Neural Coding, Computation and Dynamics 2017

17-20 September 2017, Capbreton, France

Program

Sunday, September 17, 2017

5:00 pm - 7:00 pm	Registration and Welcome
7:00 pm - 7:05 pm	Opening remarks - Organizers
7:05 pm - 7:50 pm	Directing variability for learning in birdsong - Adrienne Fairhall (Seattle)
8:00 pm - 10:00 pm	Dinner

Monday, September 18, 2017

9:30 am - 10:15 am	Flexible statistical inference for mechanistic models of neural dynamics - Jakob Macke (Bonn)
10:15 am - 10:45 am	Dimensionality Reduction of Neural Dynamics Within and Across Trials By Tensor Decomposition- Alex Williams*, Tony Kim, Fori Wang, Saurabh Vyas, Krishna Shenoy, Mark Schnitzer, Tamara Kolda, Surya Ganguli
10:45 am - 11:15 am	Coffee break
11:15 am - 12:00 pm	Efficient coding strategies in populations of sensory neurons - Julijana Gjorgjieva (Frankfurt)
12:00 pm - 12:30 pm	Multiplexed computations in retinal ganglion cells of a single type - Olivier Marre*, Stephane Deny, Ulisse Ferrari, Emilie Mace, Pierre Yger, Romain Caplette, Serge Picaud, Gasper Tkacik
12:30 pm - 2:30 pm	Lunch
2:30 pm - 3:15 pm	Brain-computer interfaces for basic science - Byron Yu (Pittsburgh)
3:15 pm - 3:45 pm	Modeling within and across area neuronal variability in the visual system - Chengcheng Huang*, Douglas Ruff, Ryan Pyle, Robert Rosenbaum, Marlene Cohen, Brent Doiron
3:45 pm - 4:15 pm	Coffee break
4:15 pm - 5:00 pm	Dissecting striatal circuitry for reward seeking behavior. - Ilana Witten (Princeton)
5:00 pm - 5:30 pm	Noise-robust “codewords” of the retinal population code: applicability, geometry, and correspondence to neuronal communities - Adrianna Loback*, Jason Prentice, Mark Ioffe, Michael Berry
5:30 pm - 8:00 pm	Poster Session I: P1-P18
8:00 pm - 10:00 pm	Dinner

Tuesday, September 19, 2017

9:30 am - 10:15 am	Motor cortex embeds muscle-like commands in an untangled population response - Mark Churchland (Columbia)
10:15 am - 10:45 am	Task engagement induces shift from sensory to behavioral representations in primary auditory cortex - Sophie Bagur*, Martin Averseng, Diego Elgueda, Stephen David, Jonathan Fritz, Pingbo Yin, Shihab Shamma, Yves Boubenec, Srdjan Ostojic
10:45 am - 11:15 am	Coffee break
11:15 am - 12:00 pm	Discovering dynamic computations in the brain from large-scale neural recordings - Tatiana Engel (Cold Spring Harbor)
12:00 pm - 12:30 pm	Learning probabilistic representations with randomly connected neural circuits - Ori Maoz*, Gasper Tkacik, Roozbeh Kiani, Elad Schneidman
12:30 pm - 2:30 pm	Lunch
2:30 pm - 3:15 pm	Changing dynamical roles of different cortical cell-types with learning - Maneesh Sahani (UCL)
3:15 pm - 3:45 pm	High cellular and columnar variability underlies the absence of early orientation selectivity - David Whitney, Gordon Smith, Bettina Hein, David Fitzpatrick, Matthias Kaschube*
3:45 pm - 4:15 pm	Coffee break
4:15 pm - 4:45 pm	A motor readout of visual perception: Deciphering cuttlefish camouflage at single-chromatophore resolution - Sam Reiter*, Philipp Hülsdunk, Theodosia Woo, Marcel Lauterbach, Jessica Eberle, Daniel Noll, Matthias Kaschube, Gilles Laurent
4:45 pm - 5:15 pm	A sensorimotor hub driving phototaxis in zebrafis - Alexis Dubreuil*,Sébastien Wolf, Georges Debrégeas, Rémi Monasson
5:30 pm - 8:00 pm	Poster session II: P19-P36
8:00 pm - 10:00 pm	Dinner

Wednesday, September 20, 2017

9:30 am - 10:15 am	Concentration invariant odor coding - Dmitry Rinberg (NYU)
10:15 am - 10:45 am	A model circuit describing spatial integration and its modulation by running of four neuronal classes recorded in V1. - Mario Dipoppa*, Matteo Carandini, Kenneth Harris
10:45 am - 11:15 am	Coffee break
11:15 am - 11:45 am	Distinct regulation of history-dependent responses by two cortical interneuron populations - Elizabeth Phillips, Christoph Schreiner, Andrea Hasenstaub
11:45 am - 12:30 pm	Stimulus quality vs stimulus intensity in sensory discrimination: Behavior and some neural correlates - Alfonso Renart (Lisbon)
12:30 pm - 2:30 pm	Lunch

Table of contents

Dimensionality Reduction of Neural Dynamics Within and Across Trials By Tensor Decomposition

Alex Williams * [1], Tony Kim [1], Fori Wang [1], Saurabh Vyas [1], Krishna Shenoy [1], Mark Schnitzer [1], Tamara Kolda [2], Surya Ganguli [1]

[1] Stanford University – Stanford, CA 94305, United States
[2] Sandia National Laboratories [Livermore] – Livermore, CA 94551, United States

Decision-making, sensation, and motor behaviors occur within fractions of seconds, while memories and learned behaviors can require many days or months to mature. Recent advances enable long-term experiments that monitor all of these timescales across hundreds of neurons and behavioral trials. However, classic and commonly used dimensionality reduction techniques are ill-equipped to summarize data across multiple timescales. We represent multi-trial neural data as a tensor (i.e., a higher-order array), and apply canonical polyadic (CP) tensor decomposition to identify low-dimensional factors that separately summarize short-term and long-term changes in neural population dynamics. We show that CP tensor decomposition formally maps onto a linear network model with low-dimensional gain modulation, and thus draw a novel connection between established tensor decomposition techniques and influential theories of sensory processing and cortical circuit function. In synthetic data generated from model networks, CP decomposition precisely identifies network inputs that vary across trials, whereas classical methods (PCA and ICA) fail to recover these ground-truth signals.

In experimental datasets collected from different species, brain regions, and behavioral tasks, CP decomposition uncovers sub-populations of neurons with low-dimensional firing rate dynamics that vary across-trials according to experimental conditions, rewards, behavioral actions, and motivational states. In mice, we applied this approach to _˜500 GCaMP6m-expressing cells from the medial prefrontal cortex (mPFC) monitored via fluorescence microendoscopy during a four-armed maze task with changing reward contingencies. The low-dimensional factors revealed that mPFC neurons (a) cluster into functional cell types and exhibit limited levels of mixed selectivity for task variables, (b) preferentially respond to allocentric (place-based) over egocentric (turn-based) task variables, and (c) exhibit stable coding for task variables over multiple days even as the animal's navigational strategy changes.

In a second experiment, we examined multi-unit data collected from a Rhesus macaque making radial point-to-point reaches with a virtual cursor through a brain-machine interface (BMI). We used CP decomposition to summarize how neural dynamics responded to a perturbation of the BMI decoder by a visomotor rotation. We found that some low-dimensional factors were selectively activated after the BMI perturbation, while others were consistently present across all trials and showed little response to the visomotor rotation. Our preliminary findings did not identify factors that de-activated in response to the BMI perturbation, suggesting a significant constraint on models of motor learning.

*Speaker

Multiplexed computations in retinal ganglion cells of a single type

Olivier Marre * [1], Stephane Deny [2], Ulisse Ferrari [3], Emilie Mace [3], Pierre Yger [3], Romain Caplette [4], Serge Picaud [3], Gasper Tkacik [5]

[1] Institut de la Vision (IdV) – Institut de la Vision – France
[2] Institut de la Vision – Université Pierre et Marie Curie - Paris 6 – France
[3] Institut de la Vision – Université Pierre et Marie Curie [UPMC] - Paris VI – France
[4] Institut de la Vision (IdV) – Université Pierre et Marie Curie (UPMC) - Paris VI – France
[5] IST Austria – Austria

In the early visual system, cells of the same type perform the same computation in different places of the visual field.
How these cells code together a complex visual scene is unclear. A common assumption is that cells of the same type will extract a single stimulus feature to form a feature map, but this has rarely been observed directly. Using large-scale recordings in the rat retina, we show that a homogeneous population of fast OFF ganglion cells simultaneously encodes two radically different features of a visual scene. Cells close to a moving object code linearly for its position, while distant cells remain largely invariant to the object's position and, instead, respond non-linearly to changes in the object's speed. Cells switch between these two computations depending on the stimulus. We developed a quantitative model that accounts for this effect and identified a likely disinhibitory circuit that mediates it. Ganglion cells of a single type thus do not code for one, but two features simultaneously. This richer, flexible neural map might also be present in other sensory systems.

*Speaker

Modeling within and across area neuronal variability in the visual system

Chengcheng Huang * [1], Douglas Ruff [2], Ryan Pyle [3], Robert Rosenbaum [3], Marlene Cohen [2], Brent Doiron [1]

[1] Department of Mathematics, University of Pittsburgh – 130 University Place, PA 15260, United States
[2] Department of Neuroscience, University of Pittsburgh – A210 Langley Hall, PA 15260, United States
[3] Department of Applied and Computational Mathematics and Statistics, University of Notre Dame – 153 Hurley Hall, Notre Dame, IN 46556, United States

Shared variability among neurons (noise correlations) have been commonly observed in multiple cortical areas (Cohen and Kohn, 2011). Moreover, noise correlations are modulated by cognitive factors, such as overall arousal, task engagement and attention (Cohen and Maunsell, 2009; Doiron et al. 2016). While there is much discussion about the consequences of noise correlations on neuronal coding, there is a general lack of understanding of the circuit mechanisms that generate and modulate shared variability in the brain. Recently, simultaneous microelectrode array recordings from V1 and MT in behaving monkeys (Ruff and Cohen, 2016) suggest that attention not only decreases correlations within a cortical area (MT), but also increases correlations between cortical areas (V1 and MT). The differential modulation of between-areas and within-area noise correlations impose further constraints on circuit mechanisms for the generation and propagation of noise correlations. We develop a spiking neuron network with spatiotemporal dynamics that internally generates shared variability matching the low dimensional structure widely reported across cortex. The low dimensional variability requires slow and narrow inhibition, which are consistent with physiology. We further show that the internally generated variability results from macroscopic chaos in population rates, which correlates neurons from the same recurrent network while decoupling them from feedforward inputs. Attention is modeled as depolarizing the inhibitory neuron population, which reduces the internally generated shared variability and allows the network to better track input signal. Our model provides a much needed mechanism for how shared variability is both generated and modulated in recurrent cortical networks.

*Speaker

Noise-robust “codewords” of the retinal population code: applicability, geometry, and correspondence to neuronal communities

Adrianna Loback * [1], Jason Prentice [1], Mark Ioffe [2,1], Michael Berry [1,3]

[1] Princeton Neuroscience Institute – Princeton, NJ 08544, United States
[2] Physics Department, Princeton University – Princeton, NJ 08544, United States
[3] Molecular Biology Department, Princeton University – Princeton, NJ 08544, United States

Understanding how local brain circuits encode and transmit information using the combined activity of large neuronal populations is a fundamental challenge in neuroscience. With the advent of new experimental technologies that enable simultaneous recording from hundreds to thousands of neurons, it is becoming increasingly tractable to study this important question. Here we take a data-driven and generative statistical modeling approach to investigate an appealing new principle for neural population codes: that correlations among neurons organize neural activity patterns into a discrete set of clusters, which can each be viewed as a noise-robust population *codeword*. Previous studies assumed that these codewords corresponded geometrically with local peaks in the probability landscape of neural population responses. Here we analyze multiple datasets (obtained via multi-electrode array recordings) of the responses of up to 170 retinal ganglion cells and show that local probability peaks are absent under broad, non-repeated stimulus ensembles, which are characteristic of natural behavior. However, by applying a hidden Markov model approach, we find that neural activity in this regime still forms clusters – which we call “collective modes” – that enable a high degree of error correction. We develop a new numerical approach to characterize the geometry of these noise-robust collective modes, and find that they correspond to *ridges* in the probability landscape of population responses. Moreover, we found that these ridges are comprised of combinations of spiking and silence in the neural population such that all of the spiking neurons are members of the same neuronal *community*, a notion from network theory. Our results have several important implications. They provide support for the error-correcting population code paradigm being applicable to the retinal population code under naturalistic low-repeat stimulation, and moreover provide empirical constraints on the geometry of the codewords. Further, they motivate a new direction that we are currently exploring: whether there are biologically-plausible (i.e. online and using local learning rules) spiking network models that can implement an approximation of the Bayesian inference used here to extract the noise-robust collective modes.

*Speaker

Task engagement induces shift from sensory to behavioral representations in primary auditory cortex

Sophie Bagur * [1], Martin Averseng [2], Diego Elgueda [3], Stephen David [4], Jonathan Fritz [3], Pingbo Yin [3], Shihab Shamma [2,3], Yves Boubenec [2], Srdjan Ostojic [5]

[1] Brain Plasticity Unit, CNRS UMR 8249 – ESPCI ParisTech, PSL Research University, ESPCI ParisTech, PSL Research University – France

[2] Laboratoire des Systèmes Perceptifs, Équipe Audition – Ecole Normale Supérieure de Paris - ENS Paris – 29 rue d'Ulm, Paris, France, France

[3] Neural Systems Laboratory, Institute for Systems Research – University of Maryland in College Park, MD, USA, United States

[4] Laboratory of Brain, Hearing and Behavior,Oregon Health & Science University, OR, USA – United States

[5] Group for Neural Theory, Laboratoire de Neurosciences Cognitives – Ecole Normale Supérieure de Paris - ENS Paris – 29 rue d'ilm, Paris, France, France

The main functions of primary sensory cortical areas are classically considered to be the extraction and representation of stimulus features. In contrast, higher cortical sensory association areas are thought to be responsible for combining these sensory representations with internal motivations and learnt associations. These regions generate appropriate neural responses that are maintained until a motor command is executed. Within this framework, responses of the primary sensory areas during task performance are expected to carry less information about the behavioral meaning of the stimulus than higher sensory, association, motor and frontal cortices. Here we demonstrate instead that the neuronal population response in the early primary auditory cortex (A1) displays many aspects of responses generally associated with higher-level areas. A1 activity was recorded in awake ferrets while they were either passively listening or actively discriminating two periodic click trains of different rates in a Go/No-Go paradigm. By applying population-level dimensionality reduction techniques, we found that task-engagement induced a shift in the nature of the encoding from a balanced sensory-driven representation of the Go and No-Go stimuli, to an asymmetric behaviorally relevant representation of the two categories that specifically enhances the target stimuli. We demonstrate that upon engagement in the task, it is mainly a realignment of spontaneous activity patterns of the population that manifests this shift in encoding. Moreover, information about the target stimulus is maintained during the silent period immediately following its presentation and is correlated with the animal's behavioral response. We show that this behavioral representation of stimuli in A1 population activity bears strong similarities to responses in the frontal cortex, and appears rapidly following stimulus presentation. Analysis of neural responses recorded in various Go/No-Go tasks, with different sounds and reinforcement paradigms, reveals that this striking increased asymmetry of the population-level stimulus representation is a general property during task engagement.

*Speaker

These findings indicate that primary sensory cortices play a highly flexible role in processing of incoming stimuli and implement a crucial change in the structure of population activity in order to extract task-relevant information during behavior.

Learning probabilistic representations with randomly connected neural circuits

Ori Maoz * [1], Gasper Tkacik [2], Roozbeh Kiani [3], Elad Schneidman [1]

[1] Weizmann Institute of Science (WIS) – 234 Herzl Street, Rehovot 76100 Israel, Israel
[2] IST Austria – Am Campus 1, A-3400 Klosterneuburg, Austria
[3] Center for Neural Science [New York] – 4 Washington Place, Room 809 New York, NY 10003, United States

Despite extensive efforts to map neural connectivity in the mammalian brain, local connectivity within microcircuits often appears to be random which has been proposed to be a design principle of the vertebrate nervous system. Here we demonstrate that sparsely and randomly connected shallow feed-forward neural circuits can efficiently perform a complex generalization task: learning the probability distribution of their joint inputs such that they can predict the saliency, or surprise, of input activity patterns that were never previously encountered. Surprisingly, the models implemented by such circuits - which mathematically take the form of maximum entropy distributions of random nonlinear projections - exceed state-of-the-art methods for density estimation of neural population activity patterns and are particularly robust when learning from limited datasets. Learning the model involves modifying only the connection weights to an output neuron while leaving the remaining random connectivity intact, which allows to (1) scale the number of projections or neurons up or down as required and (2) have multiple output neurons share the randomly connected layer in order to learn multiple stimulus-conditioned models and perform classification at the cost of a single neuron per category. We furthermore present a novel biologically-plausible learning rule for efficiently learning an approximate solution to the model using weak local noise as a driving force. Finally, we show that by pruning weak synapses to the output neuron and forming new connections to unused part of the circuit, these representations not only become more efficient but also result in sparser-firing and decorrelated projections as a natural side effect.

*Speaker

High cellular and columnar variability underlies the absence of early orientation selectivity

David Whitney * [1], Gordon Smith [1], Bettina Hein [2], David Fitzpatrick [1], Matthias Kaschube *

2

[1] Max Planck Florida Institute for Neuroscience (MPFI) – 1 Max Planck Way, Jupiter, FL 33458, United States

[2] Frankfurt Institute for Advanced Studies (FIAS) – Frankfurt Institute for Advanced Studies Ruth-Moufang-Straße 1 60438 Frankfurt am Main, Germany

Selectivity for stimulus orientation is a fundamental property of primary visual cortex in primates and carnivores, where it is organized into a smoothly varying columnar map that emerges in an activity-dependent manner during early postnatal life. Despite extensive experimental and theoretical work, it remains unclear what factors limit the emergence of orientation selectivity, such as weak responsiveness to visual stimuli, high trial-to-trial variability, and/or an intermixed 'salt-and-pepper' organization of orientation preferences at the cellular level. To distinguish between these potential factors, we visualized population activity in the visual cortex of developing ferrets with longitudinal imaging of GCAMP6s at both cellular resolution with two-photon calcium imaging and columnar resolution with wide-field epifluorescence imaging. Prior to eye opening, we show that cellular and population responses evoked by single presentations of a grating stimulus surprisingly exhibit robust, modular patterns of network activity resembling activity patterns evoked by gratings in mature animals. However, the spatial location and pattern of domains activated by presentation of the identical stimulus orientation varies substantially across trials, a variability that accounts for the low orientation selectivity of individual neurons and the inability to visualize coherent maps of orientation preference. Yet variability in network activity patterns is not a general feature of the developing cortex, as the modular patterns of network activity evoked by uniform luminance steps are already selective at these ages. Furthermore, we show that trial-averaged activity patterns evoked by gratings show similarity to the mature orientation map as early as 1-2 days prior to eye opening. We conclude that the early disassociation between stimulus orientation and consistent patterns of modular network activity is a major factor underlying the absence of orientation selectivity in a developing cortical network already exhibiting highly modular functional organization.

*Speaker

A motor readout of visual perception: Deciphering cuttlefish camouflage at single-chromatophore resolution

Sam Reiter * [1], Philipp Hülsdunk [1,2], Theodosia Woo [1], Marcel Lauterbach [1], Jessica Eberle [1], Daniel Noll [1,2], Matthias Kaschube [2], Gilles Laurent [1]

[1] Max Planck Institute for Brain Research (MPIBR) – Max-von-Laue-Str. 4 60438 Frankfurt am Main Germany, Germany

[2] Frankfurt Institute for Advanced Studies (FIAS) – Frankfurt Institute for Advanced Studies Ruth-Moufang-Straße 1 60438 Frankfurt am Main, Germany

Coleoid cephalopods (cuttlefish, squid, and octopus) provide a unique opportunity to study visual perception in a large, complex nervous system that evolved independently of the vertebrate lineage. A large part of the cephalopod brain is involved in the coordination of skin patterning. In the cuttlefish *Sepia officinalis* this patterning subserves the most advanced camouflage in the animal kingdom. It is the direct expression of a neural control over the expansion and contraction of hundreds of thousands to millions of pigment-filled cells called chromatophores. How this visually driven coordination is accomplished remains a mystery. We are investigating this problem by monitoring skin patterning in awake, behaving cuttlefish, at single-chromatophore resolution, in response to artificial and controlled textures.
We developed a multi-level video-analysis pipeline to track tens of thousands of chromatophores simultaneously at 60 frames per second over weeks. Correlations of chromatophore states are used to infer elements of a hierarchical motor control strategy, starting with the identification of motor neurons driving small, overlapping collections of chromatophores. The coordination of these inferred motor neurons give rise to higher order pattern elements, which are in turn coordinated to produce macro-scale phenomena like bilateral pattern symmetry. These measurements are validated and extended by *in vitro* experiments where motor neuron innervation is assessed physiologically (through electrical stimulation) and anatomically (through super-resolution microscopy).

Non-affine image registration using small patches of skin as uniquely identifiable features allowed us to track identified chromatophores over months. We are currently using this technique to detect and study the continuous physical and functional integration of large numbers of newly developed chromatophores into the existing array. Our approach provides the first view of cephalopod skin patterning and development at the spatiotemporal scale of the nervous system. It also uses the unique features of an atypical model animal to shed light on the fundamental problems of hierarchical motor control and visual texture perception.
Funded by the Max Planck Society.

*Speaker

A sensorimotor hub driving phototaxis in zebrafish

Alexis Dubreuil * [1], Georges Debregeas [2], Sebastien Wolf [2], Rémi Monasson [1]

[1] Ecole normale supérieure (ENS Paris) – École normale supérieure - Paris – 45, rue d'Ulm75230 Paris cedex 05, France

[2] Laboratoire Jean Perrin - CNRS/UPMC (Laboratoire Jean Perrin - CNRS/UPMC) – Laboratoire Jean Perrin - CNRS – France

Animals continuously gather sensory cues in order to move towards favorable environments. Efficient goal-directed navigation requires sensory perception and motor command to be intertwined in a feedback loop, yet the neural substrate underlying this sensorimotor task in the vertebrate brain remains elusive. Here we combine virtual-reality behavioral assays, recordings of up to 50,000 neurons simultaneously via calcium imaging, optogenetic stimulation and circuit modeling to reveal the neural mechanisms through which zebrafish performs phototaxis, i.e. actively orients towards a light source. Key to this process is a self-oscillating network that acts as a pacemaker for ocular saccades and controls the orientation of successive swim-bouts. It further integrates visual stimuli in a state dependent manner, i.e. its response to visual inputs varies with the motor context, a mechanism that manifests itself in the phase-locked entrainment of the circuit by periodic stimuli. A network model is developed that mechanistically account for our observations and demonstrates how this sensorimotor processing eventually biases the animal trajectory towards bright regions.

*Speaker

A model circuit describing spatial integration and its modulation by running of four neuronal classes recorded in V1.

Mario Dipoppa * [1], Matteo Carandini [1], Kenneth Harris [1]

[1] University College London (UCL) – London WC1E 6BT, United Kingdom

A prominent feature of neurons in the primary visual cortex (V1) is selectivity for the size of visual stimuli, known as size tuning. In pyramidal (Pyr) neurons of mouse V1, size tuning is thought to be imposed by somatostatin-expressing (Sst) interneurons, whose responses integrate Pyr input over a wide cortical area (Adesnik et al, 2012). Size tuning, moreover, depends on locomotion (Ayaz et al., 2013), which can profoundly affect sensory responses in multiple neuronal classes of V1 (Polack et al., 2013; Fu et al., 2014; Pakan et al., 2016). We sought to understand the interplay between these phenomena by measuring the activity of Pyr, Sst, Parvalbumin (Pvalb), and Vasoactive Intestinal Polypeptide (Vip) cells, and describing their interactions with a circuit model.
We imaged the visual responses of superficial V1 neurons with the calcium indicator GCaMP6 using a 2-photon microscope, identifying Pyr (n = 5,050), Sst (n = 269), Vip (n = 554) and Pvalb (n = 494) neurons in transgenic mice. We measured size tuning both for cells whose receptive field was centered on the visual stimulus and for the remaining off-center cells. In contrast with a previous report (Adesnik et al, 2012), well-centered Sst neurons did exhibit size tuning, similarly to the other cell classes. Only off-center Sst neurons increased monotonically their responses with stimulus size. Additionally, locomotion increased only responses to small stimuli in Vip neurons and only responses to large stimuli in Sst neurons.

To capture these observations we devised a rate model based on the recurrent connections of these neurons classes and on feedforward excitation from thalamus. Existing recurrent connections were based on in vitro studies (Pfeffer et al., 2013) while thalamic input was directly estimated from the in vivo thalamic recordings in Erisken et al. (2014). In the model, the visual responses of each class are estimated as the sum of average responses from local Pyr and thalamic cells subtracted or divided by the average responses of functionally connected interneuron classes. The inputs were integrated over a cortical area of finite size with two-dimensional Gaussian functions.
The model provided accurate quantitative fits to the data of all cell types, but only if we incorporated the following constraints of the functional connections between V1 cell classes: (1) in addition to a broad spatial integration of inputs from Pyr neurons to Sst neurons as suggested by Adesnik et al. (2012), an even broader integration from Sst neurons to Pyr neurons; (2) thalamic visual input that impinges not only on Pyr and Pvalb neurons (Yang et al., 2013) but also on Sst cells (perhaps through Pyr neurons in deeper layers); (3) thalamic visual input that becomes stronger with locomotion perhaps due to stronger thalamocortical synapses modulated

*Speaker

by Acetylcholine (Gil et al., 1997). We conclude that a relatively simple rate model with reasonable assumptions can capture the apparently complex size- and locomotion-dependent responses of multiple classes of V1 neurons.

Distinct regulation of history-dependent responses by two cortical interneuron populations

Elizabeth Phillips [1], Christoph Schreiner [1], Andrea Hasenstaub * [1]

[1] University of California San Francisco (UCSF) – United States

Cortical responses to repeated stimuli are highly dynamic and rapidly adaptive. Such rapid changes are prominent in all sensory cortices, across which many aspects of circuitry are conserved. As an example, in the auditory cortex, preceding sounds can powerfully suppress responses to later, spectrally similar sounds – a phenomenon called forward suppression. Whether cortical inhibitory networks shape such suppression, or whether it is wholly regulated by common mechanisms such as synaptic depression or spike-frequency adaptation, is controversial. Here, we show that optogenetically suppressing somatostatin-positive interneurons reveals facilitation in neurons that are normally forward-suppressed. This is accompanied by a weakening of forward suppression, suggesting that these interneurons regulate the strength of forward interactions. In contrast, inactivating parvalbumin-positive interneurons does not change suppression strength, but does alter its frequency-dependence. By combining computational modeling with experimental controls, we further show that these features likely depend on cell type specific differences in synaptic dynamics, rather than on intrinsic adaptation or global effects on responsivenses. These results establish a role of cortical inhibition in forward suppression and link specific aspects of rapid sensory adaptation to genetically distinct interneuron types.

*Speaker

P1. A Design Principle for Population Neural Codes

Michael Berry * [1], Mark Ioffe [1], Adrianna Loback [1], Jason Prentice [1]

[1] Princeton Neuroscience Institute – Princeton, NJ 08544, United States

I will present an overview of recent work from my lab suggesting that a common principle for population codes might be that neural activity is automatically organized into a discrete set of ***clusters***. This result has emerged from the analysis of populations of retinal ganglion cells using two different frameworks. In all cases, we discretize spikes trains using small time bins (20 ms) and form binary words over N cells in the same time bin. We adopt the standpoint of an *activity model*, which attempts to describe the probability distribution over all possible activity patterns and understand the structure of this distribution. This approach is important, as it matches the task that the brain must solve – namely, to understand the structure of the visual world based on many observations of retinal spike trains and without any supervision or stimulus labels.

First, we show that a hidden Markov model (HMM) can accurately capture the statistics of population activity. In this HMM, each latent state is a well-separated cluster. When we repeat the same visual stimulus, we find that neural activity patterns are highly variable but that they map onto just one or a few clusters. Thus, clusters serve as population codewords that exhibit strong *error correction*. In addition, the receptive field of each cluster is qualitatively different than receptive fields of its constituent neurons. Thus, clusters form a different basis set of visual features. Finally, we show that there exist simple, biologically plausible decoding algorithms that can readout cluster identity. We will describe how these properties can be combined together to constitute a system of hierarchical pattern recognition.
Second, we use a maximum entropy model (MaxEnt) to show that the specific heat of the neural population exhibits a robust peak at a temperature slightly above that of the real neural population. We interpret this to mean that the real neural population is in a *low temperature state*. We present further evidence that this state resembles a spin glass in statistical physics, where each basin in the energy landscape is a cluster of neural activity. This finding is consistent with and complementary to the result that neural activity patterns are likely to generically exhibit Zipf's Law in the large N limit.

Because the properties of any MaxEnt model depend entirely on the constraining statistics (such as firing rates and pairwise correlations), the above results may also apply to neural populations elsewhere in the brain that have the same low-order statistics. To this end, we show that if we scale down all of observed pairwise correlations by a factor > 2, the retina remains in a low temperature state. This suggests that other population codes with pairwise correlations as strong as the retina may also exist in a low temperature limit, where neural activity is organized into clusters. We next hope to test these ideas by analyzing data from other brain areas, alone or in collaboration.

*Speaker

P2. The capacity of perturbation-based cerebellar learning

Johnatan Aljadeff * [1], Vincent Hakim [2], Boris Barbour [3], Nicolas Brunel [1]

[1] The University of Chicago – United States
[2] Laboratoire de Physique Statistique de l'ENS (LPS) – CNRS : UMR8550, Université Pierre et Marie Curie (UPMC) - Paris VI, Université Paris VII - Paris Diderot, École normale supérieure [ENS] - Paris – France
[3] École Normale Supérieure – CNRS : UMR8197 – France

Learning abstract input-output associations is a fundamental problem which is solved by multiple brain regions, including the cerebellum, cortex and hippocampus. The perceptron is a useful model to study this problem, providing a framework within which one can analytically compute the capacity - the maximal number of associations that can be stored - as a function of the properties of the neuronal substrate and the statistics of the associations. Finding a set of synaptic weights which implement a prescribed set of associations is traditionally done using "perceptron learning," a gradient-based rule that is guaranteed to converge if a solution exists. A major hurdle in relating this class of models to neurobiology is that perceptron learning requires that each neuron receives full error information: updates of each synapse are based on the error, or the difference between the output using the current synaptic weights and the desired output. In the cerebellum, climbing fibre input to Purkinje Cells (PC) is thought to convey error information used for learning in Parallel Fibre (PF) to PC synapses. Though this input is correlated with error in some cases, such a relationship does not seem to be universal. Thus, it is still unclear how the cerebellum solves the credit assignment problem, especially during acquisition of fine motor skills where the relationship between the error and the appropriate correction is complex. Recently we showed how plasticity rules measured in slice under physiological conditions can be mapped to a Stochastic Gradient Descent (SGD) learning algorithm (Bouvier et al., 2016 bioRxiv). In the proposed algorithm, spontaneous complex spikes (CS) perturb movements, and the existence or absence of a CS following the movement provides an evaluation. Here we study the capacity of SGD learning and show that under certain conditions the theoretical limit is achieved. Our analysis suggests that plasticity in the Deep Cerebellar Nuclei (DCN), which allows the system to maintain a running estimate of the errors, and PC inhibition of DCN neurons, play a crucial role in ensuring that different inputs do not interfere with each other during learning.

*Speaker

P3. Why is it hard to FORCE spikes?

Alexander Antrobus * [1], Peter Latham [2]

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

To drive movement, neural activity must trace out trajectories in a relatively low dimensional space – on the order of 10 or so dimensions, corresponding to the dimensionality of the relevant space of joint angles. It has been known for some time how to train networks of (non-spiking) rate neurons to exhibit low dimensional trajectories (Sussillo2009). However, it has been much harder to train networks of spiking neurons; the first such network appeared only two years ago (Huh2016). A somewhat surprising result from that, and other more recent work (Nicola2016) is that linear integrate and fire (LIF) networks are very difficult to train while quadratic integrate and fire (QIF) networks are relatively easy. Here we ask why. To address that question in a tractable setting, we return to rate networks, and investigate the relationship between single neuron properties and learnability. More specifically, we ask: what gain function (what transformation from input current to firing rate) allows networks to learn to exhibit low dimensional dynamics? We find that the main determinant of learnability is the steepness of the gain function: if it is too steep, the network becomes chaotic, and cannot exhibit low dimensional dynamics. This is consistent with the LIF/QIF result for spiking networks: the LIF gain function is much steeper than that of the QIF. Thus, the fact that LIF networks are so difficult to train may simply be because at low firing rate they are very sensitive to changes in synaptic drive.

*Speaker

P4. Local transformations of entorhinal grid cells in polarized enclosures

Marius Bauza * [1], Julija Krupic [2], Stephen Burton [1], John O'keefe [1]

[1] Sainsbury Wellcome Center (SWC UCL) – 25 Howland St, London W1T 4JG, United Kingdom
[2] Salk Institute (Salk) – 10010 N Torrey Pines Rd, La Jolla, CA 92037, United States

Entorhinal grid cells are active in multiple locations arranged in a hexagonal lattice [1]. They are thought to represent a universal metric for space based on their remarkable periodicity and invariance of spatial firing pattern [1, 2]. Recent experiments challenge this notion by demonstrating that grid cell symmetry is non-homogeneously disrupted in polarized enclosures [3]. Currently it is not clear whether grid distortions are correlated between different cells, which would potentially preclude their usage as a metric system. Hence it is crucial to simultaneously record from a large number of grid cells to evaluate how correlated their transformations are.
Here we present large scale grid cell recordings collected using Neuropixels, a new generation silicon probe [4]. This probe is capable of simultaneously recording from 384 channels that can address 960 recording sites. Neural signals are processed directly on the probe reducing cabling requirements and allowing chronic recordings from an unprecedented number of cells from freely moving animals.
Up to 50 spatially periodic cells were recorded simultaneously in different polygonal enclosures for more than 10 weeks after the implantation. We found that local changes in the geometry of the environment induce local transformations in grid cell pattern by causing individual grid fields to change their locations. The co-localized grid field shifts were strongly correlated and were independent of the scale. Finally, we will propose a computational mechanism explaining local transformations in the grid based on the structure of the enclosure.

1. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436,** 801–806 (2005).
2. Buzsáki, G. & Moser, E. I. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nat. Neurosci. 16, 130–138 (2013).
3. Krupic, J., Bauza, M., Burton, S., Barry, C. & O'Keefe, J. Grid cell symmetry is shaped by environmental geometry. *Nature* **518,** 232–235 (2015).
4. Bauza, M., Krupic, J., Burton, S., Harris, T., & O'Keefe, J. Large-scale single unit activity recorded with high-density silicon probes in awake freely behaving animals. SfN abstracts 848.09/MMM58, (2016)

*Speaker

P5. Population coding and correlated neural variability in early mouse vision

Alex Cayco Gajic * [1], Severine Durand [2], Joel Zylberberg [3], Michael Buice [2], Clay Reid [2], Eric Shea-Brown [4]

[1] University College, London – United Kingdom
[2] Allen Institute for Brain science – United States
[3] University of Colorado Denver – United States
[4] University of Washington – United States

How neural populations encode external stimuli is a fundamental question of neuroscience. Correlations have been shown to play an important role in population coding, either increasing or decreasing stimulus information depending on their structure. However, the transformation of correlated spiking between different brain regions is not yet understood. To address this we studied the spiking activity of neural populations in awake mice while presenting drifting gratings at various orientations and spatiotemporal frequencies. LGN and V1 populations were simultaneously recorded, allowing for a comparison of how correlated spiking is transformed through the mouse visual pathway. We analysed the performance of decoding stimulus orientation based on LGN and V1 spiking patterns over varying timescales, both before and after a trial-shuffling procedure that removes inter-neural correlations. We find that correlated spiking on short timescales increases decoding accuracy in LGN, but not V1 populations. Finally, we examine how the structure of the population code is transformed between LGN and V1.

*Speaker

P6. Learning-related reorganisation of cortical inhibition

Angus Chadwick * [1], Adil Khan [2], Jasper Poort [3], Antonin Blot [2], Thomas Mrsic-Flogel [2], Sonja Hofer [2], Maneesh Sahani [1]

[1] Gatsby Computational Neuroscience Unit, University College London – United Kingdom
[2] Biozentrum, University of Basel – Switzerland
[3] Sainsbury Wellcome Centre, University College London – United Kingdom

During learning, cortical networks undergo selective reorganisation in order to enhance processing of task-relevant information. However, an understanding of these changes in terms of cortical cell types and their interactions is lacking. Multiple distinct inhibitory cell types shape cortical network dynamics, including parvalbumin (PV), somatostatin (SOM) and vasointestinal-peptide (VIP) expressing interneurons. To investigate the role of these interneuron networks in learning, we simultaneously imaged the activity of pyramidal (PYR), PV, SOM and VIP neurons in primary visual cortex of mice as they learned a visual discrimination task. We observed complex and widespread changes in single-neuron responses and pairwise noise correlations over the course of learning, including a striking increase in the selectivity of PV responses to task-relevant stimuli. To assess the underlying network changes driving these response changes, we fit a linear dynamical system (LDS) model to the experimental data. This allowed us to infer functional interactions amongst simultaneously imaged neurons together with their stimulus-driven and locomotion-related inputs. Using the LDS, we found that PV selectivity changes were linked to the emergence of selective interactions within the PYR-PV subnetwork with learning. Moreover, increases in PV selectivity contributed to PYR selectivity increases, highlighting an important functional role for learning-dependent changes in inhibitory interactions. Thus, we find that learning modifies excitatory-inhibitory functional interactions, leading to enhanced representations of behaviourally-relevant stimuli.

*Speaker

P7. Bayesian inference of neuronal input dynamics from single-trial spike trains using mechanistic models

Christian Donner * [1,2], Manfred Opper [2,1], Josef Ladenbauer [1,2,3]

[1] Department of Software Engineering and Theoretical Computer Science, Technische Universität Berlin – Marchstraße 23, 10587 Berlin, Germany
[2] Bernstein Center for Computational Neuroscience Berlin – Philippstraße 13, Haus 6, 10115 Berlin, Germany
[3] Group for Neural Theory, Laboratoire de Neurosciences Cognitives, École Normale Supérieure, Paris – Ecole Normale Supérieure de Paris - ENS Paris – 29 rue d'Ulm, 75005 Paris, France

Technical advances over the past decades have led to a rapidly growing collection of neuronal spiking activity data recorded in vivo. These data often exhibit rich dynamics over multiple timescales. To quantitatively describe the observed spike trains using models that allow to study the underlying system mechanistically is a challenging task. However, for mechanistic neuron models of the integrate-and-fire (IF) type and fluctuating inputs, recently developed analytical techniques allow to efficiently calculate the likelihood of a given spike train (see contribution by Ladenbauer & Ostojic). These techniques at hand we consider as generative model the leaky IF neuron subject to a white noise input current whose mean varies over time. We infer the dynamics of mean input, spike rate and mean membrane voltage from single-trial spike trains. Assuming Markov process priors for the dynamics of the mean input an Expectation-Maximization (EM) algorithm is employed, which iteratively infers the distribution over the mean input trajectories (E-Step) and optimizes the hyperparameters for the underlying process (M-step). Three different prior dynamics are considered: a Gaussian diffusion-, a Markov jump-, and an Ornstein-Uhlenbeck-process. Via Bayesian model comparison our method correctly identifies the dynamics of the underlying process and accurately estimates the parameters from simulated data. We apply our method to extracellular spike train recordings from monkey visual cortex in a spontaneous state and characterize the dynamics of driving input and spiking activity for multiple cells. Our method allows for reliable classification and accurate inference of the neuronal input as well as spike rate dynamics using a well-established mechanistic model, enabling a quantitative link between its biophysical properties and the observed data.

*Speaker

P8. Leveraging heterogeneity for neural computation with fading memory

Renato Duarte * [1,2,3], Abigail Morrison [1,4]

[1] Institute of Advanced Simulation (IAS-6) & Institute of Neuroscience and Medicine (INM-6) & JARA BRAIN Institute I, Jülich Research Center (INM-6 / IAS-6) – Forschungszentrum Jülich GmbH Wilhelm-Johnen-Straße 52428 Jülich, Germany
[2] Bernstein Center Freiburg (BCF) – Hansastrasse 9a 79104 Freiburg Germany, Germany
[3] Institute for Adaptive and Neural Computation (ANC) – Institute for Adaptive and Neural Computation School of Informatics University of Edinburgh Informatics Forum 10 Crichton Street Edinburgh United Kingdom EH8 9AB, United Kingdom
[4] Institute of Cognitive Neuroscience, Faculty of Psychology, Ruhr-University Bochum (IKN) – Germany

Computational studies addressing the dynamics and computational properties of biologically inspired spiking neurons and networks tend to assume (often for the sake of analytical tractability) a great degree of homogeneity in both neuronal and connectivity parameters (e.g. [1]). The biophysical reality, however, is radically different from a homogeneous system and multiple levels of complex heterogeneous properties co-exist and shape a local circuit's emergent collective dynamics and information processing properties. Despite their varying molecular, morphological and physiological features, cortical modules can be seen as variations on a common theme [2]. The combined complexity of the circuit's heterogeneous building blocks can be used to provide a rich dynamical space, suitable for online information processing.
In this study, we set out to systematically evaluate the role played by different sources of heterogeneity (structural, neuronal and synaptic) in the characteristics of population dynamics and the circuit's capacity for online stimulus processing with fading memory, using cortical layer 2/3 microcircuits as a core inspiration for the circuit specification. We cross-reference various sources of experimental data regarding the composition and patterning of these microcircuits, accounting for different phenomena of interest (e.g. neuron types and corresponding sub-threshold characteristics, conductance properties of different receptor types, circuit-level connectivity and activity statistics, etc.), across different cortical regions, assuming a certain degree of generalization is possible. The methods applied in this study to quantify the dynamics and generic processing properties, being system-independent, can provide a valuable set of tools for microcircuit benchmarking. As carefully curated and organized datasets become increasingly available, it will become possible in the near future to apply increasingly realistic constrains and comparatively study the properties of realistic microcircuits, built to model specific cortical regions and input-output relations.

References:

- Brunel, Nicolas. **Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons.** *Journal of computational neuroscience* 8.3 (2000): 183-208

*Speaker

- Harris KD, Shepherd GMG: **The neocortical circuit: themes and variations**. *Nature Neuroscience* 2015, 18:170–181.

P9. Closed-loop estimation of retinal network sensitivity by local empirical linearization

Ulisse Ferrari * [1], Christophe Gardella [2], Olivier Marre [3], Thierry Mora [4]

[1] Institut de la Vision – CNRS : UMR7210, Inserm, Université Pierre et Marie Curie (UPMC) - Paris VI : UM80 – 17 Rue Moreau, 75020 Paris, France

[2] Institut de la Vision; Laboratoire de Physique Statistique ENS (IdV; LPS-ENS) – Institut de la Vision – France

[3] Institut de la vision – CNRS : UMR7210, Inserm, Université Pierre et Marie Curie (UPMC) - Paris VI : UM80 – 17 rue Moreau 75012, Paris, France

[4] Laboratoire de Physique Statistique de l'ENS (LPS) – CNRS : UMR8550, Université Pierre et Marie Curie (UPMC) - Paris VI, Université Paris VII - Paris Diderot, École normale supérieure [ENS] - Paris – 24, rue Lhomond 75005 Paris, France

Understanding how sensory systems process information is an open challenge mostly because these systems are non-linear. Here we present a novel perturbative approach to understand the non-linear processing performed by sensory neurons, where we linearize locally their responses to a stimulus. Starting from a reference stimulus, we added small perturbations to the stimulus and tested if they triggered visible changes in the responses. We updated the perturbations according to the previous responses using closed-loop experiments. We then developed a local linear model that accurately predicts the sensitivity of the neural responses to these perturbations. Applying this approach to the retina of the Long Evans rat, we could estimate the optimal performance of a neural decoder and show that the non-linear sensitivity of the retina is consistent with an efficient encoding of stimulus information. Our approach is general and can be used to characterize experimentally the sensitivity of neural systems to external stimuli, or relate their activity to behaviour.

*Speaker

P10. Variational learning of olfactory inference

Naoki Hiratani * [1], Peter Latham [1]

[1] Gatsby Computational Neuroscience Unit, University College London – United Kingdom

The mammalian olfactory system has an amazing ability to infer the identity of odors from the stochastic activity of broadly-tuned olfactory sensory neurons. Arguably, developmental plasticity in the olfactory bulb and the piriform cortex is responsible for the acquisition of this olfactory inference, yet little is known about how local synaptic learning rules in these circuits enable learning. This sensory learning is particularly difficult because in the natural environment, odors are rarely presented in isolation, and odor identities are rarely supervised. Furthermore, for most odors, a limited number of samples are available to the animal throughout its lifetime, so the learning rule needs to be data efficient. Moreover, even an adult animal can learn and update its olfactory representation when a novel odor is presented, so the learning rule should be adaptive. To address a potential mechanism of olfactory learning, we constructed a model of olfactory learning through a variational Bayesian method. The derived synaptic plasticity rule enables a neural circuit to acquire olfactory inference from a small number of samples. The model also suggests that the developmental change of membrane resistance observed in olfactory neurons enables neurons to process inputs based on its confidence in its sensory representation.

*Speaker

P11. A metric space approach to information theory for multi-unit recordings.

Conor Houghton * [1]

[1] University of Bristol – United Kingdom

For spike trains, information theory quantities, like mutual information, are typically calculated by discretizing time and converting the spike trains into sequences of ones and zeros. This approach suffers from a sampling problem, there are far too many words to sample properly. This problem is particularly acute for multi-unit recordings. There are clever methods to ameliorate the problem, here, however, we present a more straight-forward approach: the mutual information is calculated on a metric space of spike trains. The idea is simple: two nearby stimuli, for example, two nearby positions on a maze, should produce similar spike responses. This idea can be formalized to give a definition of mutual information which is not nearly so greedy for data as discretization.

*Speaker

P12. Effects of short-term plasticity on the memory lifetime of recurrent neural circuits

Sofia Jativa * [1], Taro Toyoizumi [2]

[1] Gatsby Computational Neuroscience Unit, UCL – 25 Howland St, Fitzrovia, London W1T 4JG, United Kingdom

[2] RIKEN Brain Science Institute – 2-1 Hirosawa, Wako, Saitama 351-0198, Japan

It is a substantial open challenge to understand how recurrent neural circuits can act as a memory buffer, despite fast forgetting by individual neurons. Previous work has primarily focused on understanding memory lifetimes of recurrent circuits with conventional static synapses. However, synapses are rich dynamical systems in their own right, and recent experimental findings [1,2] along with previous theoretical proposals [3,4] implicate these characteristics as supporting short-term memory.
Recently, an information-theoretic upper bound for memory lifetime was derived for linear recurrent networks. Specifically, it was shown that any linear network can, at most, achieve a memory life time proportional to the number of neurons in the network. Furthermore, it was shown that only a delay line, or any network that is equivalent to a delay line up to a unitary transformation, can saturate this bound [5]. Here, we extend this information-theoretic analysis to understand the role of dynamic synapses with short-term plasticity on memory performance. By linearizing a non-linear network, we study how short term plasticity modifies the effective connectivity matrix of the network to change the memory performance.

We tested this framework in different architectures, concentrating on networks with very poor memory performance i.e. normal networks. We show that dynamical synapses modify the internal structure of these networks and improve their memory performance. We will analyze the conditions under which memory performance is improved and under which the effective connectivity matrix approximates an effective delay line. We expect that the short-term plasticity of synapses might be key for understanding how recurrent neural circuits buffer temporal signals during cognitive processing, and furthermore it suggests a different way for synapses to be considered as a neural substrate of working memory.

1. Mark G. Stokes.'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. Trends in Cognitive Sciences, vol. 19, issue 7, pages 394–405. Cell i ii Press, July 2015.

2. M. Lundqvist, J. Rose, P. Herman, SL. Brincat, TJ. Buschman, and EK. Miller. Gamma and Beta Bursts Underlie Working Memory. Neuron, vol. 90, pages 152-64, April 2016.

3. Gianluigi Mongillo, Omri Barak, and Misha Tsodyks. Synaptic Theory of Working Memory.Science, vol. 319, issue 5869, pages 1543-1546, Mar 2008.

*Speaker

4. Dean V. Buonomano, and Wolfgang Maass. State-dependent computations: spatiotemporal processing in cortical networks. Nature Reviews Neuroscience, vol. 10, pages 113-125, Feb. 2009.
5. Surya Ganguli, Dongsung Huh, and Haim Sompolinsky. Memory Traces in Dynamical Systems. PNAS, vol. 105 no. 4, pages 18970-18975, Oct. 2008.

P13. Connectivity regimes of the stabilized supralinear network leading to bistable, persistent and oscillatory activity

Nataliya Kraynyukova * [1], Tatjana Tchumatchenko [1]

[1] Max-Planck-Institute for Brain Research, Frankfurt am Main – Germany

The neural circuits perform many nonlinear computations such as sub- and superlinear summation of inputs, memory storage and oscillations. All these fundamental computations are thought to be generated by nonlinear interactions in recurrently connected networks of excitatory and inhibitory neurons. Therefore, it is plausible to assume that a single recurrent network model could reproduce many of these computations. In the presented work we consider a 2D stabilized supralinear network (SSN) model (Persi et al. 2011, Ahmadian et al. 2013). Recently the SSN model has been shown to reproduce a variety of nonlinear computations such as normalization, surround suppression (Rubin et al. 2015) and stimulus induced variability suppression (Hennequin et al. 2016). In its simplest form the SSN model is a set of two coupled nonlinear differential equations that describe the activity of excitatory and inhibitory populations. The power-law activation function is motivated by the experimental studies (Priebe and Ferster, 2008) as well as theoretical evidence showing that the power-law is the only function consistent with contrast invariance (Miller and Troyer, 2002, Hansel and van Vreeswijk, 2002). Even though the SSN is one of the simplest nonlinear network models, it turns out few methods are available to systematically predict what type of steady state solutions to expect for all connectivity matrices and inputs to the network. We present a new method that allows to map 2D steady states of the SSN model to the zero crossings of a 1D characteristic function. This method allowed us to derive a number of new computational insights. First, we have shown that at most two stable steady states can coexist in the SSN model. Second, we have outlined connectivity and time scale regimes for the emergence of a persistent state in the SSN model. Third, we have proven that the SSN model can undergo a Hopf bifurcation and lead to stable oscillatory attractors.

*Speaker

P14. Classifiers with limited connectivity

Lyudmila Kushnir * [1], Stefano Fusi [2]

[1] Group for neural theory - Laboratoire de neurosciences cognitives, Département d'études cognitives, École normale supérieure, INSERM, PSL Research University (GNT-LNC-DEC-ENS-PSL) – École normale supérieure [ENS] - Paris : GNT-LNC-DEC – 29, rue d'Ulm 75005 Paris, France
[2] Center for Theoretical Neuroscience — Columbia University – Columbia University College of Physicians and Surgeons 1051 Riverside Dr, New York, NY, 10032, United States

For many neural network models based on perceptrons, the number of activity patterns that can be classified is limited by the number of plastic connections each neuron receives, even when the total number of neurons is much larger.
This poses a problem of how the biological brain can take advantage of its huge number of neurons given that the connectivity is extremely sparse, especially when long range connections are being considered.

One possible way to overcome this limitation in the case of feed-forward networks is to combine multiple perceptrons together, as in committee machines. In the committee machines the number of classifiable random patterns grows linearly with the size of the network even when each perceptron has limited connectivity. However, the problem is moved to the downstream readout neurons, which would require a number of connections that is as large as the number of perceptrons.

Instead of committee machines we propose to implement the readout by connecting multiple perceptrons in a recurrent attractor neural network (recurrent readout). We show by analytical calculations that for the recurrent readout the number of random classifiable patterns still grows unboundedly with the number of units, even when the connectivity of each perceptron remains finite. Importantly, both the recurrent connectivity and the connectivity of a downstream readout are also finite.

We construct such a network of linear threshold units with limited connectivity and compute its classification capacity for multiple regimes of operation. The regimes differ by the amount of dynamical noise and by the sparsity of input representations. We show that when the recurrent noise is large on the scale of the feedforward input, the capacity drops very fast when the input representations become sparse. On the contrary, for the lower amounts of noise, even for very sparse representations, (typical number of active inputs to a perceptron is around 1), we obtain the capacity close to the maximum for the given expansion ratio.

We also describe a somewhat paradoxical regime for sparse input representations, when the recurrent readout outperforms the majority vote of the committee machine. This is achieved by making the perceptrons receiving zero feedforward input for the given pattern (free perceptrons) to follow the consistent signal from those perceptrons that receive non-zero feedforward input.

*Speaker

This signal is relatively more important because it is made consistent by strong feedforward connections, while the initially random recurrent input from other free perceptrons is subjected to the dynamical noise. This advantage in principle allows to make the input representations even sparser while keeping the capacity in the reasonable range, but this requires tuning the amount of noise.

The analytical calculation is carried out for the case of binary classification task, but we show with simulations that the same linear scaling of the capacity holds for multi-way classification as well.
Our study shows that feedforward neural classifiers with numerous long range connections between different layers can be replaced by networks with sparse long range connectivity and local recurrent connectivity without sacrificing the classification performance.

P15. Internal Global Gain Modulations but not Stimulus Contrast Changes Preserve the Neural Code

Sangkyun Lee [1], Jiyoung Park [2], Stelios Smirnakis

1

[1] Brigham&Womens Hospital, Harvard Medical School – United States
[2] Baylor College of Medicine (BCM) – United States

External stimuli are represented by the firing of populations of neurons in neocortex. However, the strength of neuronal responses depends not only on the external stimulus that is presented but also on ubiquitously occurring spontaneous fluctuations of internal inputs. Internal input fluctuations modulate sensory responses strongly, altering firing rates of individual neurons by as much as an order of magnitude, even on repeated presentations of identical stimuli. *This raises the important unresolved question how neural population codes are preserved on the face of such fluctuations of internal input.*
The same question can be asked about external stimulus parameters, which strongly modulate the firing rate of individual neurons. Therefore, we studied whether population codes for direction are preserved under *1)* changes in stimulus contrast (external input) and under *2)* the spontaneous modulation of internal inputs (internal input).

We performed two-photon calcium imaging in mouse V1 layer 2/3 that expressed GCamP6 and observed responses of neural ensembles to moving oriented-gratings at different contrasts. Surprisingly, we found that the neural code is not preserved between different contrasts. This is because, while the shape of direction tuning function is preserved across contrasts at individual neurons, contrast response-gains are highly heterogeneous across neurons.

To observe whether population codes are preserved across different internal input levels, we partitioned the trials within each stimulus condition based on whether overall population activity level was high or low, and studied whether the code for direction was maintained in the two conditions. In contrast to visual contrast, we found that the population code for direction was largely invariant across the two conditions. This was true, even though changes in internal input caused larger firing rate modulations at the single cell level than the changes in stimulus contrast we employed. The reason for this appears to be that gain responses across cells are highly homogeneous during spontaneous internal input fluctuations, thereby preserving the code.
The principles outlined above have potentially important implications for neural circuit mechanisms of information encoding.

P16. Fitting mechanistic models to spike trains using analytical methods

Josef Ladenbauer * [1,2,3], Srdjan Ostojic [1]

[1] Group for Neural Theory, Laboratoire de Neurosciences Cognitives, École Normale Supérieure, Paris – Ecole Normale Supérieure de Paris - ENS Paris – 29 rue d'Ulm, 75005 Paris, France
[2] Department of Software Engineering and Theoretical Computer Science, Technische Universität Berlin – Marchstraße 23, 10587 Berlin, Germany
[3] Bernstein Center for Computational Neuroscience Berlin – Philippstraße 13, Haus 6, 10115 Berlin, Germany

Neuronal spike train data obtained from extracellular (in-vivo) recordings are being collected at an increasing pace. These data are typically analyzed by fitting parametric phenomenological models that describe statistical dependencies but allow only for limited biophysical/mechanistic insights. On the other hand, mechanistic models of neuronal spiking activity are useful to dissect observed phenomena, but they are rarely quantitatively matched to experimental spike train data due to methodological challenges.
Here we present an analytical approach to efficiently fit coupled spiking model neurons to observed (single-trial) spike trains. In particular, we employ a Bayesian framework to infer the mean and variance of the input, neuronal adaptation properties as well as connectivity for small networks of integrate-and-fire neurons. We compute the likelihood of a given spike train using the Fokker-Planck formalism: for constant input moments or weak perturbations and adaptation we obtain a solution via the first passage time problem; for strong input perturbations and adaptation we derive an approximation under the assumption of Poisson spiking statistics. Accuracy and efficiency of the proposed estimation method are assessed using simulated ground truth data.

We exploit our method to analyze extracellular recordings from primary auditory cortex of awake behaving ferrets. The animals were either passively listening or actively discriminating periodic click train stimuli of different rates. By fitting integrate-and-fire models to the data, we find that distinct strengths and time constants of the input currents are needed to account for the activity in the different behavioral states.
Altogether, our work provides useful techniques to quantitatively link mechanistic models of coupled spiking neurons with the statistics of measured spike train time series.

*Speaker

P17. Dynamical activity patterns in the macaque posterior parietal cortex during path integration

Kaushik Lakshminarasimhan * [1], Eric Avila [1], Roozbeh Kiani [2], Greg Deangelis [3], Xaq Pitkow [1,4], Dora Angelaki [1,4]

[1] Baylor College of Medicine – United States
[2] New York University – United States
[3] University of Rochester – United States
[4] Rice University – United States

Neural circuits evolved to deal with the complex demands of a dynamic and uncertain world. Binary-decision tasks with rigid trial structure are inadequate to fully reveal the rich structure of neural representations and computations that mediate fluid behaviour in the real world. To understand dynamic neural processing underlying natural behaviour, we trained macaque monkeys on a continuous-time foraging task in which they used a joystick to steer and catch fireflies in a virtual environment devoid of landmarks. Fireflies appeared briefly, one at a time, at a random location within the field of view. In order to solve the task, monkeys had to dynamically update their position estimates by integrating optic flow generated by self-motion. We used laminar probes to sample the activity of a large number of neurons in the posterior parietal cortex and found that different neurons were active during different epochs of integration. This was confirmed in large-scale simultaneous recordings over several days using a chronic multielectrode array: Neurons exhibited rich temporal diversity such that the integration dynamics were embedded in the dynamical pattern of population activity. We are currently applying statistical techniques to characterise the precise dynamics of population activity to understand the associated neural computations.

*Speaker

P18. Sparse Predictive Coding in Balanced Spiking Networks

Mirjana Maras * [1], Sophie Deneve [1]

[1] Ecole normale supérieure (ENS Paris) – École normale supérieure - Paris – 29, rue d'Ulm, 75005 Paris, France

Recent theoretical and experimental studies are challenging the widely held notion that neural populations represent sensory input via time averaged firing rates. Experimentally observed tight balance between excitation and inhibition [1] supports the non-traditional view that neurons encode information in their precisely timed spike trains. While theoretical studies show that this tight balance can be achieved in both sparsely and densely connected networks with random weights, they do not address why the brain would operate in this regime [2,3]. On the other hand, the predictive spike-coding framework shows that the tight balance between excitation and inhibition can be seen as a direct consequence of a coding efficiency principle [4,5]. The predictive coding model is an all-to-all connected spiking network with fast inhibition and symmetric recurrent weights. It reconstructs sensory variables with an optimal number of spikes, and outperforms the Poisson code due to its faster population coding error decrease with the number of spikes. Relaxation of the fast inhibition constraint has shown that the model continues to outperform the Poisson network when realistic synaptic dynamics are modeled [6,7]. Here, we remove the all-to-all symmetric connectivity constraint, and analyze the performance of the network with both realistic synaptic dynamics and sparse lateral connections. We train the recurrent weight matrix by employing a local balance restoring learning rule [8]. Furthermore, we control the population firing rate via a homeostatic adaptation of the firing threshold.

We show that the revised model successfully learns the lateral connection weights with biologically realistic connection probabilities and synaptic delays. The resulting prediction errors are greater than those of the optimal all-to-all connected network, but the sparse network outperforms the equivalent Poisson and independent leaky integrate-and-fire networks. Furthermore, the mean of the squared learned weights scales inversely with the total number of synaptic inputs, as experimentally observed [9].

References

1. M. Okun, I. Lampl. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities, Nature Neuroscience 11 pp. 535-537, 2008.
2. T.P. Vogels et al. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. Science 334 pp. 1569-1573, 2011.
3. A. Renart et al. The asynchronous state in cortical circuits. Science 327 pp. 587-590, 2010.
4. M. Boerlin, C.K. Machens, S. Deneve. Predictive coding of dynamical variables in balanced spiking networks. PLOS Computational Biology 9 e1003258, 2013.
5. S. Deneve, C.K. Machens. Efficient codes and balanced networks. Nature Neuroscience 19

*Speaker

pp. 375-382, 2016.
6. M.A. Schwemmer et al. Constructing precisely computing networks with biophysical spiking neurons. J. Neurosci. 35 pp. 10112-10134, 2015.
7. M. Chalk, B. Gutkin, S. Deneve. Neural oscillations as a signature of efficient coding in the presence of synaptic delays. eLife 5 e13824, 2016.
8. R. Bourdoukan, S. Deneve. Enforcing balance allows local supervised learning in spiking recurrent networks. NIPS pp. 982-990, 2015.
9. J. Barral, A.D. Reyes. Synaptic scaling rule preserves excitatory/inhibitory balance and salient neuronal dynamics. Nature Neuroscience 19 pp. 1690-1696, 2016.

P19. Metastable dynamics drive anticipatory neural activity

Luca Mazzucato * [1], Giancarlo La Camera [1], Alfredo Fontanini [1]

[1] Department of Neurobiology and Behavior, Stony Brook University – United States

Sensory stimuli can be recognized faster when their delivery is expected, compared to when they are unexpected. When taste stimuli are preceded by an anticipatory cue (the same for all stimuli), neural activity in the gustatory cortex encodes stimulus-related information faster than in absence of the cue. However, the mechanism linking cue responses to faster encoding is unknown. Here, we elucidate this process using a biologically plausible model based on a recurrent network with clustered architecture. Slow fluctuations in this network generate state sequences as observed in experiment. The anticipatory cue accelerates transition rates between states, leading to faster onset of stimulus-coding states. Their shorter latency mediates the faster encoding due to expectation. This effect was confirmed in ensemble recordings from the gustatory cortex of alert rats. Anticipatory neural activity was unrelated to changes in selective neurons firing rates and was absent in homogeneous networks, suggesting that a clustered architecture is necessary to mediate the expectation of stimuli in cortex. Our results demonstrate a novel mechanism for speeding up sensory coding in cortical circuits and provide a new framework to investigate sensory processing.

*Speaker

P20. Inhibitory neurons: a nidus for feature-selective sub-network organization in mouse V1?

Ganna Palagina [1], Maria Plakia [2], Ioannis Smyrnakis [3], Maria Papadopouli [4], Stelios Smirnakis * [5]

[1] Brigham and Women's H. and Jamaica Plain VA H., Harvard Medical School – United States
[2] Institute of Computer Science, Foundation for Research and Technology-Hellas – Greece
[3] Department of Computer Engineering, Technological and Educational Institute of Crete – Greece
[4] Department of Computer Science, University of Crete and Institute of Computer Science, Foundation for Research and Technology-Hellas – Greece
[5] Brigham and Women's Hospital and JP VA Hospital, Harvard Medical School – 60 Fenwood Road, Boston, MA, 02115, United States

In mammalian brains, the encoding of sensory stimuli relies on the coordinated firing of large neuronal groups rather than single isolated neurons. However, the rules by which these groups form, evolve, and encode the information remain elusive. A leading hypothesis is that pyramidal cells with similar function and tuning preferentially connect to each other, forming dynamic multicellular encoding units yoked to a common cause. The existence of such groups should inevitably affect the profile of spontaneous events observed in neocortical networks and resting-state functional connectivity [1]. We used 2-photon calcium imaging to study spontaneous population firing events (bursts) in layer 2/3 of mouse area V1 during postnatal maturation (postnatal day 8–45). During this period the size of spontaneously occurring population bursts formed a scale-free distribution obeying a power law. The same was true for the degree of functional connectivity, a measure of pairwise synchrony across cells. These observations are consistent with a hierarchical small-world-net architecture [2], characterized by groups of cells with high local inter-connectivity (”small worlds”) connected to each other via a restricted number of links between ”hub” cells [3]. To identify candidate ”small world” groups, we applied variants of the spike time tiling coefficient measure introduced by [4] to cluster pyramidal neurons into groups based on the synchrony of firing between: 1) pyramidal neuron pairs, as well as 2) pyramidal neurons and neighboring interneurons. This strategy allowed us to reliably identify clusters of pyramidal neurons ”linked” to one or more local interneurons. These ”small-world” clusters did not remain static during postnatal development: both cluster size and overlap with other clusters decreased over time as pyramidal neurons became progressively more selective, ”linking” to fewer neighboring interneurons. By juvenile adulthood pyramidal cells ”linked” on average with _˜1.4 interneurons. Notably, neuronal clusters identified during epochs of spontaneous activity turned out to have functional significance. Specifically, pyramidal neurons in a cluster show higher orientation/direction tuning function similarity than expected with each other and with their ”linked” interneurons to diverse stimuli. Our findings suggest that spontaneous population events in the visual cortex are shaped by ”small-world” networks of pyramidal neurons that share functional properties and work in concert with one or more local

*Speaker

interneurons. We argue that such groups may represent a fundamental neocortical unit of computation at the population level. [1]. Kenet T et al. Nature 425: 954 – 956 (2003). [2]. Sporns O. Networks of the Brain. (MIT press, 2011). [3]. Bonifazi P et al. Science 326(5958):1419-24 (2009). [4]. Cutts SC and Eglen CJ. J. Neurosci. 34:14288-303 (2014).

P12. Context-dependent associative memory through modular subneworks

William Podlaski * [1], Tim Vogels [1]

[1] Centre for Neural Circuits and Behaviour, University of Oxford (CNCB) – United Kingdom

One of the best known models for associative memory in the brain is the Hebbian assembly, in which recurrent connectivity between select groups of neurons forms the basis of memory patterns. The classic model of this phenomenon is the Hopfield network (Hopfield 1982). A common issue with these types of models is catastrophic interference – memories can become unstable due to interactions from other memories. We hypothesized that the effects of interference can be lessened by adding context-dependent modularization. Such modularization can be imposed by defining a set of overlapping subnetworks, each representing a different context. Only one subnetwork is active at a given time, thereby reducing interference from memories found in other subnetworks, which remain dormant. We show that these 'modular Hopfield networks' have greatly increased memory capacity which depends on two parameters – the ratio of subnetwork to full network size, and the total number of subnetworks. The capacity of the network exceeds that of the standard Hopfield model for the majority of parameter values, and, if chosen optimally, is up to an order of magnitude greater. We verified the theoretical capacity predictions with numerical simulations, and analyzed the robustness of memory patterns to noise. The increased capacity comes at the cost of limited retrieval, as only memories stored in the active subnetwork can be recalled. To overcome this, we propose a system in which a controller network dynamically switches to a desired contextual state before retrieval. Our work provides functional implications for several key ideas about neural circuits, including the sparsity of neural activity, mixed selectivity of individual neurons, and top-down inhibitory control of circuit function.

*Speaker

P22. A reservoir computing model of the interaction between cortex and basal ganglia in motor learning

Ryan Pyle * [1], Robert Rosenbaum *

1

[1] Department of Applied and Computational Mathematics and Statistics, University of Notre Dame – 153 Hurley Hall, Notre Dame, IN 46556, United States

Neural circuits generate intricate dynamical responses to simple sensory inputs. For example, after training, a target on a screen evokes a sequence of coordinated muscle contractions to reach an arm to that target. This motor output is generated, at least in part, by firing rate dynamics in populations of motor cortical neurons. The question of how the cortex generates intricate, dynamical firing rate responses is a topic of ongoing research. A compelling hypothesis is that the brain utilizes the rich, high-dimensional dynamics of recurrent neuronal networks. This hypothesis is formalized computationally by reservoir computing models in which a recurrent network serves as a "reservoir" of dynamics that a learning algorithm mines for target activity patterns.

Reservoir computing models are capable of quickly learning complex responses to simple inputs and the firing rate dynamics they generate resemble those of cortical populations. Despite their success in modeling the generation of motor output, however, most existing reservoir computing models use unrealistic learning rules that require a teacher signal already capable of generating the target output. Biological motor learning relies on lower-dimensional, reinforcement signals that are not always related to motor output in a simple way. Existing reinforcement learning rules for reservoir computing fail to converge on non-trivial tasks.

We show that functional reservoir computing with realistic learning rules is achieved by properly accounting for the role of the basal ganglia. Data from songbirds and mammals support a two-pathway model of motor learning in which initial proficiency is obtained by the basal ganglia or their homolog, while parallel cortical pathways gradually take over as the task becomes well-learned. We present a novel reservoir computing model in which a basal ganglia pathway uses noise to explore motor response space, learning through a one dimensional reward signal modeling dopamine release. Meanwhile, a parallel cortical pathway gradually learns through pure Hebbian plasticity to mimic the basal ganglia pathway. This model can learn realistic motor tasks and reproduces some widely observed behavioral and neurophysiological phenomena related to motor learning. Our results extend the reservoir computing paradigm beyond the domain of motor output generation to the domain of biological motor learning.

*Speaker

P23. Effects of monocular-deprivation induced plasticity on the dynamics of recurrent networks

Leonidas Richter * [1], Julijana Gjorgjieva [1,2]

[1] Max Planck Institute for Brain Research – 60438 Frankfurt, Germany
[2] School of Life Sciences, Technical University of Munich (TUM) – 85354 Freising, Germany

Monocular deprivation (MD) as an experimental paradigm has long been used to to study mechanisms guiding the development and plasticity of visual cortical circuits. However, in contrast to the binocular region of the primary visual cortex (V1b), the effects of MD in the monocular cortex (V1m) are less well studied, while experimental results suggest that different plasticity mechanisms operate in these regions during MD. In our modelling study we aim at understanding the dynamical aspects of network plasticity in V1m following MD. We incorporate recent experimental findings from long term *in vivo* recordings in rodents and use dynamical systems theory and simulations of large-scale spiking networks to study the effects of plasticity at thalamocortical and intracortical synapses in recurrent networks with both excitatory and inhibitory units. We find that the effects of cortical plastic reorganisation change qualitatively depending on the operating regime of the recurrent network, controlled by the recurrent coupling scale and the structure of inputs driving the network. Incorporating strong feedforward inhibition and decoupling the background input (coming from other cortical areas) from the thalamic drive modulates the nature of dynamical changes in response to MD induced cortical plasticity. Our aim is to understand how structure in cortical networks that suits the computational demands emerges during development and how it is shaped by experience. Understanding this may also shed light onto the mechanisms underlying pathological development.

*Speaker

P24. Bayesian inference of dynamic functional couplings from single-trial neuronal data

Christian Donner * [1,2], Manfred Opper [1,2]

[1] Bernstein Center for Computational Neuroscience Berlin – Bernstein Center for Computational Neuroscience Humboldt-Universität zu Berlin Philippstr. 13, Haus 6 10115 Berlin, Germany
[2] Department of Software Engineering and Theoretical Computer Science, Technische Universität Berlin – TU Berlin Fakultät IV Elektrotechnik und Informatik Marchsstrasse 23 D-10587 Berlin, Germany

Neurons concertedly code information about the external world, internal states and actions in a highly dynamic manner. Those dynamics pose a challenge for characterizing the underlying network given extracellularly and simultaneously recorded activity of multiple neurons. Previously proposed methods often discard temporal structure of the data and assume stationarity over time. Here we aim for a phenomenological characterization of the recorded network without such explicit assumptions. To do so we propose a hierarchical Bayesian model, which assumes that the recorded network switches between a set of functional states and these states are different parameterizations of a kinetic Ising model with asynchronous updates. Our main contributions are (a) Bayesian inference for the kinetic Ising model (b) with incorporating a sparsity prior and (c) embedding this in a hierarchical Bayesian model with Markov dynamics on the latent states. We develop a Variational Bayes algorithm that accurately recovers the parameters given a temporal sequence of binary activity from multiple neurons. The algorithm is able to infer time-dependent functional couplings without any explicit assumptions of stationarity and with a temporal resolution of 1 ms. Finally, by inferring the model parameters from in-vivo V4 monkey data (40 neurons) we identify high variability in the activity. Certain functional states can be related to external stimuli, while others coincide with events in the simultaneously recorded local field potential.

*Speaker

P25. Cooperative Tuning and Timing of Excitatory and Inhibitory Inputs in the Olfactory Cortex Homolog of Zebrafish

Peter Rupprecht * [1,2], Rainer Friedrich [1,2]

[1] Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland (FMI) – Switzerland
[2] University of Basel, Basel, Switzerland – Switzerland

A neuron can be described as a processing unit that receives synaptic inputs and fires in function of those excitatory and inhibitory inputs. To understand how inhibition and excitation work together on the level of single neurons, we use whole-cell voltage clamp recordings in the zebrafish homolog of olfactory cortex, Dp, together with calcium imaging and optogenetics, and connect our findings to ideas from theoretical neuroscience on balanced networks. In Dp, we find two clearly separate sub-regions that show complementary characteristics in terms of tuning and timing of inhibition and excitation. Upon odor stimulation, neurons in the anterior-most region of Dp (aDp) exhibit persistent but stochastic excitatory input currents that come directly from the olfactory bulb and that are largely unbalanced by inhibition. Neurons in the posterior part of Dp (pDp), in contrast, exhibit strong, but transient input currents that are stemming from local connections of the network. We find that pDp neurons live in a state of balanced excitation and inhibition during this transient odor response phase. Experimentally, using simultaneous whole-cell and LFP recordings, we observe that this balance is maintained on a short timescale of few milliseconds, suggesting a 'tightly balanced' state. During the first few hundred milliseconds of the stimulus response, however, the coarse balance is tipped toward inhibition, generating a window not of opportunity, but of impediment for Dp neurons to respond to inputs from the olfactory bulb. Those inputs are known to be less informative in this early phase. Together, we describe and dissect the different strategies that are employed by different neuronal populations in Dp to use or discard inputs from the olfactory bulb.

*Speaker

P26. Flexible coding by cell assemblies

Eleonora Russo * [1], Daniel Durstewitz [2]

[1] Department of Theoretical Neuroscience, Bernstein Center for Computational Neuroscience, Central Institute of Mental Health, Heidelberg University – J5, 68159, Mannheim, Germany
[2] Department of Theoretical Neuroscience, Bernstein Center for Computational Neuroscience, Central Institute of Mental Health, Medical Faculty Mannheim, Heidelberg University – J5, 68159, Mannheim, Germany

An increasing body of evidence using simultaneous recordings from many neurons suggests that information is encoded in spatio-temporally organized groups of neurons (cell assemblies) rather than by single cells. Despite the growing consensus that spatio-temporal groupings of neurons may play a pivotal role in the control and representation of behavior, cell assemblies are still not well characterized and no agreement has been reached on the temporal resolution and structure that best characterizes neural activity patterns.

We have recently developed a novel statistical algorithm for detecting assembly structure in multiple single units recordings across a variety of time scales and time lag constellations, and have used it to analyze multiple single-unit recordings from PFC, CA1 and EC during a delayed alternation and a free exploration task. Our method is not bound to a specific a-priori hypothesis about the data, but treats the temporal scale or internal organization of assembly patterns as free parameters. Using this approach, we found that assembly structure and temporal precision are not universal properties of cortical coding, but depend on the brain area recorded from and on the behavioral task demands.

While this statistical framework unifies a variety of observations on neuronal groupings reported in the literature, it raises the question of how the different assembly codes within various areas are coordinated such as to enable information integration and transfer across areas. One of the possible mechanisms is through neural oscillations. Here we study assembly patterns and their relation to oscillatory activity in PFC-CA1-PC recordings of sleeping and behaving rats.

*Speaker

P27. On the Compressibility of Brain Activity Maps

Simon Schultz * [3,2,1]

[3] MRC Brain Network Dynamics Unit – Oxford University, Mansfield Road, Oxford, OX1 3TH, United Kingdom
[2] Jesus College, Oxford – Turl Street, Oxford, OX1 3DW, United Kingdom
[1] Department of Bioengineering, Imperial College London – South Kensington, London, SW7 2AZ, United Kingdom

Imagine that you had recorded the activity of all neurons in a human brain over a certain duration (a Brain Activity Map, or BAM), and wished to transmit this information, for instance to download into a remote robotic agent. How much bandwidth would be required to send this BAM over a data link? If individual neurons fired at 1 Hz according to a Poisson process, we would require on average 11.4 bits/sec to specify the timing of each action potential to 1 ms resolution. Presupposing further that all neurons in the brain were independent (clearly far from the truth), a data rate of approximately 8 TB/min would be required (for comparison, you can send approximately 7.5 GB/min across the gigabit Ethernet connections found in many universities). Of course, we have considered only worst case scenarios here: taking into account redundancies in both time and space, the actual information content is likely to be significantly lower, suggesting that by applying an appropriate lossless compression algorithm, the BAM could be transmitted with a substantially smaller number of bits. I will present data showing that for a range of neuronal types from the cerebellum, hippocampus and neocortex, individual spike trains can be compressed beyond the amount of information required to specify a Poisson event train, but not drastically so. The picture for redundancies across space is however somewhat rosier: using large-scale multiphoton calcium imaging and electrophysiological datasets, I will show that substantial compression rates can be achieved in comparison to the independent scenario, and examine the scaling properties of compression with neuronal ensemble size. Compression is algorithm dependent, and greater compression can be achieved with algorithms that incorporate knowledge of statistical dependencies in spike train structure, in comparison to "universal" algorithms such as Lempel-Ziv. As well as being of practical utility, the development of compression algorithms optimised for brain activity maps can tell us important things about the scaling of complexity in brain dynamics, with the potential to lead to insights into brain-scale coordination of neuronal information processing.

*Speaker

P28. A theory of mesoscopic neural activity in cortical circuits that links to microscopic parameters

Tilo Schwalger * [1], Moritz Deger [1], Wulfram Gerstner [1]

[1] Laboratory of Computational Neuroscience (LCN-EPFL) – Laboratory of Computational Neuroscience, Brain and Mind Institute, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

Linking large-scale brain activity to the properties of single neurons and connectivity is vital for a mechanistic interpretation of mesoscopic and macroscopic neural data. Population models such as Wilson-Cowan equations, neural mass or field models are widely used but lack a clear link to the underlying microscopic parameters. Here we systematically derive mesoscopic activity equations for several interacting populations starting from a microscopic network of generalized integrate-and-fire (GIF) neurons or generalized linear models (GLM). Each population consists of 50 – 2000 neurons of the same type but different populations account for different neuron types. Our mesoscopic theory captures important properties of population activity such as finite-size fluctuations due to the limited number of neurons per population and pronounced spike-history effects caused by refractoriness and adaptation on the cellular level. The mesoscopic model accurately reproduces rich activity dynamics obtained from a microscopic simulations of the spiking neural network including stochastic jumps between multistable states and synchronization in balanced networks of excitatory and inhibitory neurons. We use the mesoscopic model to predict non-stationary neural activity in a multi-laminar model of a cortical microcircuit consisting of eight neuron types under thalamic inputs.

Our theory offers a quantitative framework for modeling cortical information processing on a mesoscopic level based on single cell and connectivity parameters. We expect that such a model will become useful to predict the outcome of experiments such as optogenetic stimulation of a subgroup of neurons. Vice versa, having the link to the microscopic parameters also suggests new possibilities to infer neuronal and connectivity parameters from mesoscopic data.

Reference: T. Schwalger, M. Deger, and W. Gerstner. PLoS Comput. Biol., 13(4):e1005507,

*Speaker

2017.

P29. Dynamically constrained vs unconstrained linear models of evidence integration in a contextual decision making task

Joana Soldado Magraner * [1], Maneesh Sahani [1], Valerio Mante [2]

[1] The Gatsby Computational Neuroscience Unit (GCNU) – 25 Howland St, London W1T 4JG, United Kingdom
[2] Institute of Neuroinformatics (INI) – Winterthurerstrasse 190, 8057 Zürich, Switzerland

Circuits in PFC are believed to integrate sensory evidence and carry out context-dependent computations. In a recent study by Mante et al. [1], monkeys were trained to perform a dual-context decision-making task, learning to select a contextually-relevant stimulus while ignoring an irrelevant one. Several dynamical models have been proposed [1] [2] to explain how the circuit could be capable of performing selective integration without gating the irrelevant information. In [2], we took the approach of inferring a linear dynamical system (LDS) model [3] directly from the data. The model successfully captures the main features of the population time courses, despite being restricted to evolve under a simple linear dynamical rule. Given this results, we wanted to address the following question, is a dynamically restricted model flexible enough to accurately capture the population responses? or put it in other words, how much do we loose by taking this dynamical assumption? In order to address this, we introduce a simple linear model which has the flexibility to learn a different parametrization at each point in time. The LDS model belongs to the same model class and it is nested within the new model, as it has less flexibility. We compare the performance of both models by assessing their generalization capability across different conditions. Finally, we compare the nature of their solutions with the goal of understanding which insights the dynamical constraint can provide into the selective integration mechanism.

References

1. Context-dependent computation by recurrent dynamics in prefrontal cortex. Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. Nature 503(7474):78-84 2013.

2. Linear dynamics of evidence integration in a contextual decision making task. Soldado-Magraner, J., Mante, V., and Sahani, M. Cosyne 2015.

3. Empirical models of spiking in neural populations. Macke, J. H., Busing, L., Cunningham, J. P., Yu, B. M., Shenoy, K. V., Sahani, M. NIPS 2011.

*Speaker

P30. Diverse population coupling ensures robust yet flexible stimulus representation in a recurrent network model of perceptual learning

Yann Sweeney * [1], Claudia Clopath [1]

[1] Imperial College London – South Kensington, London SW7 2AZ, UK, United Kingdom

Stimulus representation in mouse visual cortex neurons is highly plastic throughout perceptual learning. Task-relevant visual patterns become better represented by the neuronal population as behavioural performance improves. Existing theories of perceptual learning describe plasticity of stimulus representation as the contextual top-down modulation of bottom-up signals. Local synaptic connections account for a significant portion of excitatory current in layer 2/3 pyramidal cells in V1, while the number of these connections vary widely across neurons in the network. Neurons with a large number of recurrent synaptic connections are highly coupled to the population-wide activity, while relatively unconnected neurons have low population coupling. However, the impact of local recurrent connectivity on top-down contextual modulations remains unknown.
Using a computational recurrent network model of perceptual learning we study: i) how does a neurons' population coupling impact stimulus representation plasticity during perceptual learning, and ii) how can diverse population coupling emerge and be maintained throughout synaptic plasticity? We find that the distribution of neurons' population coupling within the network has a dramatic effect on the network's ability to form new stimulus associations from combining bottom-up and top-down information. Diverse population coupling ensures that there are both weakly coupled neurons with stable stimulus representations that faithfully reflect their bottom-up inputs, and strongly coupled neurons with stimulus representations that can be flexibly modulated by top-down signals during perceptual learning. Surprisingly, we find that diverse population coupling emerges in networks of neurons with diverse learning rates: neurons with slow learning rates maintain low population coupling and robust stimulus representations, whereas neurons with fast learning rates maintain a high population coupling, enabling them to flexibly learn new stimulus representations.
Our network model predicts that neurons with high population coupling will be more plastic than neurons with low population coupling, and we test this prediction using large-scale calcium imaging datasets from the Allen Brain Observatory. Our analysis reveals that neurons with higher population coupling tend to exhibit higher plasticity of their stimulus selectivity, corroborating the predicted link between population coupling and plasticity.

*Speaker

P31. Sparse connectivity for MAP inference in linear models using sister mitral cells

Sina Tootoonian * [1]

[1] University College London (UCL) – 25 Howland Street, London, W1T 4JG, United Kingdom

Sensory processing is hard because the variables of interest are encoded in spike trains in a relatively complex way. A major goal in sensory processing is to understand how the brain extracts those variables. Here we revisit a common encoding model in which variables are encoded linearly. Although there are typically more variables than neurons, this problem is still solvable because only a small number of variables appear at any one time (sparse prior). However, previous solutions usually require all-to-all connectivity, inconsistent with the sparse connectivity seen in the brain. Here we propose a principled algorithm that provably reaches the MAP inference solution but using sparse connectivity. Our algorithm is inspired by the mouse olfactory bulb, but our approach is general enough to apply to other modalities; in addition, it should be possible to extend it to nonlinear encoding models.

*Speaker

P32. A state space model for change point detection in multivariate spike count data

Hazem Toutounji * [1], Daniel Durstewitz [1]

[1] Department of Theoretical Neuroscience, Bernstein Center for Computational Neuroscience, Central Institute of Mental Health, Medical Faculty Mannheim, Heidelberg University – J5, 68159, Mannheim, Germany

Neural activity from higher cortical areas in awake, behaving animals has highly dynamic and strongly nonstationary population-wide response properties. Nonstationary events in the form of changes in the firing rate statistics of recorded spike count time series may arise from a variety of sources. They may encode features of the experimental paradigm and could, potentially, correspond to the neural computations associated with the performance of a given behavioural task. In order to identify nonstationary events, or *change points*, in neural data, and to relate these events to behaviour, a model-based approach to change point detection may bear certain advantages not shared by model-free techniques. Here, we develop such an approach for detecting and parametrising multiple changes in multivariate spike count data within the statistical framework of *State Space Models* (SSM). The model assumes a nonlinear, nonstationary, autoregressive Gaussian process, parametrised by by the onset of change and its time scale, that captures relevant features of the underlying latent neural dynamics. Given their discrete, nonnegative nature, high-dimensional spike count time series are generated from the low-dimensional latent states through a Poisson observation model. We devised an initialisation algorithm, a model selection method and an estimation procedure that makes for a practical and efficient solution to change point identification from large data sets, such as recordings from developmental studies. Model parameters are constrained in a way that assures model identifiability, which we demonstrate by estimating latent states and model parameters from synthetic data where the ground truth is known. We also show that *population-wide* change points and their time scales can be reliably estimated, even when this information is lost by the averaging effect of classical methods, such as summarising data in peristimulus time histograms. Given that individual neurons may show selective responses to certain experimental features and, consequently, change points, an issue of overestimating the relevant time scales emerges, which we successfully remedy by regularising observation model parameters. As a real data example, the model is fitted to multiple single unit recordings from rat medial prefrontal cortex neurons during an operant rule switching task. The resulting reconstruction of the underlying nonstationary dynamics allows matching the neural correlates of learning to their behavioural counterpart by relating behavioural changes to population-wide change points, as estimated by the model.

*Speaker

P33. Hierarchical inference, Helmholtz machines and distributed representations of uncertainty

Eszter Vértes * [1], Maneesh Sahani [1]

[1] Gatsby Unit, UCL – United Kingdom

Sensory systems in the brain are tasked with making sense of noisy, often ambiguous incoming stimuli. Forming a percept based on receptor activations in the periphery is challenging – the underlying computation is probabilistic in nature. This is partly because of the inherent stochasticity of the receptors (sensory noise) and more importantly because each sensory stimulus may be consistent with a number of different unobserved generative causes.
To be able to perform rich inference taking into account the natural statistics of sensory inputs, as observed in behavioral experiments, neural circuits must learn about the underlying generative model. A number of schemes have been suggested for how populations of neurons may code for uncertainty (e.g. [1,2,3]), but there has been very little work on how such representations could be acquired by neural systems.

We propose a new approach, the DDC-Helmholtz machine, to learn a generative model of sensory stimuli and simultaneously learn to accurately infer the corresponding explanatory (or latent) variables. In our model neural activity corresponds to the inferred posterior distribution over latent causes of the observations. These densities are represented in a distributed fashion by a set of expectations in a "distributed distributional code" (DDC) [3,4]. Using this representation we are able learn hierarchical generative models and a corresponding recognition model producing estimates of the posterior distribution. To learn both the generative and the recognition model we use a *wake-sleep*-like algorithm inspired by the Helmholtz machine [5]. It requires only local updates, making our approach biologically appealing. Furthermore, the posterior representation does not impose independence or a rigid parametric structure, thus it is able to capture the statistical dependencies of the latent causes faithfully.

We evaluate the quantitative performance of the algorithm by learning a hierarchical sparse representation of synthetic data and natural image patches. Our results show that the DDC-Helmholtz machine can capture the data and the underlying latent factors well, outperforming standard variational methods.

1. Patrick O. Hoyer and Aapo Hyv́arinen. Interpreting neural response variability as monte carlo sampling of the posterior. Adv Neural Info Processing Sys, volume 15, pages 277–284. MIT Press, 2003.
2. Ma, W. J., Beck, J. M., Latham, P. E., Pouget, A. (2006). Bayesian inference with probabilistic population codes. Nature neuroscience, 9(11), 1432-1438.

*Speaker

3. R S Zemel, P Dayan, and A Pouget. Probabilistic interpretation of population codes. Neural Comput, 10(2):403–430, 1998.
4. Maneesh Sahani and Peter Dayan. Doubly distributional population codes: simultaneous representation of uncertainty and multiplicity. Neural Computation, 15(10):2255–2279, October 2003.
5. Peter Dayan, Geoffrey E. Hinton, Radford M. Neal, and Richard S. Zemel. The Helmholtz Machine. Neural Computation, 7(5):889–904, September 1995.

P34. From Sparse to Balanced Activity: The Spectrum of Asynchronous Dynamics in Neocortical Networks

Yann Zerlaut * [1], Stefano Zucca [2], Tommaso Fellin [4,3], Stefano Panzeri [4,5]

[1] Neural Coding Laboratory. Istituto Italiano di Tecnologia (NCL IIT) – Italy
[2] Optical Approach to Brain Function. Istituto Italiano di Tecnologia – Italy
[4] Neural Coding Laboratory. Istituto Italiano di Tecnologia – Italy
[3] Optical Appraoch to brain Function. Istituto Italiano di Tecnologia – Italy
[5] Center for Neuroscience and Cognitive Systems @UniTn, Istituto Italiano di Tecnologia – Italy

The classical theoretical model for asynchronous dynamics in neocortical networks assumes a balance of excitatory and inhibitory currents targeting individual neurons. However, through voltage-clamp measurements during spontaneous activity in sensory cortices, in vivo studies rather reported an equal ratio of conductances implying that inhibitory currents do not balance excitatory currents because of their different driving forces. Those observations therefore question our theoretical understanding of activated cortical states. Does stable asynchronous dynamics exist beyond the balanced setting? We addressed this question by further exploring the collective dynamics emerging in recurrently connected networks of excitatory and inhibitory spiking units. We formally demonstrate that, upon the inclusion of a disinhibitory circuit, recurrently connected networks exhibit a spectrum of asynchronous dynamics ranging from sparse activity regimes up to the classically reported balanced state. We evidence the signature of this theoretical picture in the mice somato-sensory system in vivo. Further theoretical analysis suggests that moving along this spectrum enables neural networks to efficiently transmit and process sensory stimuli of various properties.

*Speaker

P35. High-density multi-unit recordings in cortex of fragile X mice

Ingrid Van Welie * [1], Jorg Scholvin [2], Edward Boyden [3], Mark Bear [1]

[1] Picower Institute for Learning and Memory, MIT – United States
[2] Media Lab, MIT – United States
[3] Media Lab, MIT – United States

Multi-unit recordings are increasingly being used in awake head-restrained animal models to study principles of neural coding and dynamics. Technological advances in silicon probe technology now allow for the recording of hundreds of neurons simultaneously. In a collaborative effort, we are implementing new ultra high-density silicon probes with flexible layouts designed for the oversampling of neurons (Scholvin et al. 2016) in the fragile X mice mouse model. Fragile X mice are hypothesized to display hyperexcitability and hypersensitivity to sensory stimuli in many cortical areas (Contractor et al. 2015). Experimental design of our silicon probes and recording setup allows for multiple high density probes to be inserted in multiple cortical areas. Here we will present preliminary data from recordings across cortical areas of fragile X mice in response to sensory stimuli using silicon probes with hundreds of channels.

*Speaker

P36.Task-dependent tuning of the mouse olfactory bulb by glomerular interneurons

Izumi Fukunaga * [1,2], Andreas T. Schaefer1 [1,3]

[1]Neurophysiology of Behaviour Laboratory, The Francis Crick Institute;
[2]
Sensory and Behavioural Neuroscience Unit, Okinawa Institute of Science and Technology Graduate University;
[3]Department of Neuroscience, Physiology and Pharmacology, University College London

The ability of sensory systems to dynamically tune their stimulus representations may be a powerful mechanism by which the brain copes with ever-changing challenges posed by animals' own intent, as well as the environment.

We investigated, using the mouse olfactory system, how distinct behavioural demands shape the olfactory responses of neurons in the earliest stages of sensory processing in the brain. Head-fixed mice were trained to perform coarse and fine olfactory discriminations, which have opposing demands on optimal stimulus representations. With simultaneous two-photon imaging, we found that inhibitory neurons in the glomerular layer show opposite changes as mice acquire the two tasks: a general decrease in odour responses during coarse discrimination learning, while an increase as mice learnt to discriminate between similar odours. By training mice to switch between the tasks on a short timescale, we demonstrate that the change is reproducible and flexible. The change may occur so as to set the sparseness of the olfactory bulb output that suits the behavioural task at hand. Pharmacological inactivation of the ipsilateral olfactory cortex blocked the task-dependent change in sparseness, suggesting the cortical feedback contributes to the task-dependent change.

Altogether, our results suggest that behavioural demands tune the earliest olfactory processing in the brain, where sensory representations sparsen and decorrelate flexibly to suit the behaviour.

*Speaker