Neural networks - from brains to machines and vice versa

11th - 12th October 2018 - Institut Pasteur, Paris, France

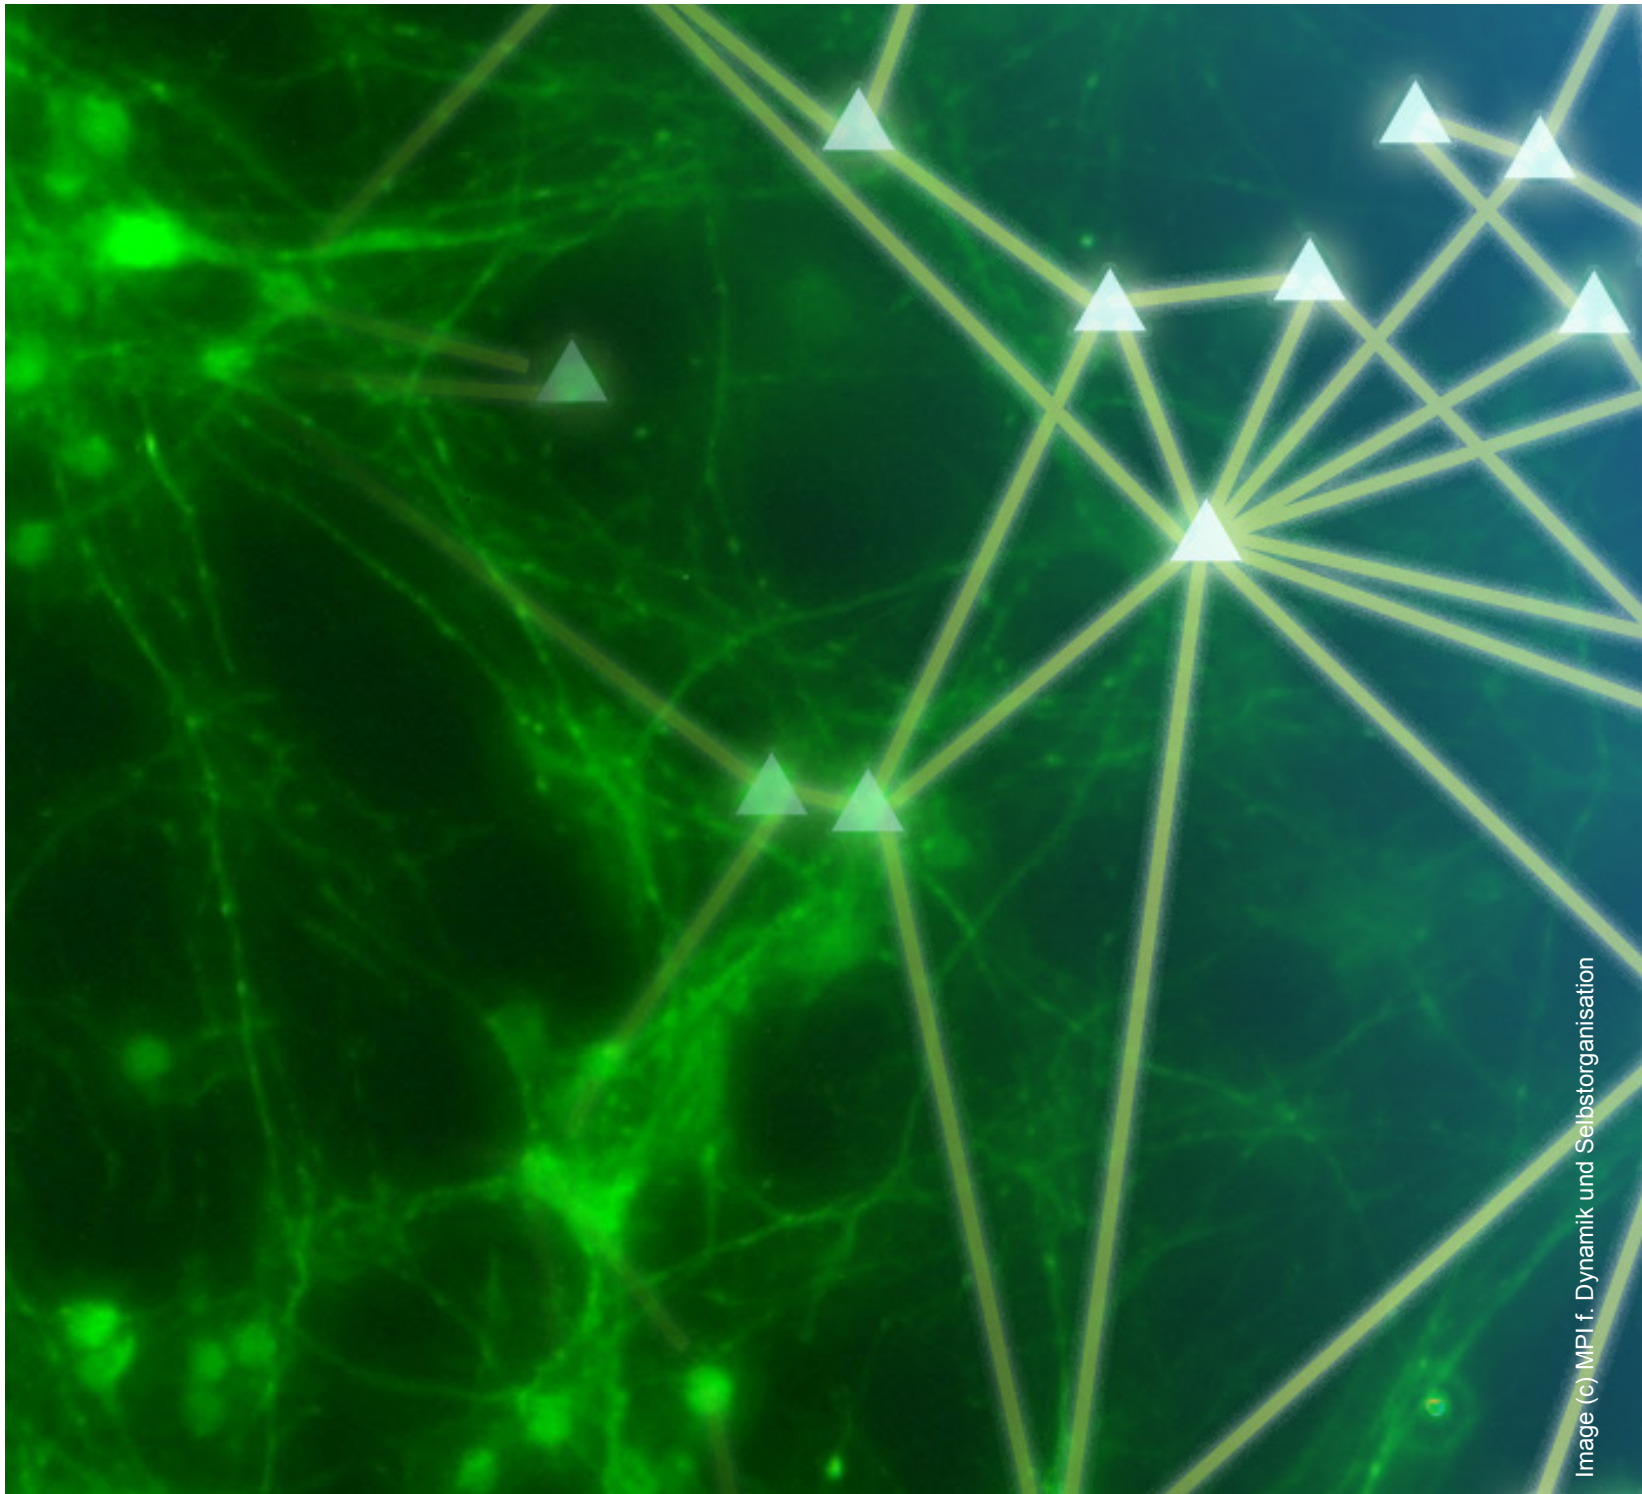Image (c) MPI f. Dynamik und Selbstorganisation

Qbio

INCEPTION

Institut Pasteur

Welcome!

Continuous advances in machine learning have enabled groundbreaking progress in diverse fields such as computer vision or strategy games by making use of artificial neural networks. In parallel, novel technologies for recording and manipulating biological neural networks allow us to probe the behavioural function of neuronal activity in unprecedented detail. By bringing together leading experts in the fields of biology, applied mathematics and physics, this symposium aims to bridge our current understanding of how biological and artificial neural networks operate.

David DiGregorio, Institut Pasteur
Jean-Baptiste Masson, Institut Pasteur
Rémi Monasson, École Normale Supérieure
Christoph Schmidt-Hieber, Institut Pasteur
Christophe Zimmer, Institut Pasteur

CÔTÉ 25

ℹ Accueil Informations	59 Henri DARRE	66 Jacques MONOD
51 Musée	60 Amédée BORREL	67 Elie METCHNIKOFF
52 Grande galerie	61 Camille GUERIN	68 Auguste FERNBACH
53 Emile ROUX	62 et 63 Albert CALMETTE	69 The 25
54 et 55 Edmond NOCARD	64 Alphonse LAVERAN	P Parking sous-sol
56 à 58 Locaux annexes	65 Charles NICOLLE	♿ Stationnement handicapés
DAE Défibrillateur		

Access only after registration & reply to follow-up email.
Bring your photo ID please.
Lectures will take place in the Amphitéâtre Monod, Building 66, basement
Poster sessions, coffee and lunch breaks will take place in "The 25", Building 69, 1st floor

Schedule

Thursday 11th October

Physics of biological neural networks (Chair: Jean-Baptiste Masson)

09h00

David DiGregorio, Institut Pasteur
Introductory remarks

09h15

Rémi Monasson, Ecole Normale Supérieure
Integration and multiplexing of positional and contextual information by the hippocampal network

10h00

Jakob Macke, TU Munich and research center ceasar
Bridging the gap between mechanistic and statistical models of neural dynamics

10h45 Coffee break

11h15

Christian Vestergaard, Institut Pasteur
Discovery of computational motifs in experimental neural connectomes using nested randomized reference models

12h00

Elad Schneidmann, Weizmann Institute
Learning the code of large neural populations by random projections

12h45 Lunch break

Keynote Lecture (Chair: David DiGregorio)

14h00

Larry Abbott, Columbia University
Lessons for Machine Learning from a Plastic Recurrent Network in Electric Fish

15h00 Coffee break

Machine learning in neuroscience (Chair: Christophe Zimmer)

15h45

Caswell Barry, University College London
Probing the spatial memory system with deep networks

16h30

Srini Turaga, Janelia Research Campus
From biological neural networks to artificial neural networks, with deep learning

17h15

Marco Zorzi, University of Padova
Deep neural networks for modeling perception and cognition

18h00 Poster session & Apéritif

Friday 12th October

Measuring neural networks in real life (Chair: Christoph Schmidt-Hieber)

09h00

Eugenia Chiappe, Champalimaud Research
Visuomotor control of walking during exploration in flies

09h45

Tihana Jovanic, Janelia Research Campus
Mapping neural circuits underlying sensorimotor decisions and sequences in Drosophila

10h30 Coffee break

11h15

Kate Jeffery, University College London
How landmarks inform the sense of direction

12h00 Lunch break

Physics of artificial neural networks (Chair: Jean-Baptiste Masson)

13h30

Carey Priebe, Johns Hopkins University
Structure discovery and exploitation in networks

14h15

Naftali Tishby, Hebrew University
The Information Bottleneck theory of Deep Neural Networks: what do the layers represent?

Poster List

Functional diversity of dynamic synapses is a substrate for temporal learning
Alessandro Barri, Martin Wiechert, François Chabrol, David DiGregorio

Effects of adaptation and synaptic filtering on randomly connected neural networks
Manuel Beiran, Srdjan Ostojic

Coding with transient trajectories in recurrent neural networks
Giulio Bondanelli, Srdjan Ostojic

Dynamic modulation of hippocampal spatial coding resolution by local virtual 3d objects
Romain Bourboulou, Geoffrey Marti, Morgane Nouguier, François-Xavier Michon, David Robbe, Julie Koenig and Jérôme Epsztein

Inferring the collective dynamics of neuronal ensembles from single-trial spike trains using mechanistic models
Christian Donner, Manfred Opper, Josef Ladenbauer

Reverse-engineering recurrent neural networks via low-rank approximations
Alexis M. Dubreuil, Srdjan Ostojic

Epigenesis of the global neural workspace: from local hebbian to global reinforcement learning
Guillaume Dumas, Jean-Stanislas Denain, Valentin Villecroze, Camille Démarre, Gauthier Guinet, Paul Jacob, Anne-Sophie Migeon, Jean-Pierre Changeux

Separating intrinsic interactions from extrinsic correlations in a network of sensory neurons
U. Ferrari, S. Deny, M. Chalk, G. Tkacik, O. Marre, T. Mora

Decoding spatial behaviour from raw hippocampal activity using deep learning
Markus Frey, Sander Tanni, Christian F. Doeller & Caswell Barry

Mice can discriminate vibrotactile frequencies simultaneously applied to neighboring whiskers, despite large overlaps in s1 evoked activity patterns.
Pierre-Marie Garderes, Dan Ganea, Mathias Günther, Florent Haiss

Prosthesis and biocompatibility: a synthetic biology approach for better integration and functionality
iGEM Team

High-throughput closed-loop tracker with online behavior detection for *drosophila* larva
Kristina Klein, Lakshmi Narayan, Michael Winding, Christopher McRaven, Jean-Baptiste Masson, Marta Zlatic

Inferring and validating mechanistic models of neural microcircuits based on spike-train data
Josef Ladenbauer, Srdjan Ostojic

Learning in PFC & striatum through shaping of recurrent dynamics
Christian David Márton, Simon R. Schultz, Bruno B. Averbeck

A geometrical description of global dynamics in trained feedback networks
Francesca Mastrogiuseppe, Srdjan Ostojic

A quantitative framework to study sub-cellular mrna localization
Aubin Samacoits, Racha Chouaib, Adham Safieddine, Abdel-Meneem Traboulsi, Wei Ouyang, Remy Dubois, Christophe Zimmer, Marion Peter, Edouard Bertrand, Thomas Walter, Florian Mueller

An analytical theory for spiking neurons driven by colored noise
Friedrich Schuessler, Stefan Rotter

Diverse spatial expression patterns emerge from unified kinetics of transcriptional bursting
Benjamin Zoller, Shawn C. Little, Thomas Gregor

FUNCTIONAL DIVERSITY OF DYNAMIC SYNAPSES IS A SUBSTRATE FOR TEMPORAL LEARNING

Alessandro Barri[1,2], Martin Wiechert[3], François Chabrol[4], David DiGregorio[1,2]
[1] Unit of Dynamic Neuronal Imaging, Institut Pasteur, Paris, France
[2] CNRS UMR 3571 Genes, Synapses and Cognition, Paris, France
[3] Department of Physiology, Universität Bern, Bern, Switzerland
[4] The Sainsbury Wellcome Centre for Neural Circuits and Behaviour, University College London, London, UK

The cerebellar cortex (CC) is thought to contribute to temporal processing and learning in the sub-second range. The prevailing hypothesis as to how this is achieved, is by the generation of diverse neuronal firing patterns within the input layer (granule cells) of the CC, which provide a temporal basis for the output neurons, Purkinje cells, to generate any output pattern necessary for controlling behaviour. However, the mechanisms by which this is achieved are unknown. Here we use computational approaches to explore theoretically whether the diversity of short-term plasticity (STP) of synapses within the input layer of the CC can provide the necessary dynamics for subsecond learning. We created a simplified rate model of the CC in which we introduced a diversity of STP at mossy-fibre-GC synapses as observed experimentally (Chabrol et al. 2015). By using analytical calculations and numerical simulations, we show that i) transient GC activity elicited by STP alone is sufficient for driving the learning of eye-blink responses in Purkinje cells, and that ii) the learning performance depends crucially on the tuning of the synaptic properties to the statistics the mossy-fibre inputs received by the CC. We further show examples of how STP in the cerebellar input layer is sufficient for enabling classical and hypothesized cerebellar computations such as the cancellation of self-generated sensory responses and the learning of prior distributions of time intervals. Our findings underscore the generality of STP as a mechanism underlying computation in the time domain.

EFFECTS OF ADAPTATION AND SYNAPTIC FILTERING ON RANDOMLY CONNECTED NEURAL NETWORKS

Manuel Beiran[1], Srdjan Ostojic[1]
[1]   Group for Neural Theory, LNC[2]. INSERM U960, ENS, PSL University, Paris, France

Cortical networks extract and represent information from the spatial and temporal structure of complex input sequences. These computations are built upon the time-dependent properties of single neurons and their inherent dynamics at the network level. Nevertheless, how slow processes at the single neuron level, such as synaptic filtering of the inputs or spike-frequency adaptation, impact the recurrent dynamics at the network level remains vastly unexplored.

To address this question, we investigate the dynamical regimes of large networks of randomly connected excitatory and inhibitory units, with the novelty that each rate unit is described by a two-variable system to account for adaptation or synaptic filtering. This additional degree of freedom modifies the network dynamics. In particular, adaptation reduces the overall mean firing rate of the units, increases the stability region of the system and, most surprisingly, changes its bifurcation structure. Beyond a critical synaptic strength, such a network in the inhibition-dominated regime undergoes a Hopf bifurcation to a state where all neurons show sustained heterogeneous oscillations at a finite frequency. In contrast, a network with synaptic filtering shows a transition to fluctuating, heterogeneous activity, and the stationary firing rate and the stability regimes do not vary with the filtering timescale.

Furthermore, we apply dynamical mean field theory (DMFT) to study the temporal statistics of the systems beyond the bifurcation. We find that the synaptic filtering time constant is proportional to the timescale of the network activity. Instead, the macroscopic dynamics of a network with adaptation do not inherit the large timescale present at the single neuron level.

CODING WITH TRANSIENT TRAJECTORIES IN RECURRENT NEURAL NETWORKS

Giulio Bondanelli [1], Srdjan Ostojic [1]
[1] Laboratoire de Neurosciences Cognitives, Ecole Normale Superieure, Paris, France

Across sensory systems, complex spatio-temporal patterns of neural activity arise in response to even simple stimuli. A sudden stimulus presentation and removal evoke strong transient neural trajectories that dynamically evolve in time and across neurons. These transient responses share remarkable similarities across sensory modalities, suggesting the possibility that they are generated by a common mechanism. At the population level, transient trajectories have been shown to be particularly informative about stimulus identity and may form the basis for computations based on transient dynamics.

The mechanisms generating strong transient responses have so far not been fully elucidated. Here we examine a network mechanism that generates transient responses through recurrent interactions. Focusing on linear recurrent networks, we determine the conditions for the existence of strong transients and examine their computational properties.

We identify a general criterion that allows us to distinguish two sharply separated dynamical regimes depending on properties of the connectivity matrix: a regime in which all inputs lead to decaying transients, and a regime in which specific inputs elicit strong transients. From a geometrical analysis of the population activity patterns in the neuronal state space, we derive the optimal set of stimuli driving strong transient responses in the network. We show that each stimulus in this set is mapped into a specific output state during the transient dynamics and can therefore be decoded from the two-dimensional transient trajectory it evokes. The number of the optimal inputs defines the capacity of the network for stimulus discrimination during transient dynamics and can be estimated for any network topology. We finally exploit our approach to build the minimal connectivity that implements specific trajectories which map a specific input into a specific output and show that such connectivity can encode an extensive number of trajectories.

DYNAMIC MODULATION OF HIPPOCAMPAL SPATIAL CODING RESOLUTION BY LOCAL VIRTUAL 3D OBJECTS

Romain Bourboulou [1,2+], Geoffrey Marti [1,2+], Morgane Nouguier [1,2], François-Xavier Michon [1,2], David Robbe [1,2], Julie Koenig [1,2§] and Jérôme Epsztein [1,2§]
[1] INMED, INSERM U901, Marseille – France
[2] UMRS 901, Aix-Marseille University, Marseille – France
[+] equally contributing authors   [§] equally contributing senior author

Animals can flexibly navigate their environment. In mammals such as rodents and humans, this ability is thought to rely on an internal cognitive map. When animals move in their environment, hippocampal place cells fire in specific locations (their place fields) and this spatial tuning is believed to provide a neuronal substrate to this cognitive map. Explored space can be internally mapped at different spatial resolutions. Nevertheless, it is still unclear if spatial coding resolution can be dynamically tuned within and between environments. In this study, we took advantage of virtual reality to selectively manipulate local visual information. Mice performed back and forth in virtual linear tracks which were differentiated by the presence or absence of virtual 3D objects. Once behavioural performance was stable, we recorded the activity of putative pyramidal cells in area CA1 of the hippocampus. Objects improved spatial coding resolution globally with a higher proportion of place cells, smaller place fields, increased spatial selectivity and stability. Spatial coding resolution was notably enhanced locally near objects and could be rapidly tuned by their manipulations. In the presence of objects, place cells also displayed improved theta phase precession and theta timescale spike coordination. These results suggest that local visual cues can rapidly tune the resolution of the hippocampal mapping system within and between environments. Such type of dynamical coding could optimize the cost of storing and using spatial information for efficient navigation.

INFERRING THE COLLECTIVE DYNAMICS OF NEURONAL ENSEMBLES FROM SINGLE-TRIAL SPIKE TRAINS USING MECHANISTIC MODELS

Christian Donner[1,2], Manfred Opper[1,2], Josef Ladenbauer[2,3]
[1]   Artificial Intelligence Group, Technische Universität Berlin, Germany
[2]   Bernstein Center for Computational Neuroscience Berlin, Germany
[3]   Laboratoire de Neurosciences Cognitives et Computationnelles, École Normale Supérieure Paris, France

Spike-train data recorded in vivo often exhibit rich dynamics over multiple timescales as well as considerable variability across neurons and trials (repetitions of identical experimental conditions). These data are frequently analyzed by fitting parametric phenomenological models that can accurately describe statistical dependencies in the data but possess only limited interpretive power because they typically do not incorporate prior biophysical constraints. Biophysically informed, mechanistic models, on the other hand, are rarely used for such quantitative analyses due to methodological challenges. Here we present a statistically principled approach based on an ensemble of doubly-stochastic integrate-and-fire neurons, taking into account basic biophysics. This model class comprises an idealized description for the dynamics of the neuronal membrane voltage in response to fast independent and slower shared input fluctuations given by Markov processes. We efficiently compute the likelihood of observed neuronal spike trains from a single trial by leveraging analytical methods for spiking neuron models combined with inference techniques for hidden Markov models, which allows us to (1) infer the model parameters, (2) compare different models from the class using established criteria, (3) reconstruct the hidden time-varying input statistics and (4) estimate instantaneous spike rates. Evaluations based on simulated ground truth data show that our method correctly identifies the dynamics of the shared input process (Markov jump versus Ornstein Uhlenbeck) and accurately estimates the model parameters. Furthermore, we successfully reconstruct the dynamics of current signals with additive noise applied to single neurons in vitro, using only their spike trains. Finally, we apply our approach to a neuronal ensemble recorded in vivo, for which we quantify how strongly individual neurons couple to the overall spiking dynamics. Our work demonstrates how statistical inference for a class of reasonably constrained, mechanistic models allows to dissect individual and ensemble neural dynamics on the basis of measured spike train data.

REVERSE-ENGINEERING RECURRENT NEURAL NETWORKS VIA LOW-RANK APPROXIMATIONS

Alexis M. Dubreuil[1], Srdjan Ostojic[1]
[1]   Group For Neural Theory, LNC, ENS, PSL Research University, Paris, France

Recurrent Neural Networks are artificial neural networks that can be trained to perform a variety of tasks. They have been proven to be useful for neurosciences as a way to find new solutions for the implementation of cognitive tasks by neural networks. Indeed they can be trained on tasks typically used in experimental neurosciences and yield functioning neural networks that are fully accessible, contrary to biological networks (see e.g. [1]). Here we take advantage of the accessibility of dynamics and connectivity of RNNs to propose a method that gives a concise mapping between cognitive mechanisms and network's structure. To do so we use recent work [2] that allows a detailed analytical description of the link between structure and dynamics in networks with low-rank recurrent connectivity matrix, i.e. that writes as outer products of a few column vectors $\vec{m}_k$ and line vectors $\vec{n}_k^T$. It has been shown that by examining the geometrical arrangements of the vectors $\vec{m}_k$'s and $\vec{n}_k$'s, together with vectors mediating network's inputs and outputs, one can understand how recurrent connectivity processes inputs and produces outputs to perform a task. Based on these results, here we adopt the following approach to analyze trained RNNs, i) we train a RNN to perform a given task, ii) find an approximate low-rank connectivity structure that preserves network's functionality, and iii) describe the cognitive mechanisms at stake by linking

network's structure and dynamics. We demonstrate the validity of this approach on two commonly used tasks in experimental neurosciences: a perceptual decision-making task (Random Dot Motion task) and a delayed discrimination task (Romo task). From this we can give a precise description, in the language of dynamical system theory, of the cognitive mechanisms involved in performing those tasks. We then discuss the application of this method to more elaborate cognitive tasks.

[1] Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503, 78–84 (2013).
[2] Mastrogiuseppe, F. & Ostojic, S. Linking connectivity, dynamics and computations in recurrent neural network. Neuron (2018).

EPIGENESIS OF THE GLOBAL NEURAL WORKSPACE: FROM LOCAL HEBBIAN TO GLOBAL REINFORCEMENT LEARNING

Guillaume Dumas[1,2], Jean-Stanislas Denain[3], Valentin Villecroze[3], Camille Démarre[3], Gauthier Guinet[3], Paul Jacob[3], Anne-Sophie Migeon[3], Jean-Pierre Changeux[1]

[1] Neuroscience Department, Institut Pasteur, Paris, France
[2] Center for Complex Systems & Brain Science, Florida Atlantic University, Boca Raton, USA
[3] Ecole Polytechnique, Palaiseau, France

The Global Neuronal Workspace hypothesis (GNW) posits a specific cortical architecture of specialised and distributed modules which perform low-level tasks at a non-conscious level, and under defined conditions, information processed by these modules becomes conscious by reaching a global network (the Global Neuronal Workspace) of distant pyramidal cells connected by long range axons. The message is then amplified and broadcasted to many brain regions through an ignition process which defines the subject's conscious access to the piece of information. GNW entails that the development of certain patterns of connectivity between territories of the global workspace is crucial in learning high-level abilities, such as social cognition, and effortful cognitive tasks (Dehaene, Kerszberg, and Changeux, 1998). We present a project of multi-scale model of the GNW through mechanisms of synaptic plasticity combining Hebbian and Reinforcement Learning (Izhikevich, 2007). The overall objective is to capture how synaptic epigenesis by selection/elimination of connections is shaping the GNW through the interaction of a genetic envelope —here embodied by the initial structure at micro-, meso-, and macro-levels—, the spontaneous neural dynamics, and the activity elicited by the interaction with the environment. We present preliminary findings about an elementary yet non-trivial task (solving the XOR problem) and delineate the future perspectives of the project.

SEPARATING INTRINSIC INTERACTIONS FROM EXTRINSIC CORRELATIONS IN A NETWORK OF SENSORY NEURONS

U. Ferrari[1], S. Deny[2], M. Chalk[1], G. Tkacik[3], O. Marre[1*], T. Mora[4*]

[1] Sorbonne Université, INSERM, CNRS, Institut de la Vision, 17 rue Moreau, 75012 Paris, France.
[2] Neural Dynamics and Computation Lab, Stanford University, California
[3] Institute of Science and Technology Austria, Klosterneuburg, Austria
[4] Laboratoire de physique statistique, CNRS, Sorbonne Université, Université Paris-Diderot and École normale supérieure (PSL), 24, rue Lhomond, 75005 Paris, France
* Equal contribution

Correlations in sensory neural networks have both extrinsic and intrinsic origins. Extrinsic or stimulus correlations arise from shared inputs to the network, and thus depend strongly on the stimulus ensemble. Intrinsic or noise correlations reflect biophysical mechanisms of interactions between neurons, which are expected to be robust to changes of the stimulus ensemble. Despite the importance of this distinction for understanding how sensory networks encode information collectively, no

method exists to reliably separate intrinsic interactions from extrinsic correlations in neural activity data, limiting our ability to build predictive models of the network response. In this work we introduce a general strategy to infer population models of interacting neurons that collectively encode stimulus information. The key to disentangling intrinsic from extrinsic correlations is to infer the couplings between neurons separately from the encoding model, and to combine the two using corrections calculated in a mean-field approximation. We demonstrate the effectiveness of this approach on retinal recordings. The same coupling network is inferred from responses to radically different stimulus ensembles, showing that these couplings indeed reflect stimulus-independent interactions between neurons. The inferred model predicts accurately the collective response of retinal ganglion cell populations as a function of the stimulus.

DECODING SPATIAL BEHAVIOUR FROM RAW HIPPOCAMPAL ACTIVITY USING DEEP LEARNING

Markus Frey[1], Sander Tanni[2], Christian F. Doeller[1] & Caswell Barry[2]
[1]   Kavli Institute for Systems Neuroscience, NTNU, Trondheim, Norway
[2]   Cell & Developmental Biology, UCL, London, United Kingdom

Spatially modulated cells in the hippocampus (CA1) and medial enthorinal cortex (mEC) provide reliable information required for successful navigation in an environment. A lower information limit of these structures can be approximated by assessing the accuracy with which self-location can be decoded from the underlying neural activity. Most traditional decoding algorithms use spike sorting methods for detecting action potentials which are subsequently used for spatial decoding. However, spike sorting potentially introduces selection biases and methods limited to action potentials inevitably discard information contained in the local field potential (LFP). Here we decode self-location from wide-band (2Hz – 6000Hz) electrophysiological signals from regions CA1 and mEC in freely moving rats, using state-of-the art deep learning algorithms. We demonstrate that our model is able to extract spatially relevant features from wide-band traces and outperforms established location-decoding methods, such as a Bayesian decoder, trained on manually sorted spikes.

MICE CAN DISCRIMINATE VIBROTACTILE FREQUENCIES SIMULTANEOUSLY APPLIED TO NEIGHBORING WHISKERS, DESPITE LARGE OVERLAPS IN S1 EVOKED ACTIVITY PATTERNS.

Pierre-Marie Garderes[1,2], Dan Ganea[1], Mathias Günther[1], Florent Haiss[1,2]
[1]   IZKF Aachen, Medical Faculty of the RWTH Aachen University, Aachen, Germany
[2]   Unit of Neural Circuit Dynamics and Decision Making, Institut Pasteur, Paris, France

Both humans and rodents can compare two simultaneously applied vibrotactile stimuli that are spatially separated. However, the ability to discriminate two proximate stimuli may be challenging as neighbouring cortical representations have a strong overlap in the upper layers of neocortex. Here, we developed a two-alternative forced choice (2-AFC) vibrotactile discrimination task using the simultaneous stimulation of two neighbouring whiskers in head-fixed mice. Stimulation of each whisker (C1 and D1) was associated with a specific-drinking spout which delivered a reward. Mice were remarkably successful in discriminating whisker frequencies, and their performance was above chance up to the smallest frequency difference tested. At the same time, neuronal activity in upper layers of primary somatosensory cortex (S1) was monitored using two-photon calcium imaging (jRGECO1a). At the single cell resolution, sample populations from layer 2/3 of the mouse barrel cortex showed sparse, and highly diverse evoked activity. Evoked responses were distributed and intermingled above the two corresponding barrels. We mapped the selectivity of the neurons to single whisker deflections and the modulatory effects from multiple whisker interactions. A subset of

neurons that were only modulated by simultaneous deflection of the two whiskers, with suppressive or additive effects, may form the substrate for the capability of the animal to discriminate simultaneous neighbouring inputs. Using simple decoding strategies on several features of neuronal encoding, including event intensity, event timing and population coupling, we could build neurometric curves in the context of either choice or disengagement from the animal and compare them to psychometric curves generated from animal behaviour. We report little modulation of neuronal patterns between correct and error responses yet large modulations of activity when animals were disengaged. This discrepancy hints a key role of S1 in shaping the sensory percept yet discards it as a direct substrate for decision-making.

PROSTHESIS AND BIOCOMPATIBILITY: A SYNTHETIC BIOLOGY APPROACH FOR BETTER INTEGRATION AND FUNCTIONALITY

iGEM Team: Etienne Axelos[1], Alice Dejoux[3], Claire Dumont[1], Antoine Ehret[4], Andreas Fyrillas[4], Aymen Ghaneme[5], Léa Guerassimoff[3], Samuel Jaoui[3], Manon Madelénat[6], Jonathan Naccache[6], Eléa Paillarès[6], Emma Picaud-Lucet[2], Sarah Porte[7], Ellyn Redheuil[4], Charlotte Richard[8], Gabriela Sachet[6], Thomas Starck[8], Florian Thomas[1], Kelly Trang[2], Anna Segu Cristina[9], Serena Petracchini[9], Guillian Graves[1], and Deshmukh N. Gopaul[9].

[1]   ENSCI-Les Ateliers, 75011 Paris, France
[2]   Faculté Jean Monnet, 92330 Sceaux, France
[3]   Chimie ParisTech, 75005 Paris, France
[4]   ESPCI Paris, 75005 Paris, France
[5]   Centrale-Supélec, 91190 Gif-sur-Yvette, France
[6]   Sup'biotech Paris, 94800 Villejuif, France
[7]   Ecole Polytechnique, 91128 Palaiseau, France
[8]   Sorbonne Université, 75006 Paris, France
[9]   Institut Pasteur, 75724 Paris, France

In the future, a due consideration will be given to people with disabilities. Amputees or people with an implant will have better access to high-performance bionic prostheses that will be connected to the amputated limb by neural pathways. Presently, some prostheses already allow amputees to recover partial autonomy and perform simple actions. However, these technologies do not allow a direct connection between the nerve and the prosthesis. In addition, a major health risk linked to prosthesis and implants is the colonization by communities of bacteria, called biofilms, which can sometimes become pathogenic.

To tackle these issues, we imagined NeuronArch, a project that aims at improving the connection between the control mechanism of the prosthesis and the nerves, using an engineered biofilm as an interface.

To do so, our interface produces substances called neurotrophins, allowing a directed and controlled growth of nerves. Moreover, it allows to pass the information towards the prosthesis thanks to enhanced conductive properties. Finally, in order to prevent the formation of pathogenic biofilms and increase the robustness of the system, our interface is able to fight against infections by perturbing the quorum sensing of *S. aureus*, responsible for the majority of implant-related infections.

HIGH-THROUGHPUT CLOSED-LOOP TRACKER WITH ONLINE BEHAVIOR DETECTION FOR *DROSOPHILA* LARVA

Kristina Klein[1, 2], Lakshmi Narayan[1], Michael Winding[1], Christopher McRaven[1], Jean-Baptiste Masson[1, 3], Marta Zlatic[1, 2]

[1]   Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, VA, USA
[2]   Department of Zoology, University of Cambridge, Cambridge, United Kingdom
[3]   Decision and Bayesian Computation, Pasteur Institute, Paris, France

Investigation of the neural mechanisms underlying learning and behavior often requires the ability to deliver stimuli to an animal directly in response to its actions. This, in turn, requires fast and reliable detection of the animal's behavior in real time. The *Drosophila* larva is an excellent model organism to study neural circuits, because genetic tools are available to selectively manipulate neural activity *in vivo*, using stimuli such as light (optogenetics) or heat (thermogenetics). Furthermore, larvae exhibit a variety of behaviors when placed on a two-dimensional surface, such as forward crawling, bending, backward crawling and rolling. We have developed a high-throughput closed-loop tracker for *Drosophila* larva with online behavior detection embedded in an upgraded multi-worm tracker software. The set-up is suited for selective optogenetic and thermogenetic stimulation of up to 16 larvae simultaneously live in response to detected behavior. Contour, head and tail of each larva are extracted from a high-resolution camera image at 20Hz using a field programmable gate array (FPGA). The data is then sent to the host computer, where features characterizing body position, velocity and direction of movement are extracted in real time and behavior is detected using a neural-network based approach. Online behavior detection for each larva is used to calculate an individual stimulus output based on a customizable protocol specified by the user. Optogenetic stimulation is obtained by projecting red light at selected larvae using a digital micromirror device (DMD). Thermogenetic activation or inhibition of neurons is achieved by heating up the animals with infrared laser beams, which can be targeted at individual larvae using two-axis galvanometers. We propose that this high-throughput tracker will facilitate the study of neural circuits underlying learning and decision-making.

INFERRING AND VALIDATING MECHANISTIC MODELS OF NEURAL MICROCIRCUITS BASED ON SPIKE-TRAIN DATA

Josef Ladenbauer[1], Srdjan Ostojic[1]

[1]   Laboratoire de Neurosciences Cognitives et Computationnelles, École Normale Supérieure Paris, France

The interpretation of neuronal spike train recordings often relies on abstract statistical models that allow for principled parameter estimation and model selection but provide only limited insights into underlying microcircuits. In contrast, mechanistic models are useful to interpret microcircuit dynamics, but are rarely quantitatively matched to experimental data due to methodological challenges. Here we present analytical methods to efficiently fit spiking circuit models to single-trial spike trains. Using the maximal-likelihood approach, we statistically infer the mean and variance of hidden inputs, neuronal adaptation properties and connectivity for coupled integrate-and-fire neurons. Evaluations based on simulated data, and validations using ground truth recordings in vitro and in vivo demonstrated that parameter estimation is very accurate, even for highly sub-sampled networks. We finally apply our methods to recordings from cortical neurons of awake ferrets and reveal population-level equalization between hidden excitatory and inhibitory inputs. The methods introduced here enable a quantitative, mechanistic interpretation of recorded neuronal population activity.

LEARNING IN PFC & STRIATUM THROUGH SHAPING OF RECURRENT DYNAMICS

Christian David Márton [1,2], Simon R. Schultz [1], Bruno B. Averbeck [2]
[1]   Centre for Neurotechnology & Department of Bioengineering, Imperial College London, London, U.K.
[2]   Laboratory of Neuropsychology, Section on Learning & Decision Making, NIMH/NIH, Bethesda, MD, USA

The brain is a non-linear dynamical system and approaching neural processing using methods from dynamical systems is beginning to change our understanding of how areas like the prefrontal cortex carry out complex cognitive tasks. Recurrent networks have been used successfully to probe neural population dynamics in simple tasks (Gallego, J. A. et. al. (2017) Neuron 94:978-984; Mante, V. et. al. (2013) Nature 503:78-84; Chaisangmongkon, W. et. al. (2017) Neuron 3 (6):1504-1517). Once trained, artificial neural network representations can be compared to neural recordings in terms of dynamical portraits which can offer insight into how a task is being represented (and solved) in neural space. We trained a recurrent network on a complex learning task involving 8 movement-sequences composed of 3 movements each. Two macaque monkeys executed the same task while recordings were made from lateral prefrontal cortex (lPFC) and dorsal striatum (dSTR) (Seo, M. et. al. (2012) Neuron 74: 947-960). The trained network learned to represent these movement sequences as distinct, non-overlapping trajectories in neural activation space. Analyzing the dynamical landscape during task execution, we found that the task is represented using distinct fixed points for each task period. We also found that uncertainty during learning is expressed in the network's gradient manifold – as learning progresses, sequence-specific gradient manifolds become more distinct from each other and, hence, more likely to push the network into the correct sequence-specific trajectory. With learning, the distance between sequence-specific gradient minima increases and sequence representations move farther apart from each other; at the same time, latent representations become more compact with learning. We analyzed spike trains from lPFC and dSTR of two macaques engaged in the same task with linear (GPFA) and non-linear (PLRNN) latent variable methods. We found that the distance between sequence representations in the neural population space increased with learning in dSTR, but not lPFC. Sequence representations in both regions also became more compact with learning. Altogether, our findings suggest task learning happens by shaping gradient manifolds at the neural population level, and that there are differences in how malleable representations are across regions with lPFC representations being more fixed and better separated than dSTR representations.

A GEOMETRICAL DESCRIPTION OF GLOBAL DYNAMICS IN TRAINED FEEDBACK NETWORKS

Francesca Mastrogiuseppe[1-2], Srdjan Ostojic[1]
[1]   Laboratoire de Neurosciences Cognitives et Computationelles, INSERM U960 and
[2]   Laboratoire de Physique Statistique, CNRS UMR 8550, École Normale Supérieure, PSL University, Paris, France

Recurrent neural networks have been extensively studied in the context of neuroscience and machine learning due to their ability to implement complex computations. While substantial progress in designing effective learning algorithms has been achieved in the last years, a full understanding of trained recurrent networks is still lacking. Specifically, the mechanisms that allow computations to emerge from the underlying recurrent dynamics are largely unknown.

Here we focus on a simple, yet underexplored computational setup: a feedback architecture trained to associate a stationary output to a stationary input. By deriving an approximate mean-field description of the global network dynamics, we show that this task admits several classes of solutions, characterized by different stability properties. These classes of solutions differ in the geometrical arrangement of the readout with respect to the input vectors, defined in the high-dimensional space spanned by the network population.

We show that our theoretical approach can be used to understand how standard training techniques implement the input-output task in finite-size feedback networks. In particular, our approximate description captures the local and the global stability properties of the target solution, and predicts training performance.

A QUANTITATIVE FRAMEWORK TO STUDY SUB-CELLULAR MRNA LOCALIZATION

Aubin Samacoits[1,2], Racha Chouaib[3,4], Adham Safieddine[3,4], Abdel-Meneem Traboulsi[3,4], Wei Ouyang[1,2], Remy Dubois[5], Christophe Zimmer[1,2], Marion Peter[3,4], Edouard Bertrand[3,4]*, Thomas Walter[5,6,7]*, Florian Mueller[1,2]*

[1] Unité Imagerie et Modélisation, Institut Pasteur
[2] C3BI, USR 3756 IP CNRS – Paris, France
[3] Institut de Génétique Moléculaire de Montpellier, University of Montpellier, CNRS, Montpellier, France
[4] Equipe labellisée Ligue Nationale Contre le Cancer
[5] MINES ParisTech,
[6] Institut Curie, 75248 Paris Cedex, France
[7] INSERM, U900, 75248 Paris Cedex, France

We present an integrated analysis framework of smFISH data to study sub-cellular mRNA localization. Non-random sub-cellular mRNA localization is important for cellular organization and its misregulation is linked to an increasing number of diseases. New imaging technologies - especially single molecule FISH (smFISH) - allow to observe single mRNA molecules in their native cellular environment and provide unique quantitative spatial information - and with recent advances even at large scale. However, as of today, there is no rigorously validated method to analyze these data. Here, we describe and validate a simulation and analysis framework for smFISH images to study mRNA localization. From simulated realistic images, we designed a set of validated features describing mRNA localization patterns inside cells, allowing their classification by supervised and unsupervised learning. With this approach, we raise the quantitative description of localization patterns to a new level, allowing to detect new localization classes. We successfully analyzed representative experimental data, which revealed a surprising degree of localization heterogeneity at the single cell level. Importantly, this flexible framework can be easily extended to new localization patterns. In ongoing work, we now investigate how deep learning can used for accurate classification without the need of manual feature design.

AN ANALYTICAL THEORY FOR SPIKING NEURONS DRIVEN BY COLORED NOISE

Friedrich Schuessler[1], Stefan Rotter[2]
[1]   Network Biology Research Group & Faculty of Medicine, Technion, Haifa, Israel
[2]   Bernstein Center Freiburg & Faculty of Biology, University of Freiburg, Germany

Existing mean field theories of spiking neural networks largely neglect the second order statistics of spike trains. This approximation hampers the understanding of real network dynamics, which can indeed be strongly influenced by related phenomena, like slow fluctuations. We develop an analytical theory for the spike train auto-covariance of leaky integrate-and-fire neurons driven by colored noise. A recently developed expansion in terms of the eigenfunctions of the corresponding Fokker–Planck equation is extended to the case of colored noise input, modeled by a Markov embedding. This theory constitutes a big step towards a self-consistent description of network dynamics including its covariance structure.

DIVERSE SPATIAL EXPRESSION PATTERNS EMERGE FROM UNIFIED KINETICS OF TRANSCRIPTIONAL BURSTING

Benjamin Zoller[1], Shawn C. Little[2], Thomas Gregor[1,3]

[1]   Joseph Henry Laboratories of Physics and the Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, USA

[2]   Department of Molecular Biology and Howard Hughes Medical Institute, Princeton University, Princeton, USA

[3]   Department of Developmental and Stem Cell Biology, Institut Pasteur, Paris, France

How transcriptional bursting relates to gene regulation is a central question that has persisted for more than a decade. Here, we measure nascent transcriptional activity in early Drosophila embryos and characterize the variability in absolute activity levels across expression boundaries. We demonstrate that boundary formation follows a common transcription principle: a single control parameter determines the distribution of transcriptional activity, regardless of gene identity, boundary position, or enhancer-promoter architecture. We infer the underlying bursting kinetics and identify the key regulatory parameter as the fraction of time a gene is in a transcriptionally active state. Unexpectedly, both the rate of polymerase initiation and the switching rates are tightly constrained across all expression levels, predicting synchronous patterning outcomes at all positions in the embryo. These results point to a shared simplicity underlying the apparently complex transcriptional processes of early embryonic patterning and indicate a path to general rules in transcriptional regulation.