00 *Ahmed El Hady and Jacob Davidson*, Mechanistic theory of foraging
- 18:00–18:15 Coffee break
- 18:15–18:45 *Elad Schneidman*, Socialtaxis and efficient collective foraging
- 18:45–19:30 Panel Discussion

1.9 Structure learning: Graphs, manifolds, and geometries

Monday, March 2, 2020

Organizers: Kim Stachenfeld, Ida Momennejad

As humans, animals, and artificial agents navigate the world they extract the structure of relationships among events. This relational knowledge is then used to guide flexible behavior. Significant theoretical and empirical advances in recent years call for rethinking how we understand learning, storing, and dynamically updating these relationships. While different mathematical approaches converge in their explanations of some empirical evidence, they diverge on other testable predictions. The goal of this workshop is to identify areas of convergence and divergence among these perspectives as well as neural and behavioral data to constrain these models. We hope the discussions will highlight the field's challenges and propose the road ahead.

Now is an opportune moment for updating our collective understanding of structure learning. Theoretical advances in computational modeling and exploratory data analysis have improved our ability to describe state representations in complex tasks without relying on spatial metaphors. In machine learning, emerging interest in improving sample efficiency has led to advances in "representation learning," a field focusing on representations of complex tasks that support efficient learning. In tandem, neuroscientific evidence has broadened our conceptualization of cognitive map-like representations in hippocampus and prefrontal cortex, showing that they need not be spatial nor necessarily Euclidean. Putative spatial regions host more domain-general representations than previously thought, path-dependent representations twist around boundaries, and grid fields are warped by reward.

Our speakers are leading experts in experimental neuroscience (rodent, NHP, human), theory, and machine learning. They represent diverse perspectives on characterizing structured representations in neural data: manifold, graph-based, hierarchical, geometric, and predictive representations. Our workshop will prioritize integrating these divergent viewpoints and formulating future directions for experiments and theory. Thus, this workshop should be of particular interest to experimental researchers testing theories of structured representations as well as theorists and machine learning researchers constraining their models with experimental data from the brain.

1.9 Structure learning: Graphs, manifolds, and geometries

Imperial ballroom

Morning session

09:00–09:15	<i>Organizers 1 and 2</i> , Introduction
09:15–09:40	<i>Kim Stachenfeld</i> , Representation Learning in hippocampal-entorhinal circuit
09:40–10:05	<i>Edward Nieh & Manuel Schottdorf</i> , Geometric representation of abstract learned knowledge by neural manifolds in hippocampus
10:05–10:30	<i>Ila Fiete</i> , Flexible representation and memory of cognitive variables with grid cells
10:30–10:45	Coffee break
10:45–11:10	<i>Danielle Bassett</i> , Graph learning: How humans infer and represent networks;
11:10–11:35	<i>Roozbeh Kiani</i> , Decision Making through Integration of Sensory Evidence at Prolonged Timescales
11:35–12:00	<i>Alla Karpova</i> , Structured representations in sequence learning

Afternoon session

16:30–16:55	<i>Becket Ebitz</i> , Rule adherence warps decision-making
16:55–17:20	<i>Doina Precup</i> , Temporal abstractions for reinforcement learning
17:20–17:45	<i>Ida Momennejad</i> , Learnig Structures: Successor maps, replay, and generalization
17:45–18:00	Coffee break
18:00–18:25	<i>James Whittington</i> , Unifying space and relational memory through generalisation in the hippocampal formation
18:25–18:50	<i>Mehrdad Jazayeri</i> , Hierarchical reasoning by neurons in the frontal cortex
18:50–19:30	Panel Discussion/wrap-up

1.10 The neurobiology of control

Monday, March 2, 2020

Organizers: Christopher Buckley, Eugenia Chiappe

The ultimate goal of neuroscience is to understand the operation of neural circuits in behaving animals, thus, fundamentally the brain should be understood as a control system facilitating the coordination of the body and environment in the pursuit of adaptive behavior. The flow of information during sensorimotor processing spans the entire brain and the activity of the circuits involved are strongly modulated by the presence of active behaviors. Thus while a partial understanding in specific circuits and restrained preparations has been achieved, a full understanding ideally requires recordings of the neural activity from across the brain in a behaving animal. While this has historically been technically prohibitive, recent technological advancements, in particular, the development of virtual reality and imaging techniques, provide an opportunity to record from large populations of neurons during active behavior. Thus there is a timely need to revisit our understanding of the neurobiological basis of sensorimotor control in light of these new data.

We argue that understanding the brain as a control system necessitates a change in both technical language and the central questions asked about neural function. Specifically, it promotes a move away from the language of information coding to stochastic dynamical systems and promotes questions not primarily about categorization and decision-making but about behavioral regulation. This workshop will bring together communities working on neurobiological control across the full gamete of model organisms (worm, fly, fish, mouse, monkey, human) together with representatives of communities where the ideas of control are much more matured, i.e., sensorimotor cognitive scientists and roboticists. The goal is to develop a shared technical perspective and identify common questions, challenges, and theories across systems. This workshop will appeal to experimentalists who want to contextualize their results in terms of modern sensorimotor theory and inspire new algorithms and mechanisms for theorists and engineers.

1.10 The neurobiology of control

Summit gallery, Peaks 9, 10

Morning session

09:00–09:10	<i>Organisers</i> , Introduction
09:10–09:35	<i>Charles Fieseler</i> , Unsupervised learning of control signals and their encodings in <i>C. elegans</i> whole-brain recordings
09:35–10:00	<i>Ann Hermundstad</i> , Operant learning through the adaptive control of behavioral variability
10:00–10:25	<i>Gaby Maimon</i> , A control-theoretic understanding of goal-oriented navigation in the fly central complex
10:25–10:45	Coffee break
10:45–11:10	<i>Megan Carey</i> , Cerebellar control of coordinated locomotion
11:10–11:35	<i>Kathleen Cullen</i> , Predictive Sensing: The Role of Motor Signals in Sensory Processing During Active Self-Motion
11:35–12:00	<i>Ruben Portugues</i> , Motor control in larval zebrafish

Afternoon session

16:30–16:55	<i>Konrad Kording</i> , A learning centric view on control
16:55–17:20	<i>Aman Saleem</i> , Visual processing during active behaviors in mouse
17:20–17:40	Coffee break
17:40–18:05	<i>Florian Engert</i> , Evidence accumulation, integration and decision making in larval zebrafish
18:05–19:30	Panel Discussion / wrap-up

2.1 Interpretable computational neuroscience: What are we modeling and what does it have to do with the brain? – Day 2

Tuesday, March 3, 2020

Organizers: Matt Whiteway, Josh Glaser

Abstract: A principal goal of computational neuroscience is to create mathematical models that lead to biological insight. While the recent development of powerful machine learning techniques has proven effective at discovering predictive structure from large amounts of data, these techniques lack the critical interpretability needed in neuroscience. This is a fundamental problem in a growing and ever-dominant segment of computational neuroscience. We therefore present a two day workshop that aims to both highlight recent advances in interpretable statistical models, and stimulate a discussion on what the “neural state” we are after even is. That is, how does the latent structure extracted mathematically correlate with the true nature of the brain? We aim to bring together computationalists, theorists, and experimentalists to define and meet this central challenge in modern neuroscience.

Day 2: Neural vs. Latent states: What is the truth and how close are we to it?

The second day of this workshop will focus on larger themes related to mathematical understanding of neural systems. Specifically, we often colloquially refer to the “neural state” of a system. However, “latent variable models”, including dynamical systems, dimensionality reduction, and manifolds, only return an unobserved latent state that serves as a descriptor of the true “neural state” of the brain. A formal definition of this neural state remains vague. This day aims to open a discussion on understanding neural states underlying many prominent and emerging computational neuroscience techniques.

2.1 Interpretable computational neuroscience: What are we modeling and what does it have to do with the brain? – Day 2

Breckenridge ballroom, Peak 14

Morning session

09:00–09:15	<i>Anqi Wu & Mikio Aoi</i> , Introduction
09:15–09:45	<i>Evren Gokcen</i> , Dissecting feedforward and feedback interactions between populations of neurons
09:45–10:15	<i>Tatiana Engel</i> , Enhancing accurate interpretation of flexible models for neural population dynamics
10:15–10:45	<i>Angela Langdon</i> , Uncovering the ‘state’: Model-based predictions for the neural circuits of reward learning
10:45–11:00	Coffee break
11:00–11:30	<i>Shreya Saxena</i> , Interpretable models of cortex-wide calcium imaging data
11:30–12:00	<i>David Tank</i> , TBA

Afternoon session

16:30–17:00	<i>Lia Papadopoulos</i> , Role of dynamical state in the response of a large scale brain network to perturbation
17:00–17:30	<i>Maryam Shanechi</i> , Multiscale neural dynamics underlying naturalistic reach-and-grasp behavior
17:30–18:00	<i>Sara Solla</i> , Neural manifolds for the control of movement
18:00–18:15	Coffee break
18:15–18:45	<i>Larry Abbott</i> , Manifesting Latent Variables
18:45–19:30	Panel Discussion/wrap-up

2.2 Scrutinizing models of brain function: from in-silico stimulus synthesis to direct brain perturbation – Day 2

Tuesday, March 3, 2020

Organizer: Kohitij Kar, Arash Afraz, Jenelle Feather, Pouya Bashivan

We are currently at a stage where neuroscience is coming out of its adolescence, moving beyond under-powered, first-generation investigations of the structure and function of the brain. Recently, much of this progress has been driven by the emergence of high-performing computational models that can predict neural activity and behavior. Nevertheless, these models are far from being perfect, and human intuitions are often insufficient to guide model improvements. Given that the space of all possible models and testable stimuli is very large, we need targeted model-centric techniques to address this issue. This workshop will explore two such methods -- model-based stimulus synthesis and in-vivo causal perturbations. It will bring together researchers who are at the forefront of developing these techniques, with the objective of establishing shared goals and common concerns that will guide the next phase of progress.

Day 2 will focus on model-based stimuli synthesis approaches. Recently, the value of predictive models has been augmented by goal-directed stimulus synthesis of images and sounds. Stimuli are generated that achieve a specified objective in a candidate model, and neural or behavioral activity is measured from these synthetic stimuli to test if the model accurately captures sensory representations. The goal of *Day 2* is to discuss how these stimuli synthesis procedures can elucidate model successes and failures. Unlike prior attempts in which relatively simple synthetic stimuli were designed based on the experimenters' intuition, current approaches directly take advantage of the model representations to generate perceptually rich images and sounds. Combined with the advent of easy-to-use tools to implement these algorithms, there is growing interest to begin probing sensory systems using such an approach.

2.2 Scrutinizing models of brain function: From in-silico stimulus synthesis to direct brain perturbation – Day 2

Breckenridge ballroom, Peak 17

Morning session

- 09:00–09:10 *Pouya Bashivan and Jenelle Feather*, Introduction
- 09:10–09:35 *Eero Simoncelli*, Testing perceptual consequences of physiological models
- 09:35–10:00 *Andreas Tolias*, Inception loops in vision
- 10:00–10:25 *Jenelle Feather*, Metamers of neural networks reveal divergence from human perceptual systems
- 10:25–10:45 Coffee break
- 10:45–11:10 *Gabriel Kreiman*, What do neurons want? Using computational models to investigate neuronal tuning preferences in an unbiased manner.
- 11:10–11:35 *Anitha Pasupathy*, Nature-inspired parametric stimuli for probing mid-level ventral visual cortex
- 11:35–12:00 *Pouya Bashivan*, Using computational models to predict and control the responses in populations of neurons

Afternoon session

- 16:30–16:55 *Reza Abbasi-Asl*, DeepTune: modeling and characterizing V4 neurons
- 16:55–17:20 *Nikolaus Kriegeskorte*, Controversial stimuli: pitting neural networks against each other as models of human recognition
- 17:20–17:40 Coffee break
- 17:40–18:05 *Sam Norman-Haignere*, Comparing natural and model-matched sounds reveals functional organization in human auditory cortex
- 18:05–18:30 *Matthias Bethge*, TBD
- 18:30–19:30 Panel Discussion

2.3 Comparative decision making workshop – Day 2

Tuesday, March 3, 2020

Organizers: Pamela Reinagel, Tim Hanks, Roozbeh Kiani, Alex Huk

There is a large body of experimental and theoretical literature on perceptual decision-making, much of which comes from studying the random dot motion visual discrimination task in primates, recording in parietal cortex, and modeling with bounded drift diffusion. Recent years have marked a period of rapid divergence in the study of perceptual decision-making, with empirical studies in many different animal species and sensory modalities, exploring evidence integration over longer timescales, and a plethora of modified or alternative theoretical models. The field is now poised to start synthesizing insights across these dimensions. This workshop will engage comparisons explicitly, with two over-arching themes:

- (1) Can the observed differences in decision-making behavior, neural mechanisms, or computational theories be understood as different manifestations of the same principle at some higher level of abstraction?
- (2) Can differences be understood as individually optimal for the unique constraints of the respective systems?

The workshop is organized around four comparative themes which weave through all the talks: comparing timescales, circuits, species, and theories. Behavioral, neurobiological, and computational approaches will be equally emphasized, and communication that is intelligible across disciplines will be stressed. In the literature each model is often tested on a different data set, and each data set is analyzed or modeled with different code. A goal of this workshop is to promote direct, controlled comparisons. We will promote this agenda in several ways. Many of the selected speakers have made direct comparisons across one or more dimension within the lab. The workshop speakers have agreed to share data and code with one another prior to the workshop, in order to collaborate and present direct comparisons bridging among the talks. A deliverable of the workshop is to create a perceptual decision-making data and code portal, collating links to publicly available data sets, analysis code and model code.

2.3 Comparative decision making workshop – Day 2

Summit gallery, Peaks 6, 7, 8

Morning session

Circuit Comparisons: Distributed circuits for decision making

- 09:00–09:15 *Alex Huk, Roozbeh Kiani*, Introduction
- 09:15–09:45 *Xiao-Jing Wang*, Distributed decision-making in a large-scale brain circuit
- 09:45–10:15 *Nuo Li*, Modularity and redundancy in multi-regional circuits underlie robustness of short-term memory: Dissecting multi-regional circuits by population recording and modeling
- 10:15–10:45 *Roozbeh Kiani*, The geometry of the representation of decision variable and stimulus difficulty across the frontoparietal cortex
- 10:45–11:00 Coffee break
- 11:00–11:30 *Christian Machens*, How many dimensions for a single decision?
- 11:30–12:00 *Long Ding*, The caudate nucleus causally contributes to monkeys' reward bias strategy for a perceptual decision task.

Afternoon session

Timescale Comparisons: Timescales of decision making

- 16:30–17:00 *Alla Karpova*, Changing environments, medial frontal circuits, randomness, sequential decisions
- 17:00–17:30 *Christine Constantinople*, OFC dynamics reflecting computations within and across trials during value-based decision-making
- 17:30–18:00 *Aaron Levi*, Explicit control of temporal weighting to dissect choice-related activity in the primate dorsal stream
- 18:00–18:15 Coffee break
- 18:15–18:45 *Gabriel Stine*, Differentiating between integration and non-integration strategies in perceptual decision making
- 18:45–19:30 Panel Discussion, Data and Code Repositories, Benchmarks; Decision-Making resource webpage

2.4 Fifteen years of grid cells – Day 2

Tuesday, March 3, 2020

Organizers: Yoram Burak, Dori Derdikman

The discovery of grid cell activity in the entorhinal cortex has provided fundamental new insights on the cortical coding of a high-level cognitive variable (an animal's position in its environment). Fifteen years later, the precise role of grid cells in spatial processing is not yet clear, but much has been learned about grid cell coding – going beyond neural representations of position in simple 2d arenas, about the anatomy of the grid cell system, and on the relationship between grid cells to other spatially modulated cells in the entorhinal cortex and related brain areas. Throughout this period experimental techniques in systems neuroscience have undergone a dramatic revolution, holding a promise to identify mechanisms of network computation in the grid cell system, and to more tightly relate the experimental work to theoretical models of network dynamics, coding, and plasticity. The goal of our proposed workshop is to bring together experimentalists and theoreticians in order to map the recent developments, discuss the key open questions, and propose ways for the field to take advantage of new experimental and theoretical tools to advance our understanding of neural computation in the high-end cortices.

The target audience of the workshop includes a wide range of experimentalists and theorists interested in computational and systems neuroscience, as well as researchers working on the entorhinal cortex and the hippocampus.

2.4 Fifteen years of grid cells – Day 2

Colorado ballroom, Peak 3

Morning session

Circuitry and dynamics

- 09:00 -09:25 *Ila Fiete*, TBD
- 09:25–09:50 *Stefan Leutgeb*, Local computations in the superficial layers of medial entorhinal cortex
- 09:50–10:15 *Menno Witter*, Identification and connectional neurobiology of grid cells: relevance of local versus extrinsic connections.
- 10:15–10:45 Coffee break
- 10:45–11:10 *Hannah Monyer*, TBD

Grid cell function: coding, memory, and navigation

- 11:10–11:35 *Lisa Giocomo*, Multiple maps for navigation
- 11:35–12:00 *Caswell Barry*, Models of grid network topography in mEC predict macroscopic activity signals

Afternoon session

- 16:30–16:55 *Matt Nolan*, Downstream functions of grid cell outputs
- 16:55–17:20 *Kim Stachenfeld*, Structured representations in MEC for efficient reinforcement learning
- 17:20–17:45 *Beth Buffalo*, Grid-like responses in the primate entorhinal cortex
- 17:45–18:15 Coffee break

Deconstructing the grid

- 18:15–18:40 *Charlotte Boccara*, Goal-distortion of grid cells: a code beyond spatial
- 18:40–19:05 *Alessandro Treves*, Grid representations beyond the age of innocence
- 19:05–19:30 *Nachum Ulanovsky*, Locally-ordered representation of 3D space in the entorhinal cortex

2.5 Modules in the brain: Compartmentalized and distributed computation

Tuesday, March 3, 2020

Organizers: Laura Driscoll, Lea Duncker

Cerebral cortex can be divided into anatomically distinct regions with particular functional roles -- a feature described as modularity. A wealth of evidence from impairment after brain trauma and fMRI studies in humans supports this view, and suggests a compartmentalized strategy for cortical information processing. Recent advances in experimental techniques such as widefield imaging, neuropixels probes, fast scanning two-photon microscopes and optogenetics, have opened up new possibilities for measuring and perturbing the flow of activity across brain areas in behaving animals. With the greater temporal and spatial resolution of these new tools, we have begun to uncover a more complicated picture of functional specialization in the brain. Recent findings suggest that many aspects of cortical processing are highly distributed. For instance, motor information is available in sensory cortices, and inactivation or ablation of functionally specialized regions may be compensated for by alternative pathways. Making progress in understanding these new findings in the context of modularity crucially relies on also developing a strong theoretical framework. What are the computational benefits of both compartmentalized and distributed information processing?

The goal of this workshop is to bring together experimentalists and theorists studying topics relating to inter-area information processing and modularity. The aim is to integrate recent experimental findings with theoretical work on hierarchical information processing, modular network architectures, and disentangling representations, which may provide insight into theoretical advantages provided through compartmentalized and distributed processing strategies. Discussions will focus on establishing a conceptual framework for interactions across multiple brain regions and address why information processing is both distributed and compartmentalized in the brain.

2.5 Modules in the brain: Compartmentalized and distributed computation

Breckenridge ballroom, Peaks 15, 16

Morning session

- 09:00–09:10 *Laura Driscoll & Lea Duncker* Introduction
- 09:10–09:45 *Nancy Kanwisher*, Functional Specificity in the Human Brain: What and Why?
- 09:45–10:20 *Nick Steinmetz*, Distributed coding of choice, action, and engagement across the mouse brain
- 10:20–10:35 Coffee break
- 10:35–11:10 *Lucas Pinto*, Task-specific whole-cortical states during decision making
- 11:10–11:45 *Laura Driscoll*, Neural dynamics shape task organization in multitask networks
- 11:45–12:00 Panel Discussion

Afternoon session

- 16:30–17:05 *Evan Schaffer*, The many dimensions of neural dimensionality
- 17:10–17:45 *Jonathan Michaels*, A modular neural network model of grasp movement generation
- 17:45–17:55 Coffee break
- 17:55–18:30 *Kanaka Rajan*, Multi-region “networks of networks” models of adaptive and maladaptive state transitions
- 18:30–19:05 *Maneesh Sahani*, Modularity in unsupervised learning
- 19:05–19:30 Panel Discussion

2.6 Learning to learn.

Novel approaches to studying synaptic plasticity.

Tuesday, March 3, 2020

Organizers: Tim Vogels, Konrad Kording, Basile Confavreux

Over the past decade it has become clear that the rules that establish, maintain and adapt neuronal networks are as important as the connectivity structure that will eventually support their function. Unlike for the field of connectomics, for which a large number of tools and techniques are being developed to investigate and map circuit structure, there are currently very few promising experimental approaches to probe the function and interaction of the multitude of concurrently active plasticity rules.

In the interim, computational and theoretical approaches offer increasingly sophisticated ideas of how to design and predict plasticity rules *in silico*, and compare their effects with the available or conceivable datasets. The aim of this workshop is to showcase the latest approaches to observe and suggest circuit-making rules. Our workshop is intended for the theorists willing to meet experimentalists for fruitful collaboration, non-expert theoreticians curious about this fashionable yet often obscure subject, and also for experimentalists working on plasticity and looking for new data analysis tools or hypotheses to test.

The day will be divided in three parts. First, major contributors from Machine Learning will come to present techniques recently pushed forward in their field that could be of interest to Neuroscience. Then, a selection of promising theoretical contributions using such novel approaches to study learning rules will be presented. Finally, experimentalists will be given the stage, hopefully leading to discussions and fruitful interactions between theorists and experimentalists interested in plasticity.

2.6 Learning to learn.

Novel approaches to studying synaptic plasticity.

Imperial ballroom

Morning session

09:00–09:15	Organizers, Introduction
09:15–09:50	<i>Niru Maheswaranathan</i> , Meta-learning biologically plausible update rules
09:50–10:20	<i>Ashok Litwin-Kumar</i> , Models of heterogeneous dopamine signaling in an insect learning and memory center
10:20–10:35	Coffee break
10:35–11:05	<i>Mitya Chklovskii</i> , Pyramidal neurons as canonical correlation analyzers
11:05–11:30	<i>Raoul-Martin Memmesheimer</i> , <i>Christian Klos</i> , Dynamical learning of dynamics
11:30–12:00	<i>Wolfgang Maass</i> , A normative model for synaptic plasticity based on gradient descent in recurrent networks of spiking neurons

Afternoon session

16:30–17:00	<i>Henning Sprekeler</i> , Short term plasticity and ML-designed synapse types
17:00–17:30	<i>Nicolas Brunel</i> , Inferring plasticity rules from data
17:30–17:55	<i>Julijana Gjorgjieva</i> , Local and global organization of synaptic inputs on cortical dendrites
17:55–18:15	Coffee break
18:15–18:40	<i>Jayeeta Basu</i> , How can theory be useful to experimentalists
18:40–19:00	<i>Konrad Kording</i> , Towards learning to learn
19:00–19:30	Discussion and closing remarks

2.7 Closing the gap between neural networks and the brain: A collaborative effort for bridging computational models and experimental data for visual processing

Tuesday, March 3, 2020

Organizers: Tiago Marques, Martin Schrimpf

During the past decade, groundbreaking advances in recording and manipulation techniques have led to a surging availability of neural, anatomical, and behavioral data in primate vision. Recently, artificial neural networks (ANNs) have started to explain some of these experimental results. However, since modeling and experimental approaches are currently conducted largely independently and focused on isolated brain areas or phenomena, our community is missing out on a potentially highly beneficial collaboration. Can we take the next step and integrate all the available experimental data to guide the development of single models that can accurately describe the whole primate ventral stream?

In this workshop, we will facilitate such collaborative efforts by bringing together recent experimental and modeling advances and discuss strategies for bridging the gap between them. We will use as a case study Brain-Score which integrates a wide range of experimental results in the form of quantitative benchmarks to evaluate how well computational models match the brain. We will identify: aspects of visual processing that models can or cannot explain; existing experimental results which have not yet been considered; future experiments that would maximally distinguish between model predictions; how to concretely map models to brain anatomy; and directions to guide the development of the next generation of models of primate vision. This approach will ultimately point the way to a deeper understanding of the cortical mechanisms underlying visual cognition.

Such endeavor can only be achieved by a large collaborative effort between experimental and modeling research groups. This workshop targets both computational neuroscientists building models of visual processing along the ventral stream and experimental neuroscientists who provide quantitative findings to evaluate models in the form of benchmarks. Additionally, we hope to inspire similar collaborations across the field of systems neuroscience, and extending to multiple model organisms.

2.7 Closing the gap between neural networks and the brain: A collaborative effort for bridging computational models and experimental data for visual processing

Summit gallery, Peaks 11, 12

Morning session

- 09:00–09:10 *Tiago Marques / Martin Schrimpf*, Introduction to morning session
- 09:10–09:30 *Tiago Marques*, Hierarchical neural networks as models of primate low-level visual processing
- 09:30–09:55 *Farran Briggs*, Corticogeniculate feedback regulates the timing and precision of feedforward visual signals
- 09:55–10:20 *Alessandra Angelucci*, Computational function of feedback connections in early visual processing
- 10:20–10:35 Coffee break
- 10:35–11:00 *Odelia Schwartz*, Incorporating flexible normalization models of visual cortex into deep neural networks
- 11:00–11:25 *Dan Yamins*, Three Big Problems in the Modern DNN Theory of the Ventral Stream, and (Tentative) Approaches to Solutions To Them
- 11:25–12:00 Panel Discussion, Improving hierarchical models of the ventral stream inspired by primate low-level vision

Afternoon session

- 16:30–16:35 *Martin Schrimpf / Tiago Marques*, Introduction to afternoon session
- 16:35–16:55 *Martin Schrimpf*, Brain-Score: a collaborative platform for evaluating computational models of vision
- 16:55–17:20 *Andreas Tolias*, Engineering a less artificial intelligence
- 17:20–17:45 *Anthony Movshon*, Cascaded cortical computation
- 17:45–18:00 Coffee break
- 18:00–18:25 *Pinglei Bao*, A unified model for understanding the functional organization of inferotemporal cortex
- 18:25–18:50 *Alice O'Toole*, Understanding high level representations in deep convolutional neural networks: How do they constrain the interpretation of neural data?
- 18:50–19:00 *James DiCarlo*, Wrap-up and outlook
- 19:00–19:30 Panel Discussion, Which data do current models not predict and which model predictions should be tested?

2.8 Learning in piriform cortex

Tuesday, March 3, 2020

Organizers: Cindy Poo, Carl Schoonover

The goal of the brain is to process streams of sensory information and produce flexible behavior through learning. Primary olfactory cortex, or piriform, has traditionally been considered a sensory processing region. However, striking resemblances in circuitry between piriform and hippocampus have generated a longstanding conjecture that they implement similar functions. This view has received increasing theoretical and empirical support in recent years and has motivated a diverse set of proposed roles for piriform:

- Implements content-addressable memory
- Implements Hopfield-style pattern completion
- Forms Hebbian assemblies
- Does / does not encode learned valence
- Forms engrams
- Supports spatial navigation
- Implements predictive coding
- Learns local structure of chemical space
- Implements fast statistical learning

This varied, and in some cases contradictory, set of proposals reflects a renaissance that the field has enjoyed in the last five years. Many of the ideas above, and the theoretical and experimental support for them, remain unpublished. This workshop will bring key players of this community in the same room--many of whom have never met. It will provide an informal setting for discussing and reconciling divergent views about piriform. And it will explore the extent to which piriform, a mere two synapses downstream of the sensory receptor layer, may serve as a simple model system for gaining traction on problems that have resisted elucidation in the more complex, multimodal hippocampus circuit.

2.8 Learning in piriform cortex

Colorado ballroom, Peak 1

Morning session

- 09:00–09:05 *Carl/Cindy*, Introduction
- 09:05–09:15 *Christiane Linster*, Anatomy introduction
- 09:15–09:45 *Rainer Friedrich*, Specific organization of inhibition in olfactory memory circuits
- 09:45–10:15 *Evan Schaffer*, Aligning neural representations in space and time
- 10:15–10:45 *Alex Fleischmann*, Genetic tagging and manipulation of odor-responsive neurons in mouse piriform cortex
- 10:45–11:00 Coffee break
- 11:00–11:30 *Anne-Marie Oswald*, Circuit and synaptic plasticity in olfactory cortex.
- 11:30–12:00 *Qi Yuan*, Patterns matter: Differential roles of locus coeruleus phasic and tonic activations in odour discrimination and valence learning

Afternoon session

- 16:30–17:00 *Kevin Franks*, Activity-dependent stabilization and destabilization in piriform cortex
- 17:00–17:30 *Robert Datta / Stan Pashkovski*, Cortex constructs a systematic representation of odor space
- 17:30–18:00 *Christiane Linster*, Changes in cortical dynamics in response to odor learning
- 18:00–18:15 Coffee break
- 18:15–18:45 *Cindy Poo*, Spatial maps in posterior piriform cortex during olfactory navigation
- 18:45–19:15 *Leslie Kay*, Piriform cortex as a dynamical controller in olfactory processing
- 19:15–19:30

2.9 Memory, modularity, and attention: Efficient information dispatching in neural computations

Tuesday, March 3, 2020

Organizers: Rosemary Ke, Anirudh Goyal, Ben Lansdell, Guillaume Lajoie

Architectural and cognitive evidence tells us the brain is organized into different subnetworks, specializing in different aspects of behavioral function, e.g., vision, motor control, and memory. Information is shared selectively between these networks in order to perform computations. For instance, only relevant parts of a visual scene for a given task are transmitted to downstream areas, and only relevant memories are recalled in order to aid decision making. Such behaviors are mediated in part by top-down attentional mechanisms, routing information between networks and selecting relevant memories. But much remains to be characterized about how each of these factors — memory, modularity and attention — interact in order to accomplish a given task. Interestingly, recent approaches in artificial intelligence and machine learning have started using exactly these mechanisms to improve their performance, suggesting these ingredients offer universal learning advantages. Both neuroscience and artificial intelligence may thus benefit from models that explore novel attentional and memory mechanisms and modular architectures.

Targeted group of participants are both neuroscientist, cognitive scientist and machine learning scientists. This workshop will bring together experimental, cognitive, and computational neuroscientists, as well as artificial intelligence researchers, to explore questions surrounding the mechanisms of attention, memory, and modularity in neural networks.

2.9 Memory, modularity, and attention: Efficient information dispatching in neural computations

Summit gallery, Peaks 9, 10

Morning session

09:00–09:10	<i>Organizers</i> , Opening remarks
09:10–09:55	<i>Yoshua Bengio</i> , Consciousness Priors for Deep Learning
09:55–10:30	<i>Anna Schapiro</i> , Rapid learning with distributed representations in the hippocampus
10:30–10:45	Coffee break
10:45–11:20	<i>David Sussillo</i> , Neural dynamics as a limited resource in multitask networks
11:20–11:55	<i>Ida Momennejad</i> , Memory for planning: predictive maps and prioritized replay

Afternoon session

16:30–17:05	<i>Michael Mozer</i> , Modularity and information transmission in neural architectures
17:05–17:40	<i>Sandro Romani</i> , Learning rules underlying hippocampal representations
17:40–18:15	<i>Bruno Averbeck</i> , Neural basis of Bayesian inference during reversal learning
18:15–18:30	Coffee break
18:30–19:00	<i>Greg Wayne</i> , Long-term Credit Assignment by Mental Time Travel
19:00–19:30	Panel Discussion/wrap-up