The Early Visual System: comprehensive and data-driven modeling

May 19-20, 2016

Organizers: Yves Frégnac (UNIC, CNRS), Olivier Marre (IDV), Cyril Monier (UNIC, CNRS), Antoine Chaillet (Centrale-Supélec) and Alain Destexhe (UNIC, CNRS).

Abstracts by alphabetical order:

Jose-Manuel Alonso (SUNY, NY, USA)

Functional organization of ON and OFF pathways in visual cortex

Visual processing in the brain is mediated by two major thalamo-cortical pathways that signal local light increments (ON) and decrements (OFF) in the visual scene. In carnivores and primates, ON and OFF channels remain segregated in thalamus and combine for first time in visual cortex. However, our work demonstrates that the ON-OFF cortical mixing is not complete: it is partial and unbalanced. OFF thalamic afferents make stronger connections and occupy more territory in primary visual cortex than ON afferents and OFF-dominated neurons largely outnumber ON-dominated neurons. Moreover, cortical responses to dark stimuli are stronger, faster, more linearly related to luminance contrast and have better spatial and temporal resolution than responses to light stimuli. Our recent work also demonstrates that ON and OFF afferents segregate in different cortical domains of primary visual cortex and that these domains run orthogonal to ocular dominance columns. Moreover, we find that the cortical organization of ON/OFF domains is closely related to the layout of visual cortical maps for different stimulus features (e.g. orientation/direction) and makes maps of visual space to become dark-centric. That is, lights rotate around darks in the cortical representation of visual space just as stars rotate around black holes in the Universe. We propose that this surprising dark-dominated/dark-centric organization of visual cortical topography is a consequence of a size distortion for lights that originates already at the photoreceptor. Taken together, these findings could explain why visual acuity is commonly measured with dark characters on light backgrounds and why visual resolution is lower for lights, as already noted by Leonardo da Vinci and Galileo Galilei when judging the size of light objects in paintings and the dark sky.

Jan Antolik (UNIC, CNRS, France)

Model constrained by visual hierarchy improves prediction of neural responses to natural scenes

Accurate estimation of neuronal receptive fields and its link to neural substrate is essential for understanding sensory processing in the early visual system. While previous work has incorporated known structural properties of the early visual system, such as lateral connectivity, or imposing simple-cell-like receptive field structure, no study has exploited the fact that nearby V1 neurons share common feed-forward input from thalamus and other upstream cortical neurons. We introduce a new method for estimating receptive fields simultaneously for a population of V1 neurons, using a model-based analysis incorporating knowledge of the feed-forward visual hierarchy. We assume that a population of V1 neurons shares a common pool of thalamic inputs, and consists of two layers of simple and complex-like V1 neurons. When fit to recordings of a local population of mouse layer 2/3 V1 neurons, our model exhibits significant improvement of prediction power over the current state-of-the-art methods. We show that the responses of a large local population of V1 neurons with locally diverse receptive fields can be described with surprisingly limited number of thalamic inputs, consistent with recent experimental findings. Our structural model not only offers an improved functional characterization of V1 neurons, but also provides a framework for studying he relationship between connectivity and function in visual cortical areas.

James A. Bednar (School of Informatics, Edinburgh, Scotland)
Adaptation, homeostasis, synaptic plasticity, lateral interactions and feedback : Sufficient to explain V1?

What would it take to explain what we know about how the primary visual cortex (V1) works? I.e., is there a basic set of mechanisms that could account for the bulk of observations that have been made about the function and operation of this cortical area? We propose that a relatively small set of fundamental mechanisms, not specific to visual processing, can account for a large fraction of the functionally relevant observations that have been made about V1.

We will consider a series of computational models containing a network of interconnected V1 neurons driven by a model early visual pathway, with behavior governed by rules for homeostatic adaptation, synaptic plasticity, lateral interactions, and feedback. Given suitably rich input patterns, these self-organizing models develop neurons, maps, and intracortical connectivity specific for all of the major visual dimensions (retinotopy, orientation preference, ocular dominance, spatial frequency, color, and disparity). They can then account for behavior across a very wide range of spatial scales and temporal scales, for both naturalistic and artificial stimuli.

Fundamentally, these models encode the statistics of their input patterns (which themselves are a transformation of the statistics of the external environment) such that (a) the response properties of neurons across the cortical area have good coverage of the patterns typically experienced, (b) the lateral connectivity patterns encode patterns of coactivation of these neurons, (c) the responses of each neuron to a given input are modulated by these self-organized connectivity patterns, and thus by prior experience, and (d) the responses of each neuron can be modulated by intercortical connections according to task-specific demands.

At least as a first approximation, these mechanisms may be sufficient to account for the bulk of the observed behavior of V1 neurons, and establishing this type of model as a baseline would help drive discoveries of novel mechanisms or phenomena. Building a single model as a synthetic V1 now seems achievable; the immediate problems are practical (complexity of the parameter spaces, computational requirements, lack of data on input patterns and developmental context), rather than conceptual or theoretical.

Dan Butts (Univ. of Maryland, MD, USA)
Hidden cortical function manifesting as variability

Visual neuron responses in awake cortex can be quite variable across repeated presentations of the same stimulus. This is to say that a large fraction of cortical activity is “unexplainable", meaning that even a perfect model of stimulus processing would give limited insight into cortical function. Rather than noise, we believe such variability reflects the integration of stimulus-processing with other to-be-determined variables that are not experimentally controlled. In this talk, I will describe recent studies in my lab inferring the underlying sources of variability in visual cortex, using both local field potentials and simultaneous recordings from large numbers of neurons. In discovering what are these variables represent, and how are they combined with stimulus processing, we expect to discover broader roles throughout sensory cortices in combining sensory information with variables relevant to behavior and cognition.

Alain Destexhe (UNIC, CNRS, France)
Mesoscopic modeling of the suppressive effect of propagating waves in monkey visual cortex

Propagating waves are widely seen in the nervous system, and their presence was disputed in primary visual cortex (V1) of the monkey. Using a combination of voltage-sensitive dye (VSD) imaging in awake monkey V1 and model-based analysis, we recently showed that virtually every visual input is followed by a propagating wave (Muller et al., Nat Comm 2014). The wave was confined within V1, and was consistent and repeatable for a given input. More recently, we showed that two propagating waves interact in a suppressive fashion, and are always sublinear. Here, we show that a mean-field model can reproduce such features based on two mechanisms, first the preferential recruitment of inhibitory cells over excitatory cells by afferent activity, and second, the conductance-based shunting effect of the two waves onto one another. This suggests that the suppressive effect is a general feature of propagating activity, in agreement with the general suppressive effect seen in ather circumstances (Blair et al. J Neurosci 2003; Reynaud et al., J Neurosci 2012).

Valentin Dragoi (University of Texas-Houston, TX, USA)
Functional maps and signal integration in V1?

A fundamental property of cortical networks in the primary visual cortex (V1) of higher mammals is its functional organization, i.e., multiple stimulus dimensions are mapped such that neuronal response properties vary systematically across the cortical surface. Whether functional maps play any role in visual information processing and the pooling of signals relevant for perception is poorly understood. I will present recent results showing that optogenetic stimulation of distinct populations of excitatory neurons in V1 of macaque monkey can enhance the detection of an oriented stimulus when the stimulated population is tuned to the stimulus orientation. In contrast, activating populations of neurons untuned to the stimulus elicited a large increase in neuronal firing rates, but did not impact perception. By examining how optical stimulation influences the information encoded in population activity, we found that the light-induced improvement in behavioral performance was accompanied by a reduction in noise correlations and an increase in the population signal-to-noise ratio. These results demonstrate that causal manipulation of the responses of an informative population of excitatory neurons in V1 can bias the animal?s perceptual reports. We conclude that the functional architecture of V1 constrains the pooling of neural signals that influence perceptual decisions.

Alexander S. Ecker (Centre for Integrative Neuroscience, Tübingen, Germany)
What's the signal in the noise?

Responses of cortical neurons are highly variable. Even repeated presentations of the same visual stimulus never elicit the same spike train. Identifying the origins of this variability remains a challenge. There is increasing evidence that it is not just noise arising from stochastic features of neuronal architecture, but at least partly represents meaningful top-down signals. One of the most prominent examples of such top-down modulation in the visual system is covert attention. I will present both theoretical and experimental results showing that trial-to-trial fluctuations of attentional state contribute significantly to response variability in primary visual cortex of awake, behaving monkeys. I will argue that much can be learned about information processing in the brain by using latent variable models of neuronal activity to help us identify and account for cognitive variables and make sense of single-trial neural population data.

Sami El-Boustani (EPFL, Lausanne, Switzerland)
Functionally heterogeneous synapses shape cell-wide plasticity of visual cortex neurons in vivo

Neuronal circuits in the developing brain are subject to dramatic changes driven by sensory inputs, causing individual cells to adapt their responses to changing levels of drive and maintain a stable level of activity. Cell-wide homeostatic plasticity was initially reported as a global mechanism for stabilizing the output firing rate of a cell by uniformly scaling up or down the effective strength of all its inputs. More recent experimental evidence in vitro has suggested the existence of local homeostatic mechanisms acting at the level of dendritic stretches or even at single synapses. However, the existence and nature of local homeostatic plasticity in vivo and its implications for the coherent reorganization of single cell responses remains unexplored. Using a combination of two-photon imaging and optogenetic manipulation, we were able to induce receptive field plasticity in single-neurons of awake mice and study synaptic reorganization along dendritic branches. We found that homeostatic interactions take place around potentiated spines in stretches of dendrites and that this heterosynaptic plasticity involves AMPA receptors removal through redistribution of the immediate early gene Arc. This work, in line with previous theoretical predictions, demonstrates how local homeostatic plasticity can orchestrate cell-wide plasticity in dendrites displaying heterogeneous functional organization.

Adrienne Fairhall (Univ. of Washington, WA, USA)
Receptive fields and their discontents

Robbe Goris (NYU, NY, USA)
Origin and function of tuning diversity in Macaque visual cortex

Neurons in visual cortex are selective for stimulus orientation, but the precision of this selectivity varies across neurons. We used novel mixture stimuli and a computational "cascade" model to dissect the mechanistic origins of this diversity. Our analysis revealed that neuron–to–neuron variations in linear receptive fields and in non-linear sharpening mechanisms create tuning diversity. We also found that this diversity is matched to the needs of visual encoding. The orientation content found in natural scenes is diverse, and neurons with different selectivities are adapted to different stimulus configurations. A diverse population of neurons therefore provides better overall discrimination capabilities for natural images than any homogeneous population.

David Hansel (Paris Descartes, Paris, France)
Theory of orientation selectivity in rodent V1

Neurons in primary visual cortex (V1) exhibit orientation selectivity (OS). This is true for animals like cats or primates in which V1 has an orientation map (OM) as well as for those without OM (salt-and-pepper organization), e.g. rodents. Hubel and Wiesel proposed that OS arises from thalamic projections to the cortex that are spatially organized such that convergence of multiple untuned inputs generate orientation selective thalamic excitation to Layer 4 neurons. This scenario has been confirmed for the thalamic input in V1 of cats and monkeys, but up to now, it was not clear whether it applies to animals without OM. In a recent study, Lien and Scanziani (2013) combined intracellular recordings and optogenetics to characterize the orientation tuning and receptive field properties of the thalamo-cortical excitation in mouse, upon stimulation with drifting gratings. Combining analytical calculations and numerical simulations, I will demonstrate that the results of Lien and Scanziani can be accounted for without assuming any specific organization of the thalamic projections. In fact, a purely random connectivity from LGN to cortex is sufficient to quantitatively explain the experimental observations. I will then show that the weak information on the stimulus orientation conveyed by the resulting thalamic input can be extracted and amplified by the Layer 4 network in such a way that the orientation tuning of L4 neurons is comparable to what is reported in mouse at early stage of development.

Jason Kerr (Caesar Research Center, Bonn, Germany)
Imaging activity in the freely moving animal: From the eye to the cortex

Motivation underlies the performance of self-determined behavior and is fundamental to decision making, especially with regard to seeking food, mates, and avoiding peril. As many decision-making based behaviors in rodents involve a combination of head movements, eye movements, vestibular driven neuronal activity to guide behavior, studying the freely moving animal is paramount. In order to achieve this we have added to our head-mounted 2-photon microscope a custom-built miniaturized ocular videography system in combination allows the imaging of neuronal populations as well as precise tracking of head and eye movements. We show that movements of the two eyes in freely moving rats differ fundamentally from the precisely controlled eye movements used by other mammals to maintain continuous binocular fusion. We show that the observed eye movements serve to keep the visual fields of the two eyes continuously overlapping above the animal during free movement, but not continuously aligned. Overhead visual stimuli presented to rats freely exploring an open arena evoked an immediate shelter-seeking behavior, while the same stimuli were ineffective when presented beside the arena. We suggest that eye movements in freely moving rats provide constant overhead surveillance which would be of substantial evolutionary benefit for predator detection by minimizing blind spots**.**

Olivier Marre (Institut de la Vision, INSERM, France)
A flexible code in the retina

A classical view of the retina is that each ganglion cell type extract one feature from the visual scene, providing an individual feature map to the brain. Using large-scale recordings of the rodent retina, we show that a single cell type can multiplex radically different features, with individual neurons switching their encoding strategy depending on the visual context. When an object was moving randomly far away from their receptive fields, neurons of a single subtype responded synchronously to its motion. We found a non-linear, cascade model that predicts accurately their responses, and showed that they did not encode the absolute position of the object, but rather its speed and acceleration. However, if a second object moved inside their receptive fields, these neurons completely suppressed these responses to the far away motion, and encoded linearly the absolute position of the object inside their receptive field. Therefore, depending on the visual context, a single type of ganglion cell performs two very different kinds of computations: either a linear one, to code for position, or a non-linear one, coding for speed. Furthermore, these two computations were not linearly mixed in the neural response, but competed for encoding resources: facing multiple moving objects, a neuron suppressed its responses to objects moving far away from its receptive field to privilege objects inside its receptive field. This finding enriched the classical view of ganglion cell types as a feature map, and uncovers the flexibility of the retinal code when stimulated with complex stimuli.

Ken Miller (Columbia University, NY, USA)
Canonical computations of the cerebral cortex

Across multiple sensory cortical areas, strong nonlinearities are seen in the summation of responses to multiple stimuli. Responses to two stimuli in a neuron's receptive field (the sensory region in which appropriate stimuli can drive spike responses) typically sum sublinearly, with the response to the two stimuli presented simultaneously typically closer to the average than the sum of the responses to the two individual stimuli. However, when stimuli are weak, responses sum more linearly. Similarly, contextual stimuli, outside the receptive field, can suppress responses to strong stimuli in the receptive field, but more weakly suppress or facilitate responses to weaker receptive field stimuli. Many of these phenomena are often referred to as "normalization". I'll present a simple circuit mechanism that explains these and many other results. Individual neurons have supralinear input/output functions, leading the gain of neuronal responses to increase with response level. This drives a transition from (i) a weak-input regime in which neurons are weakly coupled, responses sum linearly or supralinearly, and contextual stimuli can facilitate, to (ii) a stronger-input regime in which neurons are strongly coupled and stabilized by inhibition against excitatory instability, responses sum sublinearly, and contextual stimuli suppress. I'll describe this mechanism and show how it can explain a variety of cortical behaviors, including those described above as well as suppression of correlated neural variability by stimuli and other behaviors as time permits.

Cyril Monier (UNIC, CNRS, France)
Data-driven models of the visual system and natural visual statistics, a stumbling block ?

An ultimate goal of visual neuroscience is to understand the neural encoding of natural scenes. A central hypothesis concerning sensory processing is that the neuronal circuits are specifically adapted to represent these stimuli efficiently. The usual way to study neural encoding of visual information is to compute a receptive field for each neuron, that is, a static transfer function between stimulus and neuronal response. Yet most of our knowledge of neuronal receptive fields has come from studies at the single cell level using only spiking firing rates and using simple (e.g., bars, gratings, sparse noise) or too complex (dense noise) artificial stimuli, that may fail to reveal the full nature of a neuron's actual response properties. As a consequence, present models have restricted their analysis mostly to the "feedforward" dependency of the single neuron activity to local stimulus features viewed through their classical discharge field.

Very few, but more recently, address the issue of contextual modulation of sensory responses to the spatial statistics location of stimulus itself, within, across or outside the studied RF (spatial adaptation, figure-ground and completion issues) and to the past history of the sensory flow (temporal adaptation). The difficulty of the task is that these contextual modulations occur at many spatial and temporal scales and at all levels of the visual hierarchy. Their main putative mechanisms are activity-dependent modulations of synaptic strength and neuronal excitability, but also context-dependent inputs, i.e. lateral and feedback inputs. The classic way of circumventing this intractable complexity in the receptive field framework is to use a stationary stimulation context, in order to compute a receptive field that can account for the responses only in this context. But during natural vision, stimulation statistics fluctuate continuously, and therefore adaptation will occurs all the time. The input-output relationship thus never reaches a stable state, contrary to what is implicitly assumed by the receptive field approach.

Thus, we propose to switch from the current conception of a static receptive field to a more dynamic view of population-based encoding. We will outline, step by step, how to build an unified adaptative data-driven multi-scale model of the primary visual cortex (V1) constraints with naturalistic visual stimulation with an interdisciplinary approaches in order to constrain and define a functional processing architecture whose structural complexity adapts to sensory input statistics dimensionality.

Bruno Olshausen (UC Berkeley, CA, USA)
Possible neural circuits for inferential computations in V1

Although the idea of perception as an inference problem goes back to Helmholtz, it is only recently that we have seen the emergence of neural models of perception that embrace this idea. Here I shall describe two neural models of perceptual inference that give us new ways of thinking about response properties in V1: 1) sparse coding, in which neurons compete to infer a compact description of image data in terms of a set of basis patterns, and 2) dynamic routing, in which separate neural populations representing form and motion interact to infer a stable representation of the world in the face of eye-movements that occur during fixation. In contrast to traditional computational models based on deductive processes such as feature detection and classification, these inferential computations rely upon recurrent computation in which information propagates both within and between levels of representation in a bi-directional manner. The inferential framework shifts us away from thinking of 'receptive fields' and 'tuning' of individual neurons, and instead toward how populations of neurons interact to perform collective computations.