

# **Neural Coding, Computation and Dynamics 2019**

**22-25 September 2019, Capbreton, France**

# Program

Sunday, September 22, 2019

5:00 pm - 7:00 pm    Registration and Welcome  
7:00 pm - 7:05 pm    Opening remarks - Organizers  
7:05 pm - 7:50 pm    Neural sequences and dynamics - Gilles Laurent (MPI Frankfurt)  
8:00 pm - 10:00 pm    Dinner

Monday, September 23, 2019

9:30 am - 10:15 am    Neural control of vocal turn-taking in zebra finches  
                              - Daniela Vallentin (MPI Munich)  
10:15 am - 10:45 am    Functional segregation of the ferret auditory cortex probed  
                              with natural and model-matched sounds - Agnès Landemard, Célian Bimbard,  
                              Sam Norman-Haignere, Shihab Shamma, Yves Boubenec (ENS Paris)  
10:45 am - 11:15 am    Coffee break  
11:15 am - 12:00 pm    Controlling movement and coding sensation in orofacial sensorimotor pathways  
                              - David Kleinfeld (UCSD)  
12:00 pm - 12:30 pm    Dimensionality of parallel fibre population activity in awake mice  
                              - Alex Cayco Gajic, Frederic Lanore, Angus Silver (UCL)  
12:30 pm - 2:30 pm    Lunch  
2:30 pm - 3:15 pm    Learning content-independent transformations: A visual motion detector  
                              - Mitya Chklovskii (Flatiron Institute NYU)  
3:15 pm - 3:45 pm    The emergence of multiple retinal cell types through efficient coding  
                              of natural movies - Stéphane Deny (Stanford)  
3:45 pm - 4:15 pm    Coffee break  
4:15 pm - 5:00 pm    Representations and transformations in the mammalian olfactory system  
                              - Sandeep Robert Datta (Harvard)  
5:00 pm - 5:30 pm    Developmental and evolutionary principles of olfactory circuit design  
                              - Naoki Hiratani, Peter Latham (UCL)  
5:30 pm - 8:00 pm    Poster Session I: P1-P13  
8:00 pm - 10:00 pm    Dinner

## Tuesday, September 24, 2019

- 9:30 am - 10:15 am Spike-based coding and computation  
- Sophie Deneve (ENS Paris)
- 10:15 am - 10:45 am Representing a background signal in spike coding networks allows for response normalization  
- Sander Keemink, Nuno Calaim, Christian Machens (Champalimaud Lisbon)
- 10:45 am - 11:15 am Coffee break
- 11:15 am - 12:00 pm Linking changes in neurons to perception  
- Marlene Cohen (Pittsburgh)
- 12:00 pm - 12:30 pm The liminal state: neural signatures of perceptual awareness  
- Meenakshi Asokan, Kenneth Hancock, Ross Williamson, Daniel Polley (Harvard)
- 12:30 pm - 2:30 pm Lunch
- 2:30 pm - 3:15 pm The agony of choice  
- Alexandre Pouget (Geneva)
- 3:15 pm - 3:45 pm Constrained plasticity can compensate for ongoing drift in the parietal cortex  
- Michael E. Rule, Adrianna R. Loback, Dhruva V. Raman, Christopher D. Harvey, Timothy S. O'Leary (Cambridge)
- 3:45 pm - 4:15 pm Coffee break
- 4:15 pm - 4:45 pm Basic principles of neural ensemble organization for multi-neuronal coding  
- Ioannis Smyrnakis, Anna Palagina, Stelios Smirnakis, Maria Papadopouli (University of Crete)
- 4:45 pm - 5:15 pm Dynamical mechanisms of flexible timing by temporal scaling  
- Manuel Beiran, Mehrdad Jazayeri, Srdjan Ostojic (ENS Paris)
- 5:30 pm - 8:00 pm Poster session II: P14-P27
- 8:00 pm - 10:00 pm Dinner

## Wednesday, September 25, 2019

- 9:30 am - 10:15 am Title TBA - Andrea Hasenstaub (UCSF)
- 10:15 am - 10:45 am Maintenance of visual working memory by a visuo-premotor cortical loop  
- Ivan Voitov, Thomas Mrsic-Flogel (UCL)
- 10:45 am - 11:15 am Coffee break
- 11:15 am - 11:45 am Traveling waves shape neural computations in vision  
- Lyle Muller, Davis Zachary, John Reynolds, Terry Sejnowski (Salk Institute)
- 11:45 am - 12:30 pm (Re)thinking understanding the brain  
- Peter Latham (UCL)
- 12:30 pm - 2:30 pm Lunch

# Table of contents

The liminal state: neural signatures of perceptual awareness, Meenakshi Asokan	3
Dynamical mechanisms of flexible timing by temporal scaling, Manuel Beiran . .	4
Representing a background signal in spike coding networks allows for response normalization, Sander Keemink . . . . .	5
Maintenance of visual working memory by a visuo-premotor cortical loop, Ivan Voitov . . . . .	6
Traveling waves shape neural computations in vision, Lyle Muller . . . . .	7
Developmental and evolutionary principles on olfactory circuit designs, Naoki Hiratani . . . . .	8
Functional segregation of the ferret auditory cortex probed with natural and model-matched sounds, Agnès Landemard . . . . .	9
Constrained plasticity can compensate for ongoing drift in the parietal cortex, Michael E. Rule . . . . .	10
The emergence of multiple retinal cell types through efficient coding of natural movies, Stéphane Deny . . . . .	11
Dimensionality of parallel fibre population activity in awake mice, Alex Cayco Gajic . . . . .	12
Basic principles of neural ensemble organization for multi-neuronal coding, Ioannis Smyrnakis . . . . .	13
P1. Does structure in neural activity match anatomical structure?, Thomas Delaney . . . . .	15
P2. Associative memory in winner-take-all networks: from binary units to spikes, Roman Pogodin . . . . .	16

P3. Presynaptic inhibition rapidly stabilised recurrent excitation in the face of plasticity, Laura Bella Naumann . . . . .	17
P4. Statistical symmetries in connectivity shape slow activity manifolds in recurrent neural networks, Francesca Mastrogiuseppe . . . . .	19
P5. Transient coding through non-normal dynamics during OFF responses in auditory cortex, Giulio Bondanelli . . . . .	21
P6. From an ethogram to a model of neuronal computation: the optomotor response in zebrafish larva., Adrien Jouary . . . . .	22
P7. Disentangling the roles of dimensionality and cell classes in neural computations, Alexis Dubreuil . . . . .	24
P8. Processing information stored in working memory through modulations of effective connectivity, Adrian Valente . . . . .	25
P9. Predictive coding in the salamander retina, Danica Despotović . . . . .	26
P10. Disentangling encoding of stimuli features, perceptual category and behavioural choice in ferret A1 population activity, Rupesh Chillale . . . . .	27
P11. A novel framework for detecting synchronously active cell assemblies featuring extremely low non-detect error rates, Gabriel Makdah . . . . .	28
P12. Operational range of sensory response variability during neurodevelopment, Beatriz E. P. Mizusaki . . . . .	30
P13. Adaptive optimisation of visual prosthetic stimulation, Tristan Fauvel . . . . .	31
P14. Learning low-dimensional inputs for brain-machine interface control, Jorge A. Menendez . . . . .	33
P15. Counterintuitive effects of excitatory and inhibitory perturbations in spike-coding networks, William Podlaski . . . . .	35
P16. Predicting point process GLM excessively high firing rates, Dmitry Todorov . . . . .	36
P17. A Dynamic Model for Decoding Direction and Orientation in Macaque's Primary Visual Cortex, Wahiba Taouali . . . . .	37
P18. Coordination of inhibitory Golgi cell population activity in the cerebellar cortex, Harsha Gurnani . . . . .	39

P19. A Unified Theory of Early Visual Representations from Retina to Cortex through Anatomically Constrained Deep CNNs, Stéphane Deny . . . . .	40
P20. Rate-space attractors and low dimensional dynamics interact with spike-synchrony statistics in spiking neural networks, Daniel N. Scott . . . . .	41
P21. Dynamics of random networks with correlated low-rank perturbations., Friedrich Schuessler . . . . .	42
P22. Coherence states and signal transfer of communicating gamma oscillatory neural networks, Gregory Dumont . . . . .	43
P22. Computational analysis of cholinergic mechanisms for disinhibition of CA1 pyramidal cells and induction of hippocampal plasticity, Inês Guerreiro . . . . .	45
P23. Controlling Burst Activity Allows for a Multiplexed Neural Code in Cortical Microcircuits, Filip Vercauteren . . . . .	46
P24. On the spectra and connectivity properties of transiently amplifying random networks, Georgia Christodoulou . . . . .	47
P25. Functional network models from algorithmically designed synaptic plasticity rules, Basile Confavreux . . . . .	49
P26. Complementary Excitatory and Inhibitory Plasticity Rules produce Sparse Coding through EI Balance in a Spiking Neural Network Model, Nasir Ahmad . . . . .	51
P27. Learning functional properties of visual cortex from ultrafast fluorescence microscopy recordings, Sebastien Wolf . . . . .	52

**Author Index**

**53**

# The liminal state: neural signatures of perceptual awareness

Meenakshi Asokan \* <sup>1</sup>, Kenneth Hancock <sup>2</sup>, Ross Williamson <sup>2</sup>, Daniel Polley <sup>1,2</sup>

<sup>1</sup> Harvard University [Cambridge] – United States

<sup>2</sup> Harvard Medical School [Boston] – United States

Perception involves using the sensory information available to us to form a meaningful representation of the world around us. In most real world scenarios, that involves analyzing and segmenting the sensory scene presented to detect and track the signals of interest amidst potentially distracting stimuli. While scene analysis is a fundamental problem to most of the sensory systems, auditory scene analysis is particularly hard since sound sources add linearly to create the signal hitting our ear drums and even when the ear’s resolution is not a problem, the brain’s information processing capacity can be a bottleneck, especially when it is further modulated by internal state and context. In such scenarios with a complex acoustic background and high perceptual load, a target can go undetected even when it is well above pure sensory threshold. While modulation of neural responses by behavior and context have been shown in simple detection tasks in different regions of the central auditory system, the different regions have never been studied simultaneously in a task complex enough to potentially require different contributions from each region. It is therefore not known whether these modulation effects are reinvented at every processing stage or whether they are hierarchically inherited and if so, where these effects originate. Here, we trained head-fixed mice in a complex acoustic task where they have to detect a target which is a stream of repeating tones embedded in a complex stochastic multi-tone background. We made extracellular recordings of single units from three 64-channel recording probes positioned in the inferior colliculus, medial geniculate body and auditory cortex. This allows us to track dozens of single neurons from interconnected brain regions as the animal switches between task-engaged and passively listening contexts. We tease apart the varying contributions from the different processing centers along the auditory neuro-axis to solving a complex task in different contexts. Our goal is to figure out where in the auditory neuro-axis neural activity first co-varies with perceptual awareness so that we begin to understand how perceptual representations arise from purely sensory representations.

---

\*Speaker

# Dynamical mechanisms of flexible timing by temporal scaling

Manuel Beiran \* <sup>1</sup>, Mehrdad Jazayeri <sup>2</sup>, Srdjan Ostojic <sup>1</sup>

<sup>1</sup> Laboratoire de Neurosciences Cognitives Computationnelles – Ecole Normale Supérieure de Paris - ENS Paris – France

<sup>2</sup> Massachusetts Institute of Technology (MIT) – United States

Humans and other animals can flexibly produce a given action at vastly different speeds. To investigate the neural substrate of such behavior, Wang et al. [1] recorded neural activity in medial frontal cortex of monkeys performing a flexible timing task. They found that flexible timing relied on temporal scaling of neural population activity along an invariant trajectory in neural state space. In this work, we investigated the core dynamical mechanisms that allow a network of recurrently connected units to implement flexible timing through such temporal scaling along an invariant manifold.

To address this question, we relied on the framework of low-rank recurrent networks [2]. We started by training recurrent networks with constrained rank to perform flexible timing, and next reverse-engineered them to isolate the mechanisms used to solve the task. Applying mean field theory, we then analytically determined the core dynamical phenomenon and identified classes of networks that implement it.

We found that networks constrained to be unit rank generate bistable dynamics that lead to temporally scaled trajectories along a single dimension. Such solutions however show limited robustness and can only scale the responses within a restricted range of timescales. Allowing the network connectivities to be of rank-two in contrast led to trained recurrent networks that could precisely perform a flexible timing task spanning a much broader range of timescales. Similarly to experimental findings, such networks relied on an invariant, low-dimensional manifold along which the speed is controlled by the external cue. Using mean-field theory, we analytically showed that, in a large class of rank-two networks, this slow activity manifold originates from a ring attractor present closeby in parameter space, and that the overlap between the external cue and the connectivity structure adjusts the velocity of the dynamics along it. Altogether, we identified a novel dynamical mechanism for temporal flexibility based on intrinsically generated dynamics that is robust, general, and can be fully described analytically.

[1] Wang, J., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature neuroscience*, 21(1), 102.

[2] Mastrogiuseppe, F., & Ostojic, S. (2018). Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*, 99(3), 609-623.

---

\*Speaker



# Representing a background signal in spike coding networks allows for response normalization

Sander Keemink \* <sup>1</sup>, Nuno Calaim <sup>1</sup>, Christian Machens <sup>1</sup>

<sup>1</sup> Champalimaud Centre for the Unknown – Portugal

Brains are able to represent real-world information with incredible precision. However, this precision is not fixed: it can vary due to changes in stimulus strength and attention. 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 range of possible responses is shifted, usually dependent on stimulus features or attention. This allows for adjusting response sensitivity, and is also thought to be involved in attention-based response changes (Reynolds and Heeger, 2009). Neural networks based on spike coding (Spike Coding Networks, or SCNs) have been shown to be remarkably efficient in their representation of stimuli (Boerlin et al., 2013, Denève and Machens, 2016), but they do not currently have a natural way to adjust neural sensitivity or coding precision without changing the connectivity. SCNs also have several biological inconsistencies: the recurrent and cross-area connections mix inhibition and excitation from the same neuron (both violating Dale’s law), and the recurrent excitation can cause the network to be unstable in the presence of noise (giving rise to ‘epileptic’ seizures). We propose that all these problems can be solved by representing an additional background signal equally across the network. The background signal then controls the overall firing rate, but the network also gains several new properties. As the background signal increases, both the population firing rates and the decoding precision increase, mimicking experimental attention-based responses and precision changes. As stimulus strength increases, population firing rates increase but the decoding precision decreases, mimicking experimental stimulus-based responses and precision changes. Simply by representing an additional background signal we have thus effectively implemented response normalization, with a plain-vanilla network of integrate-and-fire neurons. Moreover, we find that the adjusted spike-coding networks are more stable and that all recurrent connections are inhibitory. These results also extend to cross-area connections, which are now purely excitatory instead of mixed. The background signal can be interpreted as a feedback signal from a downstream area, or a feed-forward signal based on the average stimulus magnitude. In either case the background signal causes shifts in tuning curves of individual neurons, similar to what one would expect from attentional (feedback) or contextual effects (feed-forward). We provide a geometrical intuition to all these effects by showing that the network bounds the decoding error to be within a cone-like polyhedron. Both the background signal and the stimulus strength then move the network state ‘up’ or ‘down’ in this cone, effectively changing both the precision and overall firing rate. The resulting networks present an intriguing hypothesis for brain function, and allow for new predictions on the effects of stimulus strength, attention, and response normalization.

---

\*Speaker

# Maintenance of visual working memory by a visuo-premotor cortical loop

Ivan Voitov \* <sup>1</sup>, Thomas Mrsic-Flogel <sup>1</sup>

<sup>1</sup> Sainsbury Wellcome Centre – United Kingdom

Connections between cortical areas are almost always reciprocal. One function of such cortical loops may be to maintain coherent representations between functionally distinct cortical areas. We examined this hypothesis using a novel working memory task for mice and projection target-specific two-photon calcium imaging with simultaneous optogenetic perturbations. Based on delayed match-to-sample tasks in the primate literature, our task operationalized the maintenance of visual stimuli in working memory during a seconds-long delay period, while controlling for motor preparation, visual stimulus history, and reward expectancy. Using focal optogenetic silencing during the delay period, we found that visual cortical area AM and secondary motor cortex (M2), which are reciprocally connected, both contributed to the maintenance of visual working memory. Population recordings across multiple cortical areas revealed highly distributed representations of task variables, including working memory engagement and strength. As a more direct test for the instantaneous necessity of activity in the visuo-premotor loop for memory maintenance, we silenced area AM (‘feedforward’) while imaging the M2 (‘feedback’) axons in AM during the delay periods. Preliminary data show that the delay activity was selectively inhibited during blocks of working memory engagement. Our results therefore support the role of reciprocal corticocortical interactions in maintaining dynamic cognitive representations during visual working memory.

---

\*Speaker

# Traveling waves shape neural computations in vision

Lyle Muller \* <sup>1</sup>, Davis Zachary <sup>2</sup>, John Reynolds <sup>2</sup>, Terry Sejnowski <sup>2</sup>

<sup>1</sup> Department of Applied Mathematics, University of Western Ontario – Canada

<sup>2</sup> The Salk Institute for Biological Studies – United States

New recording technologies allow neuroscientists to record from cortex with high spatial and temporal resolution. For the first time, we can visualize the complex activity patterns in cortical populations during natural sensory behaviors. Because these imaging experiments occur in intact biological systems, however, certain restrictions are inevitable. In particular, the signal-to-noise ratio (SNR) remains low relative to other scientific imaging domains.

In our research, we have developed new signal processing techniques to analyze fine-scale spatiotemporal dynamics in noisy multisite recordings from neocortex. With these tools, we have found unexpected structure in the dynamics of cortical populations during normal sensory behavior. First, by applying our algorithms to voltage-sensitive dye (VSD) imaging data in V1 of the awake macaque, we found that small visual stimuli consistently evoke waves traveling outward from the point of thalamocortical input to cortex (Muller et al., *Nature Communications* 5, 2014). These waves were previously obscured in trial-averaged data, and they are an important finding because they show local neuron populations influence networks far across V1 in a highly structured manner. This in turn has consequences for theories of sensory processing, and we have introduced a conceptual framework for thinking about computations with these waves (Muller et al., *Nature Reviews Neuroscience* 19, 2018).

In collaboration with Zachary Davis and John Reynolds (Salk Institute), we have recently introduced a signal processing approach to handle the spectrally broad, non-frequency-resolved fluctuations present during spontaneous background activity in vivo. This approach is suited to track nonstationary fluctuations in the local field potential (LFP), where frequency, amplitude, and phase can change dramatically from moment to moment. Using this approach, we find that ongoing activity in extrastriate visual cortex of awake, behaving marmosets is organized into spontaneous traveling waves. These waves are spectrally broad, modulate spiking activity, and directly impact the processing of sensory inputs. In monkeys trained to report the appearance of a visual stimulus near perceptual threshold, the timing and position of spontaneous traveling waves - before the target is presented - predicts the monkey's ability to detect the target.

Taken together, these results indicate that traveling waves shape neural computations during normal vision. This unexpected spatiotemporal structure has general implications for the way we think about noise in the brain. These results will be discussed in the context of ongoing experimental collaborations and computational modeling work.

---

\*Speaker

# Developmental and evolutionary principles on olfactory circuit designs

Naoki Hiratani \* <sup>1</sup>, Peter Latham <sup>1</sup>

<sup>1</sup> Gatsby Computational Neuroscience Unit, University College London – United Kingdom

A fundamental goal of neuroscience is to understand the designing principle of neural circuits. Ideally, we would like to be able to predict the anatomical configuration of a circuit by optimizing its performance. Here we apply this approach to the olfactory system. We first show that a scaling law holds between the number of glomeruli and the number of mushroom body Kenyon cells among seven species of invertebrates, through a literature survey. Notably, the exponent of the scaling is more than cubic, as opposed to the  $3/2$  law previously observed among the mammals. Secondly, we model the olfactory system as a three-layered nonlinear neural network, and analytically derive the scaling laws by estimating the network size that optimizes, over the lifetime of the animal, the ability to identify odors. Although having more neurons increases the information capacity, having too many neurons makes developmental tuning of synaptic weights difficult due to potential overfitting. Applying this tradeoff, the estimated optimal population sizes robustly follow a scaling law similar to the one observed among the mammals. The scaling emerges in full batch optimization and stochastic gradient learning, under various choices of non-linearity. We further extend the framework to the case when a fraction of the olfactory circuit can be genetically specified, not developmentally learned, and numerically demonstrate that this evolutionary constraint makes the scaling steeper when there are a small number of glomeruli, as is observed among the invertebrates. Overall, our study provides an important step for the theoretical understanding of neural circuit designs.

---

\*Speaker

# Functional segregation of the ferret auditory cortex probed with natural and model-matched sounds

Agnès Landemard \* <sup>1</sup>, Célian Bimbard <sup>1</sup>, Sam Norman-Heigneré <sup>1,2</sup>,  
Shihab Shamma <sup>1</sup>, Yves Boubenec <sup>1</sup>

<sup>1</sup> Laboratoire des Systèmes Perceptifs, Équipe Audition – Ecole Normale Supérieure de Paris - ENS  
Paris – France

<sup>2</sup> Zuckerman Institute of Mind, Brain and Behavior, Columbia University – United States

Sensory systems are adapted to extract and encode relevant features from natural stimuli despite their complexity. However, how this information is organized spatially in the brain 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 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 in head-fixed ferrets.

We first mapped the classical tonotopy of ferret auditory cortex, highlighting core and belt regions. We then used a computational approach to confront auditory cortex responses with acoustic models of different complexity levels, based on cochlear and spectro-temporal modulation filters. We contrasted brain responses to original natural sounds and to synthetic sounds matching either part or all low-level original acoustic features. Doing so, we were able to reveal functionally distinct subregions in auditory cortex, based on either temporal or spectral features.

We compared fMRI responses in humans and fUS recordings in ferrets to speech/music and their model-matched counterparts. Interestingly, we observed speech selective regions in the ferret auditory cortex. However, and contrary to the real speech- and music-selective response components observed in human non-primary regions (Norman-Haigneré, 2015/2018), model-matched stimuli evoked similar responses in the ferret. Because speech and music are not ecologically relevant sounds for ferrets, we wanted to test whether ferret auditory cortex could discriminate between ferret pup vocalizations and their corresponding model-matched versions. We observed differences in animal motor activity for original vocalizations compared to model-matched stimuli, indicating that the animal is able to perceptually discriminate these two classes of sounds. We are currently investigating the neural correlates of this capability in auditory cortex responses. Follow-up work will test if ferrets can innately discriminate original vs synthetic speech, or whether perceptual learning is necessary to do so.

---

\*Speaker

# Constrained plasticity can compensate for ongoing drift in the parietal cortex

Michael E. Rule <sup>\*</sup> <sup>1</sup>, Adrianna R. Loback <sup>1</sup>, Dhruva V. Raman <sup>1</sup>,  
Christopher D. Harvey <sup>2</sup>, Timothy S. O’leary <sup>1</sup>

<sup>1</sup> University of Cambridge [UK] – United Kingdom

<sup>2</sup> Harvard Medical School – United States

Recent experiments reveal that neural populations underlying behavior reorganize their tunings over days to weeks, even for routine tasks. How can we reconcile stable behavioral performance with ongoing reconfiguration in the underlying neural populations? We examine drift in the population encoding of learned behaviour in posterior parietal cortex of mice navigating a virtual-reality maze environment. Over five to seven days, we find a subspace of population activity that can partially decode behaviour despite shifts in single-neuron tunings. Additionally, directions of trial-to-trial variability on a single day predict the direction of drift observed on the following day. We conclude that day-to-day drift is concentrated in a subspace that could facilitate stable decoding if trial-to-trial variability lies in an encoding-null space. However, a residual component of drift remains aligned with the task-coding subspace, eventually disrupting a fixed decoder on longer timescales. We illustrate that this slower drift could be compensated in a biologically plausible way, with minimal synaptic weight changes and using a weak error signal. We conjecture that behavioral stability is achieved by active processes that constrain plasticity and drift to directions that preserve decoding, as well as adaptation of brain regions to ongoing changes in the neural code.

---

\*Speaker

# The emergence of multiple retinal cell types through efficient coding of natural movies

Stéphane Deny \* <sup>1</sup>

<sup>1</sup> Stanford University – Stanford University – France

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.

---

\*Speaker

# Dimensionality of parallel fibre population activity in awake mice

Alex Cayco Gajic \* <sup>1</sup>, Frederic Lanore <sup>2</sup>, Angus Silver <sup>1</sup>

<sup>1</sup> University College, London – United Kingdom

<sup>2</sup> Interdisciplinary Institute for Neuroscience – Université de Bordeaux, Centre National de la Recherche Scientifique : UMR168 – France

Classic work by Marr and Albus posited that the cerebellar cortex separates overlapping input patterns by projecting them onto a much larger number of granule cells, thereby increasing the dimensionality of their representation. 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, and whether granule cell populations can support high-dimensional representations, is poorly understood. To address this, we expressed GCaMP6f in granule cells in the Crus 1 region of the cerebellar hemisphere. We used a high-speed random-access 3D 2-photon microscope to simultaneously monitor the Ca<sup>2+</sup> activity in hundreds of parallel fibres (PFs) of awake animals for the first time. We find that PF population activity transitions between separate, orthogonal coding spaces representing periods of quiet wakefulness vs. active movement. Finally, we characterise the dimensionality and spatial properties of PF population activity.

---

\*Speaker



# Basic principles of neural ensemble organization for multi-neuronal coding

Ioannis Smyrnakis \* <sup>1</sup>, Anna Palagina <sup>2,3</sup>, Stelios Smirnakis \* <sup>4,5</sup>, Maria Papadopouli \* <sup>1,6</sup>

<sup>1</sup> Foundation for Research and Technology - Hellas - Greece

<sup>2</sup> Jamaica Plain VA hospital, Harvard Medical School - United States

<sup>3</sup> Brigham and Women's Hospital, Harvard Medical School - United States

<sup>4</sup> Brigham and Women's Hospital, Harvard Medical School [Boston] - United States

<sup>5</sup> Jamaica Plain Veterans Administration Hospital, Harvard Medical School, Boston MA - United States

<sup>6</sup> University of Crete - Greece

This work presents a simple modeling framework that explores the principles that govern how multi-neuronal ensembles encode sensory stimuli in order to achieve responses that are definite, with high reliability, robustness, and information-capacity. The existence of single neurons that receive unreliable input from many cells and integrate this into a definite response ("grandmother-cells") does not appear to be a general feature of the neural code. Such cells have not been spotted in early-visual areas, like V1, and are elusive even in higher areas. Hence, one is led to believe that definitive information is carried by a group of unreliable neurons whose aggregate, approximately simultaneous, output responds reliably to a given stimulus (e.g. an orientation grating in the case of area V1). We call this group of neurons an information-pathway or simply pathway. Neurons in such pathways respond probabilistically to the stimulus encoded by the pathway, while different pathways consist of sets of neurons that partially overlap. For the pathways to reliably encode the information two conditions need to be satisfied: (i) The encoding must be definite, i.e., the probability of a pathway to be active, when its preferred-stimulus is present, should be close to 1, versus close to 0, when the corresponding stimulus is not present, and (ii) There should be no significant interference among overlapping pathways tuned to cover the range of different stimuli that can be perceived. When constraints (i) and (ii) are satisfied but locality is not enforced, i.e. neurons of a pathway are selected randomly with equal probability among the total population, all pathways overlap with each-other and there is no well-positioned pathway-center structure (lattice). In this model, the number of non-interfering pathways grows as a power in the number of neurons at the cost of being unable to maintain a topographical mapping. A more realistic arrangement assumes that pathways are local, i.e. consist of neurons that belong to the same approximate cortical-vicinity. At one limit, taking local pathways to be dense (i.e. all neurons within a radius from the pathway center belong to the pathway) severely limits the number of pathways that are non-interfering, reducing information-capacity. Relaxing this requirement, to require only that a fraction of neurons in the pathway-neighborhood belong to the pathway (sharp-cutoff constant-density "dilute-pathway" model), results in a linear dependence of the number of pathways on the number of neurons and, as "diluteness" increases, in an exponential dependence on the maximum permissible overlap between pathways, significantly increasing capacity without compromising topography. We compute the maximum number of pathways for the different models and comparatively analyze their

---

\*Speaker

advantages and disadvantages in terms of reliability, robustness, and information-capacity. In general, when pathways become sufficiently sparse, the number of distinct information channels increases dramatically, boosting information-capacity. Finally, we examined potential realizations of pathway-models and corresponding predictions in the context of datasets collected from area V1 of adult mice using two-photon imaging[Palagina\_et\_al\_Journal\_of\_Neuroscience\_2019]. A special pathway-arrangement of interest consists of a set of pyramidal neurons functionally-connected with each other as well as with 1-2 inhibitory-interneurons, which provide control[Palagina\_et\_al\_Journal\_of\_Neuroscience\_2019].

# P1. Does structure in neural activity match anatomical structure?

Thomas Delaney \* <sup>1</sup>, Cian O'donnell <sup>1</sup>

<sup>1</sup> Faculty of Engineering [Bristol] – United Kingdom

Information in the brain is carried in correlated network activity. Previous research has established that these correlations play a crucial role in representing sensory information. For example, the onset of visual attention has been shown to have a greater affect on the correlations in the macaque V4 than on the firing rates in that region. In order to understand the representation of sensory information we must understand the interactions between neurons. Because of limitations in recording technology almost all research has explored correlations between neurons within a given brain region. Relatively little is known about correlations between neurons in different brain regions. However, the recent development of ‘Neuropixels’ probes has allowed extracellular voltage measurements to be collected from multiple brain regions simultaneously, routinely, and in much larger numbers than traditional methods. In this project we used a publicly-available Neuropixels dataset to analyse correlations between different brain regions.

Using two probes, spiking activity was simultaneously collected from over 800 neurons in the brain of an awake mouse for a period of 84 minutes. During this period, the mouse and was shown various visual stimuli. The 800 neurons were distributed across 5 different brain regions: V1, hippocampus, thalamus, motor cortex, and striatum. Using these data, we measured single neuron firing rates across all regions. We also measured pairwise correlations and mutual information between neurons within the same region, and between neurons in different regions. We compared the distribution of these measurements across regions.

We then used spectral clustering to cluster the neurons based on the pairwise correlations and mutual information. We compared the structure of these cluster networks to the anatomical structure of the cells.

---

\*Speaker

## P2. Associative memory in winner-take-all networks: from binary units to spikes

Roman Pogodin \* <sup>1</sup>, Peter Latham <sup>1</sup>

<sup>1</sup> Gatsby Computational Neuroscience Unit, University College London – United Kingdom

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". The model was potentially powerful: the memory capacity scaled with the number of connections per neuron,  $K$ , and inversely with the sparsity of individual memories (i.e. how many neurons encode it). However, as shown theoretically by Roudi and Latham (2007), Hopfield networks implemented with spiking neurons cannot store arbitrarily sparse memories, mainly because it leads to unrealistically high firing rates. As a result, memory capacity is proportional to the number of connections per neuron  $K$ , and simulations with spiking networks have never, to our knowledge, yielded more than several hundred memories; far too few to explain human performance. Building upon a recent model of Gripon and Berrou (2011), we overcome this problem with a network of coupled winner-take-all subnetworks. We extend their work on binary networks in several directions: we derive explicit expressions for capacity as a function of the size of the basin of attraction and the signal-to-noise ratio, and we find that for a reasonable choice of parameters the number of memories is  $KN / 10^4$ , where  $N$  is the number of neurons (large enough to explain human level memory); we show that high capacity can be achieved with almost no spurious memories; and we explain intuitively why the model has such high capacity. To test for biological plausibility, we translate the binary model to a balanced excitation-inhibition spiking network. We confirm the theoretical predictions in large scale simulations with Poisson neurons. Extrapolating these results for  $K=1000$ , this implies that  $0.1N$  memories can be embedded, comparable to the vanilla Hopfield network with (unrealistic) all-to-all connectivity.

---

\*Speaker

# P3. Presynaptic inhibition rapidly stabilised recurrent excitation in the face of plasticity

Laura Bella Naumann \*<sup>1,2</sup>, Henning Sprekeler<sup>1,2</sup>

<sup>1</sup> Technische Universität Berlin – Germany

<sup>2</sup> Bernstein Center for Computational Neuroscience [Berlin] – Germany

Hebbian plasticity, a mechanism believed to play a key role in learning and memory, detects and further enhances correlated neural activity. In recurrent networks this constitutes an inherently unstable positive feedback loop and therefore requires additional homeostatic control [1]. Recent computational work indicates that slow homeostasis, as observed in experiments, is insufficient to compensate the instabilities arising from Hebbian plasticity in recurrent neural networks [2].

We suggest presynaptic inhibition as a compensatory process, which does not suffer from this discrepancy of timescales. Experimental studies have revealed that excess network activity can trigger inhibition of transmitter release at excitatory synapses through activation of presynaptic GABAB receptors [3]. This effectively and reversibly attenuates recurrent interactions on timescales of 100s of milliseconds, thus serving as a candidate mechanism for the rapid compensation of elevated recurrent excitation induced by Hebbian changes.

To study the network effects of presynaptic inhibition, we analyzed a rate-based recurrent network model, in which presynaptic inhibition is mimicked by a multiplicative reduction of recurrent synaptic weights in response to increasing firing rates. Using analytical and numerical methods, we show that presynaptic inhibition ensures a gradual increase of firing rates with growing recurrent excitation, even for very strong recurrence, whereas classical subtractive postsynaptic inhibition is unable to control recurrent excitation once it has surpassed a critical strength. Moreover, we find that presynaptic inhibition stabilizes firing rates in a recurrent population subject to Hebbian plasticity, while allowing synaptic homeostasis to operate on biologically plausible timescales.

In summary, the multiplicative character of presynaptic inhibition provides a powerful compensatory mechanism to rapidly reduce effective recurrent interactions. As it conserves the underlying network connectivity, presynaptic inhibition might therefore set the stage for stable learning without interfering with plasticity at the level of single synapses.

[1] Abbott LF and Nelson SB. Synaptic plasticity: taming the beast. *Nat Neurosci* 2000, 3:1178-1490.

[2] Zenke F and Gerstner W. Hebbian plasticity requires compensatory processes on multiple timescales. *Phil Trans R Soc B* 2017, 372(1715): 20160259.

---

\*Speaker

- [3] Urban-Ciecko J, Fanselow EE, Barth AL. Neocortical somatostatin neurons reversibly silence excitatory transmission via GABA<sub>B</sub> receptors. *Curr Biol* 2016, 25(6):722-731
- [4] Zenke F, Hennequin G and Gerstner W. Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS computational biology* 2013, 9(11), e1003330.

# P4. Statistical symmetries in connectivity shape slow activity manifolds in recurrent neural networks

Francesca Mastrogiuseppe <sup>\*</sup> <sup>1</sup>, Srdjan Ostojic <sup>2</sup>

<sup>1</sup> Gatsby Computational Neuroscience Unit, UCL – United Kingdom

<sup>2</sup> LNC2 – Département d'études cognitives, École normale supérieure, PSL Research University, CNRS  
– France

To process information and produce adaptive behavior, the brain represents the external world in terms of abstract quantities such as value, position, or orientation. Increasing experimental evidence suggests that neural circuits encode such continuous, topologically-organized quantities by means of the collective organization of neural activity along non-linear, low-dimensional manifolds in the space of possible network states [1,2,3]. In higher order brain areas, these manifolds persist in absence of sensory stimuli [3], and are therefore presumably generated by intrinsic recurrent interactions. How recurrent connectivity gives rise and shapes activity manifolds is however not fully understood.

The most prominent models of recurrently-generated manifolds are continuous attractor networks. In these models, the emergence of activity manifolds typically relies on strong and highly ordered structure in the synaptic connectivity. For instance, in the classical bump attractor model [4,5,6] a ring-like manifold of fixed points relies on a distance-dependent, bell-shaped connectivity, which is itself ring-like. While such tightly structured connectivity has recently been identified in the fly brain [7], it remains challenging to reconcile classical attractor networks with circuits in the mammalian cortex, where low-dimensional activity organization co-exists with highly heterogeneous connectivity and single-cell activity. In this work, we asked how much structure is required and expected in the connectivity and in the activity of a recurrent neural network which generates low-dimensional activity manifolds.

We considered a large class of recurrent networks in which the connectivity can be expanded in terms of rank-one components [8]. By studying analytically the emergent dynamics, we found that hidden statistical symmetries in the distribution of connectivity weights generate a fundamental degeneracy in the dynamics that leads to the appearance of slow activity manifolds in the neural state space. In the specific case of classical ring models, the connectivity is fully ordered and specified by the symmetry itself; more in general, though, the connectivity can include strong additional variance along irrelevant directions which are orthogonal to the symmetry. Statistical symmetries can arise in absence of precise constraints, as in the example of spherical symmetry that emerges from iid Gaussian variables, and therefore require very little fine-tuning. We found that connectivity symmetries fully specify the shape and the topology of activity manifolds in the high-dimensional neural state space. The intrinsic dimensionality of the manifold is determined by the number of parameters defining the symmetry, while the embedding dimensionality is determined by the symmetry matrix representation. Importantly,

---

<sup>\*</sup>Speaker

the variance of the connectivity distribution along irrelevant directions introduces significant heterogeneity in population activity and tuning curves. As a result, the symmetry which generates the manifold does not prominently manifest itself neither in the synaptic connectivity, nor in the single-unit activity.

- [1] Gallego, Solla et al., *Neuron* (2017)
- [2] Low, Tank et al., *bioRxiv* (2018)
- [3] Chaudhuri, Fiete et al., *bioRxiv* (2019)
- [4] Ben-Yishai, Bar-Or and Sompolinsky, *PNAS* (1995)
- [5] Burak and Fiete, *PNAS* (2012)
- [6] Wimmer, Compte et al., *Nat Neurosci* (2014)
- [7] Kim, Jayaraman et al., *Science* (2017)
- [8] Mastrogiuseppe and Ostojic, *Neuron* (2018)



# P5. Transient coding through non-normal dynamics during OFF responses in auditory cortex

Giulio Bondanelli \* <sup>1</sup>, Srdjan Ostojic <sup>2</sup>

<sup>1</sup> Laboratoire de Neurosciences Cognitives et Computationnelles – Ecole Normale Supérieure de Paris (FRANCE) – France

<sup>2</sup> Laboratoire de Neurosciences Cognitives et Computationnelles – Ecole Normale Supérieure de Paris (FRANCE) – France

Across sensory systems, complex spatio-temporal patterns of neural activity can arise following the appearance (ON) and disappearance (OFF) of simple stimuli. While ON responses have been widely studied, the mechanisms generating OFF responses in cortical areas have so far not been fully elucidated. Recent studies have argued that OFF responses, may reflect strongly transient sensory coding at the population level, suggesting that transient OFF responses might be generated by a collective, network mechanism. Here we examine the hypothesis that OFF responses are single-cell signatures of non-normal dynamics in the underlying recurrent network and propose a network model that generates transient OFF responses through recurrent interactions. To test this hypothesis, we analyse the responses of a large population of neurons in primary auditory cortex recorded using two-photon calcium imaging in awake mice passively listening to different auditory stimuli. Focusing on linear recurrent networks, we determine the conditions for the existence of strong OFF responses. We identify a general criterion that allows us to distinguish two sharply separated dynamical regimes depending on properties of the connectivity matrix: a regime in which OFF responses cannot occur, and a regime in which specific stimuli elicit strong OFF responses based on non-normal network dynamics. Directly fitting a linear dynamical system to recorded neural activity, we find that the criterion for network-generated OFF responses is borne out by the connectivity matrix inferred from the data. Moreover, the population dynamics elicited by individual stimuli are low-dimensional. Combining these two observations, we analytically show that OFF responses generated by low-dimensional, non-normal network dynamics can transiently encode an extensive number of stimuli.

---

\*Speaker

# P6. From an ethogram to a model of neuronal computation: the optomotor response in zebrafish larva.

Adrien Jouary \* <sup>1</sup>, Christian Machens <sup>1</sup>, Michael Orger <sup>1</sup>

<sup>1</sup> Champalimaud Centre for the Unknown – Portugal

Taking advantage of recent advances in optical imaging and genetics, much has been learned about the sensorimotor circuits underlying the optomotor response in the larval zebrafish. Yet, to avoid overfitting, the relation established between task parameters (e.g. visual stimuli or motor actions) and neuronal activity is often restricted to an instantaneous and linear mapping. This mapping fails to account for two essential characteristics of neuronal computation: non-linearity and recurrent dynamics. These ingredients are essential to dissociate sensorimotor transformations from simpler reflex chains. Our aim is to build a model of sensorimotor computation that uses recurrent dynamics and non-linearities in order to account for the influence of past actions, sensory experiences, or internal states on the animal's actions.

We first collected a very large behavioral dataset (> 400 hours). Whole field motions was presented below the larva at specific speeds and orientations with respect to the animal's field of view while using automatic tracking, stereotypical movements were identified. Second, we developed a model capable of predicting the motor activity recorded in response to visual stimuli. To model the probability of observing a given tail movements, we used a generalized linear model (GLM). This model has been previously applied to study diverse neuronal populations such as retinal ganglion cells during visual processing or lateral intraparietal cortex (LIP) during perceptual decision-making. We designed a variation of this model suitable to account for motor actions instead of spike trains. It is composed by the following elements:

- A stimulus filter for each of the stereotypical maneuvers. This spatio-temporal receptive field reduces the dimensionality of visual input to the components most likely to elicit a given action. Preliminary analysis of responses to whole-field motion revealed that each of the most frequent maneuvers was associated with a distinct spatial filter.

- A set of feedback coupling filters, that capture dependencies on recent actions performed by the larva. This component is critical to introduce recurrent dynamic and generate sequential motor actions with complex temporal dynamics as observed in the fish behavior.

- For each maneuver, the sum of filtered responses are passed through a non-linearity to obtain the instantaneous probability of each tail movements.

This model presents two advantages: first, the influence of each filter can be meaningfully interpreted and second, its convexity guaranties a unique mapping between the network model and the observed behavior. We validated this algorithm and its ability to converge on a synthetic

---

\*Speaker

dataset and we are currently integrating the behavioral dataset collected. We expect two main results from this study: First, we hypothesized that the spatio-temporal receptive fields in the model will reveal the selectivity of neuronal pathways between groups of early visual neurons and specific motor behaviors. Second, we await to learn how much a decision to select a motor action is influenced by the contribution of external sensory cues relative to the weight of internal causes (reflected in the model by the history terms).

# P7. Disentangling the roles of dimensionality and cell classes in neural computations

Alexis Dubreuil \* <sup>1</sup>, Francesca Mastrogiuseppe <sup>2</sup>, Srdjan Ostojic <sup>3</sup>

<sup>1</sup> Ecole normale supérieure (ENS Paris) – École normale supérieure - Paris – 45, rue d’Ulm 75230 Paris cedex 05, France

<sup>2</sup> Gatsby Computational Neuroscience Unit, UCL – United Kingdom

<sup>3</sup> Group for Neural Theory, Laboratoire de Neurosciences Cognitives – Ecole Normale Supérieure de Paris - ENS Paris – 29 rue d’ilm, Paris, France, France

The description of neural computations in the field of neuroscience relies on two competing views: (i) a classical single-cell view that relates the activity of individual neurons to sensory or behavioural variables, and focuses on how different cell classes map onto computations; (ii) a more recent population view that instead characterises computations in terms of collective neural trajectories, and focuses on the dimensionality of these trajectories as animals perform tasks. How the two key concepts of cell classes and low-dimensional trajectories interact to shape neural computations is however currently not understood. Here we address this question by combining machine-learning tools for training RNNs with reverse-engineering and theoretical analyses of network dynamics. We introduce a novel class of theoretically tractable recurrent networks: low-rank, mixture of Gaussian RNNs. In these networks, the rank of the connectivity controls the dimensionality of the dynamics, while the number of components in the Gaussian mixture corresponds to the number of cell classes. Using back-propagation, we determine the minimum rank and number of cell classes needed to implement neuroscience tasks of increasing complexity. We then exploit mean-field theory to reverse-engineer the obtained solutions and identify the respective roles of dimensionality and cell classes. We show that the rank determines the phase-space available for dynamics that implement input-output mappings, while having multiple cell classes allows networks to flexibly switch between different types of dynamics in the available phase-space. Our results have implications for the analysis of neuroscience experiments and the development of explainable AI.

---

\*Speaker

# P8. Processing information stored in working memory through modulations of effective connectivity

Adrian Valente \* <sup>1</sup>, Alexis Dubreuil <sup>1</sup>, Srdjan Ostojic <sup>1</sup>

<sup>1</sup> Laboratoire de Neurosciences Cognitives et Computationnelles – Ecole Normale Supérieure de Paris - ENS Paris – France

Working memory is a cognitive process responsible for holding information about items for further processing. In a typical laboratory task such as the delay-match-to-sample task, an animal is asked to memorize a first stimulus, and, as a second stimulus is presented, to compare the identity of these two stimuli. While classical works have examined mechanisms for the working memory part of the task, how the stored information is extracted and exploited for the task has so far not been elucidated. Here we aim at providing a mechanistic description that accounts for the processing of information during the comparison phase of DMS tasks. To this end, we use machine learning tools to train low-rank recurrent neural networks on these tasks. Exploiting theoretical results linking low-rank synaptic structure and network dynamics, we reverse-engineer trained networks and describe their function in the framework of dynamical system theory. We show that when the system performs the comparison operation, the effect of inputs corresponding to the second stimulus is to change the effective connectivity of the network, so as to transiently redefine the phase portrait of the neural system, and allow relevant rotations towards fixed-points encoding Match or Non-Match decisions.

---

\*Speaker

# P9. Predictive coding in the salamander retina

Danica Despotović \* <sup>1</sup>, Olivier Marre <sup>1</sup>, Matthew Chalk <sup>1</sup>

<sup>1</sup> Institut de la Vision – Institut National de la Santé et de la Recherche Médicale : U968, Sorbonne Université : UM80, Centre National de la Recherche Scientifique : UMR7210 – France

A central question in sensory neuroscience is how information about the environment is represented in the brain. The efficient coding hypothesis proposes that early sensory neurons transmit maximal information about sensory stimuli, given internal constraints (e.g. energy, wiring). This theory predicts that neurons should invest most resources into encoding stimuli that are unexpected / surprising. To investigate this, we recorded the responses of retinal ganglion cells (RGCs) to repeated sequences of full-field flashes. Previous results (Schwartz et al) showed that some ganglion cells can respond to an omitted flash in a periodic sequence of flashes (termed the omitted stimulus response, OSR), suggesting that neural activity codes for the unexpected stimuli rather than for the physical luminance. However, it is not clear if ganglion cells responses could be predicted by how surprising the stimulus is. We hypothesized that neurons would fire most strongly in response to the stimuli that were unexpected. We defined surprise rigorously and found that ganglion cell responses were predictable based on how surprising the stimulus was. Surprisingly, our results were that each ganglion did not estimate how unexpected a stimulus was according to the stimulus statistics, but given their internal model of the stimulus statistics. We found that a simple model with only 4 parameters was able to capture many aspects of recorded RGC responses. Specifically, our model could explain (i) the fact that RGCs responded strongly to the omitted stimulus i.e. the OSR, (ii) the dependence of the OSR on the number of consecutive flashes. Our results suggest that encoding surprise could be a functional objective of retinal ganglion cells.

---

\*Speaker

# P10. Disentangling encoding of stimuli features, perceptual category and behavioural choice in ferret A1 population activity

Rupesh Chillale \*<sup>1,2</sup>, Shihab Shamma<sup>1,3</sup>, Yves Boubenec<sup>1</sup>, Srdjan Ostojic<sup>2</sup>

<sup>1</sup> Laboratoire des systèmes perceptifs, Département d'études cognitives, École normale supérieure, PSL Research University, CNRS, 75005 Paris, France.

<sup>2</sup> Laboratoire de Neurosciences Cognitives, Département d'études cognitives, École normale supérieure, PSL University, INSERM, 75005 Paris, France – France

<sup>3</sup> University of Maryland – United States

In a goal directed task, sensory features of the stimulus are transformed into abstract perceptual representation for appropriate behavioural decisions. This process is classically believed to be implemented by hierarchical processing through a series of cortical areas within which primary sensory areas are the first step and supposed to encode only sensory features associated with the stimuli. Recent studies have however challenged that picture and argued that primary sensory areas may already be extracting the task-dependent, behavioral meaning of stimuli. To further determine to which extent auditory cortex represents task-dependent vs sensory properties, here we examine population representations in primary auditory cortex (A1) during a categorization task. We trained two ferrets under a Go/NoGo paradigm to classify regular click trains into Target vs. Reference categories (low and high rates with the stimulus-meaning association reversed between the two ferrets). After training, we chronically recorded A1 population activity while ferrets were either passively listening or actively discriminating stimuli. Using population-level analyses, we exploit the structure of the task to contrast the representation of stimuli features (click-rates), perceptual categories (Target/Reference) and behavioural choice (Go/No-Go). We found that neural responses in A1 encode purely sensory features during passive listening, while during active discrimination the encoding shifts from sensory to categorical. Task-engaged coding dynamically evolves between the different trial epochs, with mixed (sensory and categorical) encoding during stimulus presentation and purely categorical encoding during the delay period separating stimulus offset and the response window. During the delay period and response window, the category decoder is moreover highly correlated to the choice decoder, whereas the correlation with the sensory dimension is decreased. Altogether, our results suggest a picture in which stimuli features, perceptual category and behavioural choice are represented along different directions in the high-dimensional neural state space, so that overlaps between these directions at different trial epochs represent transformation of sensory information into categorical percepts and finally behavioural output.

---

\*Speaker

# P11. A novel framework for detecting synchronously active cell assemblies featuring extremely low non-detect error rates

Gabriel Makdah \*<sup>1,2,3</sup>, Daniel Bennequin<sup>4</sup>, Michaël Zugaro<sup>3</sup>, Ralitsa Todorova<sup>3</sup>, Sidney Wiener<sup>3</sup>, Marco Pompili<sup>3,5</sup>

<sup>1</sup> Hospices Civils de Lyon – Hospices Civils de Lyon – France

<sup>2</sup> Faculté de Médecine Lyon Est – Université Claude Bernard Lyon 1 – France

<sup>3</sup> Centre interdisciplinaire de recherche en biologie – Collège de France, Institut National de la Santé et de la Recherche Médicale : U1050, Centre National de la Recherche Scientifique : UMR7241 – France

<sup>4</sup> Géométrie et dynamique – Université Paris Diderot - Paris 7 – France

<sup>5</sup> Institut de psychiatrie et neurosciences de Paris – Université Paris Descartes - Paris 5 : UMS 1266, Institut National de la Santé et de la Recherche Médicale : UMS 1266, Hôpital Sainte Anne – France

Since higher brain functions involve the cooperation and interaction of large numbers of neurons, many tools have been developed to assess the temporal dynamics of large neuronal populations. Existing algorithms treat the assembly detection problem as a dimensionality-reduction task. As such, these methods combine or fail to detect many cell assemblies, limiting the scope and accuracy of subsequent analyses of the temporal properties, interactions, and behavioral correlates of cell assemblies.

Here we present ICAC (Iterative Cell Assembly Clustering), a novel mathematical framework which detects virtually all cell assemblies in artificially-generated data sets. ICAC can detect cell assemblies across multiple timescales and is characterized by its high sensitivity and specificity to cellular co-activations. Cell assemblies are extracted from a data set in three defining steps. First, the data set is fragmented into smaller subsets based on the eigenvectors of a non-parametric cross-correlation matrix. These subsets contain fewer activation patterns and spurious correlations. Second, these subsets are clustered with a novel, parameter-free, clustering algorithm designed to maximize the within-cluster similarity of features. Each cluster yields a candidate cell assembly which is then retained only if all of its members are significantly correlated ( $p < 0.05$ ). Third, previously detected patterns are removed from the original dataset, and the algorithm is iteratively restarted. This removal process primes the detection of weaker correlations, and the subsequent identification of progressively sparser assemblies. These iterations are carried out across multiple time windows, which allow for the detection of cell assemblies across a range of timescales.

ICAC was tested on a series of artificially-generated datasets comprised of 100 to 250 cells, which contained from 100 to 1000 cell assemblies, with some cells participating in multiple assemblies. Assembly activations occurred on a range of timescales between 0 and 200 milliseconds. Across these data sets, ICAC correctly and consistently recovered over 95% of cell assemblies.

---

\*Speaker



False positive detection rate was very low ( $< 5\%$ ). In preliminary testing on a real dataset of 285 cortical neurons recorded in a behaving rat, ICAC detected 373 assemblies, a seven-fold increase compared to conventional algorithms. These assemblies were comprised of 3 to 7 cells, and were active in timescales ranging between 5 and 150 milliseconds. On average, individual neurons participated in five assemblies, with 40% of cells participating in at least two assemblies. This algorithm's major advantages are its high detection rate, including sparsely active and overlapping assemblies as well as its capacity to detect patterns over different timescales. The mathematical basis of conventional methods limits the number of assemblies that can be detected, and our approach is consistent with the participation of individual neurons in multiple and diverse processes. Subsequent analyses should determine the functional significance of the multitude of assemblies detected with this method in terms of their interactions and behavioral correlates.

## P12. Operational range of sensory response variability during neurodevelopment

Beatriz E. P. Mizusaki <sup>\*</sup> <sup>1</sup>, Nazim Kourdougli <sup>2</sup>, Anand Suresh <sup>2</sup>, Carlos Portera-Cailliau <sup>2</sup>, Cian O'donnell <sup>1</sup>

<sup>1</sup> University of Bristol – United Kingdom

<sup>2</sup> UCLA – United States

In the context of neuropsychiatric diseases not many studies have explored the link between affected components and effects over information representation at population level. In this work we address Fragile X Syndrome, which is the leading single gene cause of autism spectrum disorder. A common symptom includes enhanced reactions to sensory stimuli, a behaviour also present in a mouse model of this condition (knock out of the Fmr1 gene). We aimed to test a recent hypothesis that it could be caused by a higher variability among sensory representations in a hyperexcitable circuit. Our group performed longitudinal in vivo 2-photon calcium imaging of spontaneous and sensory-evoked activity of layer 2/3 neurons in S1 of lightly anaesthetized mice at two ages, postnatal day (P) 14 and P19, using GCaMP6s. These ages encompass the end of a developmental, critical period for the experience-dependent maturation of S1 that coincides with peak FMRP expression in the wild type (WT) mice. The same imaging protocol was performed in both WT and Fmr1 KO animals, recording activity both at rest and under different types of whisker stimulation. We were therefore able to track the activity of individual neurons in both time points (10-40 cells), which enabled us to assess individual contributions to population patterns. Our basic approach to analysing the data was to fit and compare various statistical models, specially the coupling between single cell activity and population representations. We contrasted spontaneous and evoked activity in both ages in order to compare representation stability for WT and KO. For this we tested whether the resulting models were compatible with decoding from data samples. Finally, each recording was repeated twice in a row, which enabled us to also analyse trial-to-trial variability.

---

\*Speaker

# P13. Adaptive optimisation of visual prosthetic stimulation

Tristan Fauvel \* <sup>1</sup>, Olivier Marre <sup>1</sup>, Matthew Chalk <sup>1</sup>

<sup>1</sup> Sorbonne Université, INSERM, CNRS, Institut de la Vision, 17 rue Moreau, F-75012 Paris, France  
tristan.fauvel@inserm.fr

Retinal prosthetics are a promising strategy to restore vision in patients suffering from retinal degenerative diseases. These implantable devices are designed to replace photoreceptors in the eye by providing direct electrical stimulation to retinal ganglion cells (RGCs; the cells that form the optic nerve) [1]. However, currently, the visual resolution that can be achieved with such implants is very low, limiting their use in real world tasks. This is due to e.g., the relatively small number of electrodes which each stimulate many different neighbouring cells. Here, we asked whether it is possible to improve the performance achieved by these devices by optimising how visual inputs are converted into electrical stimulation (the ‘encoder model’). Our idea was to use subjects’ responses in a range of visual tasks to optimise the encoder model online.

To see whether this could work in principle, we used a previously developed model [2] to simulate the visual percept experienced by patients fitted with a prosthetic device. This ‘simulated percept’ was displayed to healthy subjects, who were required to perform a range of visual tasks (e.g. locating a presented visual bar), and their responses used to optimise the encoder model online.

A major challenge was to optimise the encoder model (consisting of thousands of parameters) using a limited number of subjects’ responses. To do this, we learned a gaussian process model of subjects’ responses [3] which allowed us to: (i) efficiently use subjects’ responses to optimise the encoder model, while avoiding over-fitting; (ii) present only ‘maximally informative stimuli’ [4], that were most useful for optimising the encoder model. Our initial results indicate that our framework can be used to improve subjects performance in specific visual tasks (e.g. discriminating the position of presented visual bars). In the future, we will investigate how these improvements could transfer to other visual tasks, that were not explicitly presented.

[1] Henri Lorach, Olivier Marre, José-Alain Sahel, Ryad Benosman, and Serge Picaud. Neural stimulation for visual rehabilitation: Advances and challenges. *Journal of Physiology-Paris*, 107 (5) : 421-431,2013 .

[2] Michael Beyeler, Devyani Nanduri, James D Weiland, Ariel Rokem, and Ione Fine. A model of ganglion axon pathways accounts for percepts elicited by retinal implants. *bioRxiv*, 2018.

---

\*Speaker

[3] Carl Edward Rasmussen and Christopher K.I. Williams. Gaussian processes for machine learning. 2006.

[4] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian Active Learning for Classification and Preference Learning. *arXiv*, 2011 .

# P14. Learning low-dimensional inputs for brain-machine interface control

Jorge A. Menendez \* <sup>1,2</sup>, Peter Latham <sup>1</sup>

<sup>1</sup> Gatsby Computational Neuroscience Unit, University College London – United Kingdom

<sup>2</sup> CoMPLEX, University College London – United Kingdom

Underlying the extraordinary diversity of behaviors exhibited by many mammals is a motor system remarkably capable of adapting to completely novel environments. A dramatic example of this versatility is provided by brain-machine interfaces (BMIs), which demonstrate that animals can learn seemingly arbitrary experimenter-imposed mappings from motor cortical (M1) activity to movements.

Previous models of this phenomenon (Legenstein et al '10, Warnberg & Kumar '19) have invoked synaptic plasticity within the local M1 circuit to account for the gradual improvement in animals' ability to control a BMI. But several recent observations stand at odds with this hypothesis. Firstly, because the BMI decoder is not available to the animal's motor system, any learning algorithm it employs must use gradient-free optimization, which is known to scale badly with the number of parameters being learnt (Werfel et al '04). Learning the synaptic strengths of the many connections in the local circuit is thus incompatible with the fast timescales of learning observed under certain decoders (Sadtler et al '14, Chase et al '12, Jarosiewicz et al '08). Secondly, statistical structure in the local population activity is largely maintained throughout learning. For certain BMI decoders, it has been observed that the correlation structure between neurons is remarkably conserved (Hwang et al '13, Golub et al '18). Moreover, single neuron tuning to limb movements is conserved after learning a BMI, even for random decoders (Ganguly & Carmena '09). These observations suggest that the recurrent circuitry in the local M1 circuit remains intact.

Here we propose that learning is confined to a set of low-dimensional internal variables encoded in the upstream inputs to motor cortex. We show analytically that in linear networks the success of this learning strategy hinges on the "alignment" between the decoder and the network's dynamics. Through analysis and simulations of linear and non-linear networks, we show that good reaching performance under this model is only achievable for certain decoders. In particular, we find that constraining learning to a one- or two- dimensional space of inputs allows for low reaching error only under those decoders for which fast learning has been previously observed (Chase et al '12, Sadtler et al '14). Moreover, our model predicts that the resulting network activity during BMI control will maintain its correlational structure, as has been observed in primates controlling such decoders (Hwang et al '13, Golub et al '18, Hennig et al '18). These results suggest that a low-dimensional learning strategy may explain primates' ability to rapidly learn them, while also accounting for the remarkably preserved underlying population activity.

We additionally show that, when learning in this model is allowed to operate over 10s of input

---

\*Speaker

dimensions rather than only 1 or 2, proficient BMI control is achievable even for random decoders. We thus consider whether our model may serve as a general theory of BMI control, by making two testable experimental predictions. First, we predict that neural activity over learning should initially be confined to a few dimensions. Second, we predict that learning speed and asymptotic performance should differ substantially across reach directions.

# P15. Counterintuitive effects of excitatory and inhibitory perturbations in spike-coding networks

William Podlaski \* <sup>1</sup>, Nuno Calaim <sup>1</sup>, Christian Machens <sup>1</sup>

<sup>1</sup> Champalimaud Centre for the Unknown [Lisbon] – Portugal

Perturbation studies, in which neural activity in an area of interest is selectively manipulated during a task, have become widespread in experimental neuroscience research. However, the interpretation of such experiments may not always be straightforward, and rely upon certain assumptions about the underlying neural representations, as well as the relevant temporal and spatial scales (Jazayeri & Afraz 2017; Wolff & Ölveczky 2018). In this work, we compare the effects of excitatory and inhibitory perturbations on predictive spike-coding networks (Deneve & Machens 2016) as well as more traditional rate-coding networks. We characterize network robustness by observing how neurons respond to external manipulations depending upon their selectivity, and how the resulting performance is affected during simple tasks such as population coding and sensory input integration. Interestingly, we show that excitatory and inhibitory perturbations can have qualitatively different (non-opposing) effects on neural activity and coding performance – while excitatory perturbations generally bias network output proportional to the strength of the manipulation, spike-coding networks are much more resistant to inhibitory perturbations, which may have little to no effect. Critically, this depends upon the proportion of perturbed neurons and the relationship between their receptive fields. Our work may have important implications for the understanding of normal and perturbed network function, and could provide alternative explanations for some puzzling experimental findings related to the effects of photosuppression during decision making tasks (e.g., Fetsch et al. 2018).

---

\*Speaker

# P16. Predicting point process GLM excessively high firing rates

Dmitry Todorov \* <sup>1,2</sup>, Wilson Truccolo <sup>1</sup>

<sup>1</sup> Brown University – United States

<sup>2</sup> Centre de Recerca Matemàtica – Spain

In recent years, nonlinear Hawkes processes implemented as point process generalized linear models (PPGLMs) have been proven to be a useful tool for analyzing microelectrode array recordings of neuronal ensemble spiking activity. They are also related to classical models of neuronal dynamics, but unlike ODE-based neuron models, PPGLMs can be fitted directly to the spike-time data, using standard optimization tools.

An ongoing research problem is to further the understanding of the stability and neural dynamics in both univariate and multivariate (network) data-driven PPGLMs. Our previous work in Gerhard et al. (2017, PLoS Comp Biol), showed that often simulation of these models can be unstable, leading to non-physiologically high firing rates. Thus to make nonlinear Hawkes PPGLMs useful for long-term prediction of neuronal activity and simulation studies, it is important to understand which model features cause the firing rates to become excessively large (“runaway excitation” phenomena). Despite several existing approaches based on statistical physics-inspired methods developed for the analysis of similar systems, their performance when dealing with data-driven PPGLMs, in particular multivariate models, has not been assessed yet.

Here, we compare the accuracy of several theoretical approaches for predicting the occurrence of runaway excitation in multivariate Hawkes processes. The approaches are based on the following theoretical approaches: mean field approximation, 1-loop fluctuation expansion based on stochastic path integral formulations, quasi-renewal approximation (Gerhard et al., 2017) and the regular spiking limit test (Gerhard et al., 2017). These approaches are quite different conceptually, having been introduced in different settings and having limitations in different aspects.

Based on simulation studies, we identify model features that make some approaches work much better than others. In addition, we show that, in some cases, the different approaches can complement each other. Furthermore, we demonstrate how these theoretical approaches work when applied to multivariate PPGLMs fitted to nonhuman primate cortex data. Finally, we also describe an algorithm for efficient simulation of arbitrary nonlinear Hawkes process networks.

---

\*Speaker



# P17. A Dynamic Model for Decoding Direction and Orientation in Macaque's Primary Visual Cortex

Wahiba Taouali \* <sup>1</sup>, Giacomo Benvenuti <sup>2</sup>, Frédéric Chavane <sup>3</sup>, Laurent Perrinet \*

1

<sup>1</sup> Institut de Neurosciences de la Timone – CNRS : UMR7289, Aix Marseille University. – France

<sup>2</sup> University of Texas at Austin – United States

<sup>3</sup> Institut de Neurosciences de la Timone – CNRS : UMR7289, Aix Marseille University. – France

When objects are in motion, the local orientation of their contours and the direction of motion are two essential components of visual information which are processed in parallel in the early visual areas. Generally, to probe a neuron's response property to moving stimuli, bars or gratings are drifted across neuron's receptive field at various angles. The resulting tuning curve will reflect the "confound" selectivity to both the orientation and direction of motion orthogonal to the orientation. Focusing on the primary visual cortex of the macaque monkey (V1), we challenged different models for the joint representation of orientation and direction within the neural activity. Precisely, we considered the response of V1 neurons to an oriented moving bar to investigate whether, and how, the information about the bar's orientation and direction could be encoded dynamically at the population activity level. For that purpose, we used a decoding approach based on a space-time receptive field model that encodes jointly orientation and direction. Then, using this model and a maximum likelihood paradigm, we inferred the most likely representation for a given network activity [1, 2]. We tested this model on surrogate data and on extracellular recordings in area V1 of awake macaque monkeys in response to oriented bars moving in 12 different directions. Using a cross-validation method we could robustly decode both the orientation and the direction of the bar within the classical receptive field (cRF). Furthermore, this decoding approach shows different properties: First, information about the orientation and direction of the bar is emerging before entering the cRF. Second, when testing different orientations with the same direction, our approach unravels that we can "unconfound" the information about direction and orientation by decoding them independently. Finally, our results demonstrate that the orientation and the direction of motion of an ambiguous moving bar can be progressively decoded in V1. This is a signature of a dynamic solution to the aperture problem in area V1, similarly to what was already found in area MT [3].

1-M. Jazayeri and J.A. Movshon. Optimal representation of sensory information by neural populations. *Nature Neuroscience*, 9(5):690–696, 2006.

2-W. Taouali, G. Benvenuti, P. Wallisch, F. Chavane, L. Perrinet. Testing the Odds of Inherent versus Observed Over-dispersion in Neural Spike Counts. *Journal of Neurophysiology*,

---

\*Speaker

2015.

3-C. Pack, R. Born. Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. *Nature*, 409(6823), 1040–1042. 2001.

This work was supported by "BrainScaleS" (EC FP7-269921) and "FACETS-ITN" (FP7-237955).

# P18. Coordination of inhibitory Golgi cell population activity in the cerebellar cortex

Harsha Gurnani \* <sup>1</sup>, R Angus Silver <sup>1</sup>

<sup>1</sup> University College, London – United Kingdom

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 shape 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). Additionally, GoC pairwise activity correlations remained high across distances of at least 200  $\mu\text{m}$ . Despite this slow-timescale (hundreds of milliseconds) GoC network coherence, there was faster heterogeneous modulation superimposed on the envelop of activation. We are quantifying the spatial organisation of these functional differences.

This suggests a mixed role of shared and tuned inhibition within the GCL. Global inhibition may perform adaptive gain control and support pattern separation in the GC layer by increasing decorrelation and sparseness of GC activity [Billings 2014, Cayco-Gajic 2017], while functional inhibitory subnetworks may shape temporal and plasticity profile of GCs.

References:

- [1] Maex and De Schutter 1998, *J Neurophysiol.* 1998 Nov; 80(5):2521-37
- [2] Vervaeke et al 2010, *Neuron.* 2010 Aug 12; 67(3): 435-51
- [3] Nadella et al 2016, *Nat Methods.* 2016 Dec; 13(12):1001-1004
- [4] Billings et al 2014, *Neuron.* 2014 Aug 20; 83(4):960-74
- [5] Cayco-Gajic et al 2017, *Nature Comm.* 2017 Oct 24; 8 (1), 1116.

---

\*Speaker

# P19. A Unified Theory of Early Visual Representations from Retina to Cortex through Anatomically Constrained Deep CNNs

Stéphane Deny \* <sup>1</sup>

<sup>1</sup> Stanford University – Stanford University – France

The visual system is hierarchically organized to process visual information in successive stages. 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, typical RFs are sharply tuned to a precise orientation. There is currently no unified 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 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 cortical networks, visual representations at the retinal output emerge as nonlinear and lossy feature detectors, whereas they emerge as linear and faithful encoders of the visual scene for more complex cortices. This result predicts that the retinas of small vertebrates 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.

---

\*Speaker

# P20. Rate-space attractors and low dimensional dynamics interact with spike-synchrony statistics in spiking neural networks

Daniel N. Scott \*<sup>1,2</sup>, Michael J. Frank<sup>1,2</sup>

<sup>1</sup> Department of Cognitive, Linguistic, and Psychological Sciences, Brown University – United States

<sup>2</sup> Carney Institute for Brain Science, Brown University – United States

The dimension of a neural network's rate space constrains the set of states it can access and thereby its behavior. As such, the investigation of rate dynamics and their dimensionality is a major area of research spanning theoretical neuroscience, the statistical analysis of neural data, and the development of brain-computer interfaces. An alternative description of spiking networks treats time-binned spiking activity as a sequence of binary vectors and provides a fundamentally different constraint on a network's states. Such a description also maintains precise spike-timing information and has strong links with statistical physics. These approaches are likely to be complementary, but the interaction between rate space dynamics and spike distributions has received little attention. Our research investigates this relationship. Specifically, we analyze spike co-occurrence probabilities and rate-space dynamics via mathematical derivation and neural network simulation. We consider spiking neural networks with varying spike-generation assumptions and establish quantitative relationships between rate-space dynamics and so-called spike-word distribution statistics. We find that, even under the assumption of independent Poisson-like spiking, attractors in rate-space dynamics can modify spike-synchrony rates. These modifications are the result of a global integration over the eigenvalues of the linear dynamics local to attracting subspaces. Several assumptions are implicit in the Poisson model, including a separation of time-scales between spikes and rates, the input of uncorrelated neural membrane noise, and independence of spiking activity across neurons. We relax all of these assumptions and analyze the resulting models as well. Effects on the spike-word distribution entropy, a key quantity in information theoretic and statistical physics settings are also discussed. Overall, our results elaborate how signatures of rate-dynamics are found in spike-word distributions and vice versa, along with some of their consequences.

---

\*Speaker

# P21. Dynamics of random networks with correlated low-rank perturbations.

Friedrich Schuessler \* <sup>1</sup>, Alexis Dubreuil <sup>2</sup>, Francesca Mastrogiuseppe <sup>3</sup>,  
Srdjan Ostojic <sup>2</sup>, Omri Barak <sup>4</sup>

<sup>1</sup> Department of Medicine, Technion, Haifa – Israel

<sup>2</sup> Group for Neural Theory, Laboratoire de Neurosciences Cognitives – Ecole Normale Supérieure de Paris - ENS Paris – France

<sup>3</sup> Gatsby Computational Neuroscience Unit, UCL – United Kingdom

<sup>4</sup> Department of Medicine, Technion, Haifa – Israel

A given neural network in the brain is involved in many different tasks. This implies that, when considering a specific task, the network's connectivity contains a component which is related to that task and another component that is, in a sense, random. Understanding the interplay between the structured and random components, and their effect on network dynamics and functionality is an important open question. Connectivity structures that were hand-designed to solve specific tasks are often low-rank. Similarly, trained networks have shown to involve low-rank structures. Finally, experimental recordings showed both low-dimensional activity as well as large heterogeneity – consistent both with low-rank structures as well as random connectivity. Previous work addressed the co-existence of random and structured connectivity, but without correlation between them, limiting the number of possible solutions. We characterize the effect of correlations in linear and nonlinear networks. The outliers in the eigenvalue spectra of random matrices and additional low-rank structures are computed. For nonlinear networks, we characterize emerging fixed points. Finally, for Gaussian structure vectors, we uncover a degeneracy in the type of solutions and derive closed form solutions for the fixed point stability.

---

\*Speaker

# P22. Coherence states and signal transfer of communicating gamma oscillatory neural networks

Gregory Dumont \* <sup>1</sup>, Boris Gutkin

<sup>1</sup> Gregory Dumont – Ecole Normale Supérieure de Paris - ENS Paris – France

Macroscopic oscillations of different brain regions show multiple phase relationships that are persistent across time [6]. Such phase locking is believed to be implicated in a number of cognitive functions and is key to the so-called Communication Through Coherence theory for neural information transfer [8]. Multiple cellular level mechanisms influence the network dynamic and structure the macroscopic firing patterns. Key question is to identify the synaptic properties that permit such motifs to arise and how the different coherence states determine the communication between circuits.

We use a semi-analytic approach to investigate the emergence of phase locking within two bidirectionally delayed-coupled spiking circuits with emergent gamma oscillations. Internally the circuits consist of excitatory and inhibitory quadratic integrate-and-fire neurons coupled synaptically in an all-to-all fashion [7]. The circuits can show global pyramidal-interneuron or interneuron gamma rhythms. Using a mean-field approach and an exact reduction method [3,9], we break down each gamma network into a low dimensional nonlinear system. We then derive the macroscopic phase resetting-curves [2,4] that determine how the phase of the global oscillation responds to incoming perturbations.

We then study the emergence of macroscopic coherence states of two weakly synaptically-coupled gamma-networks [1]. We derive a phase equation that links the synaptic mechanisms to the coherence state of the system. We show that the delay is a necessary condition for symmetry breaking, i.e. a non-symmetric phase lag between the macroscopic oscillations. We find that a whole host of phase-locking relationships exist, depending on the coupling strength and delay, potentially giving an explanation to the experimentally observations [8]. Our analysis further allows us to understand how signal transfer between the gamma circuits may depend on the nature of their mutual coherence states [5].

## Funding

GD: CNRS, INSERM, ANR-ERMUNDY and IDEX ANR-10-IDEX-0001-02 PSL\*. BG: supported by Russian Science Foundation grant (Contract No. 17- 11-01273).

## References

- [1] Ashwin, P., Coombes, S., Nicks, R. *The Journal of Mathematical Neuroscience* 6(1), 2 (2016).
- [2] Brown, E., Moehlis, J., Holmes, P.: *Neural Computation* 16(4), 673-715 (2004).
- [3] Deco, G., Jirsa, V.K., Robinson, P.A., Breakspear, M., Friston, K. *PLoS Comput Biol* 4(8), 1-35 (2008).
- [4] Dumont, G., Ermentrout, G.B., Gutkin, B. *Phys. Rev. E* 96, 042,311 (2017).

---

\*Speaker

- [5] Dumont, G., G.B., Gutkin, B. Accepted in Plos Computational Biology.
- [6] Fries, P.: Annual review of neuroscience 32, 209–224 (2009)
- [7] Izhikevich, E.M. The MIT Press (2007)
- [8] Maris, E., Fries, P., van Ede, F. Trends in Neurosciences 39(2), 86-99..
- [9] Montbrió, E., Pazo, D., Roxin, A. Phys. Rev. X 5, 021,028 (2015).



# P22. Computational analysis of cholinergic mechanisms for disinhibition of CA1 pyramidal cells and induction of hippocampal plasticity

Inês Guerreiro \* <sup>1</sup>, Zhenglin Gu <sup>2</sup>, Jerrel Yakel <sup>2</sup>, Boris Gutkin <sup>1</sup>

<sup>1</sup> Laboratoire de Neurosciences cognitives – École normale supérieure - Paris : U960 – France

<sup>2</sup> Neurobiology Laboratory, NIEHS – United States

Studies of induction of hippocampal plasticity have shown that blockade of GABA inhibition can greatly facilitate the induction of LTP in excitatory synapses (Wingstrom and Gustafsson, 1983).

It was shown experimentally that repeated inhibition of hippocampal CA1 somatostatin-positive interneurons can induce lasting potentiation of Schaffer collateral (SC) to CA1 EPSCs, suggesting that repeated dendritic disinhibition of CA1 pyramidal cells plays a role in the induction of synaptic plasticity. It was also shown experimentally that repeated cholinergic activation enhanced the SC-evoked EPSCs through  $\alpha 7$ -containing nicotinic acetylcholine receptors ( $\alpha 7$  nAChRs) expressed in oriens lacunosum-moleculare (OLM $\alpha 2$ ) interneurons.

We used a biophysically-realistic computational model to examine mechanistically how inhibitory inputs to hippocampal pyramidal neurons can modulate the plasticity of the SC-CA1 excitatory synapses. We found that locally-reduced GABA release (disinhibition) paired with SC stimulation could lead to increased NMDAR activation and intracellular calcium concentration sufficient to upregulate AMPAR permeability and potentiate the excitatory synapse. Repeated disinhibition of the excitatory synapses could lead to a larger increase of the AMPAR permeability. This results in the potentiation of the SC-CA1 excitatory synapse, which can be maintained when disinhibition period is over through repeated stimulation of the SC that keeps a balance between the down and upregulation of the AMPARs.

We then used our model to show how repeated cholinergic activation of  $\alpha 7$  nAChR in stratum oriens OLM $\alpha 2$  interneurons paired with SC stimulation can induce synaptic plasticity at the SC-CA1 excitatory synapses. Activation of pre-synaptic  $\alpha 7$  nAChRs in OLM cells activates these interneurons which, in turn, inhibit fast-spiking stratum radiatum interneurons that provide feed-forward inhibition onto pyramidal neurons after SC excitation, and thus disinhibiting the CA1 pyramidal neurons and inducing synaptic plasticity.

---

\*Speaker

# P23. Controlling Burst Activity Allows for a Multiplexed Neural Code in Cortical Microcircuits

Filip Verduyck \* <sup>1</sup>, Paola Suarez <sup>1</sup>, Henning Sprekeler <sup>1</sup>, Richard Naud <sup>2</sup>

<sup>1</sup> Technical University Berlin – Germany

<sup>2</sup> University of Ottawa – Canada

The presence of special mechanisms for generating bursts in pyramidal cells (PCs) suggests that bursts are likely to be an important temporal feature of neural signals. Bursts seem correlated with sensory processing and perception, induce synaptic plasticity, and have been proposed as a cellular mechanism to combine sensory and internal information. In layer 5 PCs, bursts occur at a low, but consistent rate, and are thought to originate from active dendritic processes. Because burst activity relies on dendritic threshold mechanisms, it appears likely that low burst activity requires an intricate homeostatic control. We hypothesized that this control is mediated by inhibitory plasticity of connections from Martinotti cells, which are known to control dendritic activity. To investigate this hypothesis, we studied a computational network model comprising layer 5 pyramidal cells with a somatic and dendritic compartment that was fitted to *in vitro* data, as well as different subclasses of interneurons. Our results show that a simple Hebbian plasticity rule on inhibitory synapses leads to robust and self-organized control of dendritic and burst activity. The dendritic learning rule we propose is based on a homeostatic rule that was previously proposed to control somatic spiking activity and therefore inherits properties such as a balance of excitation and inhibition. We demonstrate that this E/I balance is necessary for realistic burst firing patterns in biologically inspired cortical microcircuits with inhibitory neurons and recurrent connections. Furthermore, we show that the self-organized control of somatic and dendritic activity in pyramidal cells enables a multiplexed burst code suggested recently, by alleviating the need to tune input or noise levels. Finally, we show in simulations that the self-organising properties of inhibitory plasticity rules can be used to multiplex sensory and decision-related signals in decision-making networks, allowing us to decode behavioral decisions from burst activity in populations of sensory neurons.

---

\*Speaker

# P24. On the spectra and connectivity properties of transiently amplifying random networks

Georgia Christodoulou \* <sup>1</sup>, Friedemann Zenke <sup>1,2</sup>, Everton J. Agnes <sup>1</sup>,  
Tim P. Vogels <sup>1</sup>

<sup>1</sup> Department of Physiology, Anatomy and Genetics, University of Oxford – United Kingdom

<sup>2</sup> Friedrich Miescher Institute for Biomedical Research (FMI) – Switzerland

Cortical networks can exhibit a dynamical behaviour that is characterized by strong, multiphasic responses to specific inputs, and stable baseline activity. Such transiently amplifying circuits are thought to be prevalent in many neural circuits, but the link between their structure and function is not well understood.

Transient amplification in network models satisfying Dale’s law requires strong excitatory weights, which are counterbalanced by precisely tuned inhibition to retain stability. Such Stability Optimised Circuits (SOCs) are typically created through numerical optimization, and we find that the resulting SOC’s possess eigenspectra with a characteristic shape: (i) the imaginary part of the spectrum is spread along the imaginary axis, (ii) the real parts of the bulk of the eigenvalues are constrained to a narrow regime along the real axis, (iii) with the exception of one large negative real part.

Here, we identify the reason for these spectral properties, and study how variations affect the spectra and dynamics of the resulting SOC’s. Finally, we explore ways to design strongly amplifying circuits that obey Dale’s law from first principles, without the need for numerical optimization algorithm.

First, we examine why the imaginary part of spectrum needs to be spread along the imaginary axis. It is known that the spectral radius depends on the strong connections, which have been empirically linked to transient amplification. Here we formalise this fact mathematically and explore the role of the imaginary part further with numerical simulations. Adding to the known property that the upper bound of the real part is necessary for stability, we prove that the lower bound depends on the existence of auto-loops in the connectivity matrix. We show that the spectrum is drastically altered due to this constraint. Finally, we demonstrate how the position of the negative outlier depends on the extent of how much the dynamics are dominated by inhibition. We track the position of this outlier as a function of the network size and of the ratio of the inhibitory and excitatory strengths and characterize how it constrains the distribution of the remaining eigenvalues.

Finally, to understand how stability and amplification translate into properties of the connectivity matrix we quantify the effect of these two requirements on the architectural characteristics

---

\*Speaker

of the resulting networks, i.e, their weight distributions, in and out degrees, loops, etc., and further analyse this system within the context of a reduced mean field model. Our goal is to exploit the above findings to construct stable, amplifying networks from scratch without the need for numerical optimisation, and to shed light on putative learning rules which may serve this purpose in neurobiology.

# P25. Functional network models from algorithmically designed synaptic plasticity rules

Basile Confavreux \* <sup>1</sup>, Friedemann Zenke <sup>1,2</sup>, Everton J. Agnes <sup>1</sup>, Tim P. Vogels <sup>1</sup>

<sup>1</sup> Centre for Neural Circuits and Behaviour, University of Oxford – United Kingdom

<sup>2</sup> Friedrich Miescher Institute for Biomedical Research – Switzerland

Synaptic plasticity is the basis of learning. In models, synaptic plasticity is usually expressed in (activity dependent) rules that describe the temporal evolution of connections strengths between neurons. A common approach to find such rules has been to identify the minimal (simplest) expression that captures a given dataset from synaptic plasticity experiments. However, the rules produced by this approach are typically under-constrained and only provide an incomplete description of the various mechanisms at play. Indeed, the faithful implementation of any of these rules in neural network models rarely produces functional or even just stable networks. Here, we sought to find more comprehensive, parsimonious and biologically plausible plasticity rules by way of machine learning.

To that end, we developed an algorithmic meta-learning approach to search a set of all possible learning rules and identify effective candidates. We begin with a general and flexible plasticity rule that is expressed by a set of polynomials that determine the impact of synaptic strength, pre- and post synaptic activity, neighbouring input currents and other factors. These candidate rules are initially populated with random parameters, and evaluated in how effectively they transform a network with an arbitrary initial connectivity into a network with a specific, pre-set function after learning.

To test the validity of our approach, we initially aim to (re-) discover previously known plasticity rules, with specific function for the postsynaptic neuron. On such function is, e.g., to find the direction of largest variability in a cloud of points, or, in other words, to find the first principal vector of a dataset. It can be solved in a two-layer neural network by a simple Hebbian plasticity rule named Oja’s rule. By comparing the connectivity of a network trained with Oja’s rule to the connectivity of a network trained with the current candidate rule, we define a loss function which we minimize. To solve this optimization problem, we explored various optimizers adapted from the machine learning literature. Specifically, we show and compare several successful approaches which robustly find plasticity rules similar to Oja’s rule and discuss the details of the optimization landscape.

In summary, we present an optimization technique that selects plasticity rules based on their ability to generate a desired network architecture and function. It serves as a proof of principle for future work in which the optimizer will be tested on more realistic network models and

---

\*Speaker

functions.

# P26. Complementary Excitatory and Inhibitory Plasticity Rules produce Sparse Coding through EI Balance in a Spiking Neural Network Model

Nasir Ahmad \* <sup>1</sup>, Tim Vogels <sup>2</sup>

<sup>1</sup> Department of Experimental Psychology, University of Oxford – United Kingdom

<sup>2</sup> Centre for Neural Circuits and Behaviour, University of Oxford (CNCB) – United Kingdom

Reliable unsupervised learning in spiking neural network models is non-trivial, in particular when aiming for biological plausibility and using ecological stimulus sets. Here we combine an existing inhibitory synaptic plasticity rule [1] with a novel formulation of an excitatory synaptic plasticity rule (based upon van Rossum et al. 2000, [2]) and we produce networks which converge to a sparse distributed encoding of stimulus features. This learning rule arrangement is applied in a 1-layer feedforward network architecture with Poisson spiking input neurons, leaky integrate-and-fire output neurons, lateral inhibitory synaptic connections, and all spiking dynamics. The success of this learning framework is shown in a model of V1 Simple Cell (Gabor-like) receptive field development after presentation of patches from a natural image dataset, reproducing a classic feature of sparse coding models [3]. An analysis of the effective dynamics of the combined learning rules show that changes in synaptic weight depend upon the post-synaptic neuron firing rate in a sliding-threshold manner. The EI balancing nature of our choice of inhibitory synaptic plasticity rule is key to the production of this sliding-threshold behaviour. We also find that the lateral inter-inhibitory neuron connectivity (which provides competition between inhibitory neurons) can produce oscillating receptive fields if randomly initialised. Instead, by initialising the network with bidirectional inter-inhibitory neuron synaptic connections, the formation of receptive fields is stabilised. Our approach ultimately provides a circuit level description of a computationally simple and robust learning process in spiking neural network models with a predictable sparseness of neural responses.

[1] Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*. 2011 Dec 16;334(6062):1569-73.

[2] Van Rossum MC, Bi GQ, Turrigiano GG. Stable Hebbian learning from spike timing-dependent plasticity. *Journal of neuroscience*. 2000 Dec 1;20(23):8812-21.

[3] Olshausen BA, Field DJ. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*. 1996 Jun;381(6583):607.

---

\*Speaker

# P27. Learning functional properties of visual cortex from ultrafast fluorescence microscopy recordings

Sebastien Wolf <sup>\*</sup> <sup>1</sup>, Walther Ackeman <sup>1</sup>, Laurent Bourdieu <sup>1</sup>, Rémi Monasson <sup>2</sup>

<sup>1</sup> Ecole Normale Supérieure, Paris – Ecole Normale Supérieure de Paris - ENS Paris – France

<sup>2</sup> École normale supérieure – Ecole Normale Supérieure de Paris - ENS Paris – France

Here, we take advantage of a new optical technique to demonstrate the efficiency of restricted Boltzmann machines to model and understand cortical activity in V1 of mouse.

The data sets

Ultrafast fluorescence microscopy using calcium dyes combined with spike inference algorithms demonstrated recently [1] that it is now possible to record simultaneous spiking activity of large neuronal populations at high framerates (several hundreds of Hz). This 2-photon microscopy technique based on acousto-optic deflectors (AOD) allowed us to record simultaneously at high framerates (400Hz) the neuronal activity within several layers in the visual cortex (V1) of awake mice under different behavioral contexts.

Modelling

In order to understand and model the recorded activity of V1, we followed model-based approach. We used a class of models called the Boltzmann and Restricted Boltzmann machines (BM and RBM). They are graphical models that learn a probability distribution and, jointly in the case of RBM a representation of the data sets. We prove the efficiency of these machine to fit cortical data sets in the case of our calcium recordings. We show that RBM outperforms BM in terms of likelihood and computational speed capturing very precisely the high-order correlations of the cortical columns in our data sets.

The representation: the case of RBM

Then, we tried to address the question of representation in machine learning in the particular case of RBM applied to neural data. How can we interpret the networks fitted on data sets by machine learning approaches? Could these machines help us to extract the principal properties of the recorded brain regions?

RBM are two layer bipartite neural nets, simple to manage analytically and in terms of algorithm. Beyond its simplicity, it has been shown that these machines have a compositional phase of their parameters [2] i.e. a phase in which they are easily interpretable. Indeed, in

---

<sup>\*</sup>Speaker



the case of protein data sequence analysis, RBM have proved to be very efficient in extracting structural and functional properties of proteins directly from its sequences [3].

Here, fitting RBM to our large V1 data sets of mouse under different behavioral paradigms, we show that these machines are able to extract the principal properties of the network such as (i) the dependency of responses in V1 to the angle of the bar stimulus, (ii) the layered organization of the correlation (layer II, III and V), (iii) the presence of purely spontaneous correlations in V1, (iv) the delayed responses to stimulus of sub-population of neurons in V1 and even (v) a propagation process occurring between layers. We compare these results to traditional PCA and ICA analysis and demonstrate the power of RBM in jointly fitting and building a comprehensive representation of cortical data.

[1] Akemann et al, 2019 (submitted)

[2] Tubiana and Monasson, Phys. Rev. Let., 2017

[3] Tubiana et al., eLife, 2019