# Linking neural dynamics and coding: correlations, synchrony, and information

Brent Doiron (University of Pittsburgh)
Kresimir Josic (University of Houston)
Nancy Kopell (Boston University)
Andre Longtin (University of Ottawa)
Alex Reyes (New York University)
Eric Shea-Brown (University of Washington)

October 3, 2010 – October 8, 2010

## 1 Overview

Understanding the mechanisms by which the nervous system represents and processes information is a fundamental challenge for mathematical neuroscience. It has long been known that information is carried in the intensity of *individual* neurons' responses to stimuli. As a consequence, many mathematical tools have been developed to describe populations of statistically independent neurons. However, new experimental techniques show the prominence of correlations and synchrony in neural activity – and understanding whether and how these *collective* dynamics encode information has become a major challenge for mathematical neuroscience.

This question was the focus of our workshop. We brought together international experts working in network dynamics and network information theory to forge new connections between underlying biological mechanisms and their consequences. Thus, the week was spent seeking bridges among three mathematical disciplines: (1) dynamical systems, (2) statistical mechanics, and (3) probability and information theory. These three branches of the mathematical sciences coincide with three central sub-disciplines in theoretical neuroscience whose focus is the study of collective nervous system activity. The first two concern how correlations and synchrony develop through network interactions, and the third seeks to quantify their information-theoretic impact on the neural code:

1. **Dynamical systems and network oscillations:** Recurrently (feedforward-feedback) coupled networks of spiking neurons often show synchronous activity. Mathematical analysis has revealed the mechanisms by which asynchronous activity loses stability and synchronous population rhythms arise. These mechanisms – and the specific patterns of synchronous rhythms that emerge – depend on rich interactions between network structure, coupling type, and single-oscillator dynamics. Moreover, recent research has shown that rhythms with distinct frequencies appear to interact. Unraveling the dynamical mechanisms of such interactions poses a new set of challenges that are only beginning to be addressed. How synchronous patterns are modified, created, and destroyed when networks are driven by external stimuli (e.g., sensory inputs) is another essential question that is being addressed using these mathematical tools.

2. **Statistical mechanics of network correlations:** Correlations can develop due to overlapping input in purely feedforward networks with irregular, stochastic activity. This is of particular importance for *layered* network architectures ubiquitous in neuroscience, where the propagation and amplification of correlated activity has been studied in systems ranging from cultured neural circuits to intact brains. Here, mathematical analysis seeks to quantify how correlated activity – modeled via multivariate point processes – is transferred among layers. This is a critical challenge; while it is evident from neural recordings that weak correlations are often present and presumably play an important part in normal brain function, excessive correlations are associated with neurological diseases, such as Parkinson's disease and epilepsy. Other current challenges focus on higher-order (beyond pairwise) correlations, and on how these correlation patterns depend on the spatiotemporal structure of stimuli.

3. **Information theory of network coding:** Neuroscience observations from high density electrode arrays are becoming more prevalent, posing the challenge of interpreting data recorded simultaneously from approximately 100 spatial locations. At the same time, results from information theory show that even weak correlations and synchrony can have strong effects on stimulus coding. However, whether these effects improve or degrade coding depends on the spatiotemporal structure of the collective activity. The primary challenge is to develop a systematic framework that predicts the impact of correlations in specific cases, and generalizes to allow an intuitive understanding of the underlying mechanisms of information encoding and decoding.

The experimental neuroscientists attending the workshop have been selected, in part, because of their collaborations with theorists. Their input was essential in guiding our discussion.
The question that unites these three areas is:

> **What are the information-theoretic consequences of the correlation and synchrony patterns that arise through the dynamics of prototypical neural circuits?**

Below, we report on progress toward the answer that was covered at our meeting.

**Integration of graduate students and postdocs:** A number of graduate students and postdoctoral fellows participated in the meeting. Nearly all of these participants also gave talks during the meeting. It is important to note that, while many of these students came from either side of the mathematics/neuroscience divide, they had no trouble in communicating their ideas to the diverse audience attending the workshop. All talks contained non-trivial mathematics, but presented in a way understandable to the participating experimentalists (admittedly, a selected group). We also observed, that while some of the presented research made use of fairly sophisticated mathematical ideas, all of it was well motivated by questions pertinent to neuroscientists.

# 2 Mathematical and scientific content of the meeting

The lectures and discussions at our meeting fell into four main themes:

1. Linking oscillations and signal processing
2. Correlations in specific circuit architectures
3. Defining useful metrics for encoding and decoding collective network dynamics
4. Linking feedforward and recurrent mechanisms for collective network activity

We next give a brief discussion of each, together with selected abstracts contributed by participants after the meeting that summarize the thrust of their talks (in some cases these were also edited by the organizers).

## 2.1 Linking oscillations and signal processing

**There is increasing evidence that synchronous oscillatory activity is controlled by both stimulus characteristics and the specific context of the recordings, e.g. during sleep vs. tasks requiring attention. While the mechanisms that underly oscillatory activity in the brain are being uncovered, little is known about the impact that oscillations have on the processing of sensory inputs. An especially challenging, and fascinating, question is: how does the brain make use of coexistent, multi-frequency, interacting rhythms? Especially intriguing are opportunities to apply information theoretic metrics to network models on either side of the transition from asynchronous behavior to different patterns of synchronous oscillations, as one step toward linking oscillations and neural coding.**

**Jonathan Rubin:** Rhythms in central pattern generators

Central pattern generators (CPGs) drive rhythmic movements such as respiration and locomotion. CPG outputs are rhythmic and repetitive, featuring multiple phases of activity with abrupt transitions between phases. Many different sets of intrinsic dynamics and connections between neurons can yield similar rhythms, yet these may involve different phase transition mechanisms. Although these transition mechanisms may not be discernible from direct examination of CPG output patterns, which mechanisms are present can have significant implications for CPG responses to external perturbations. In this talk, Rubin presented some theoretical analysis of this principle in an abstract, simplified rhythmic circuit. He subsequently illustrates the implications of transition mechanisms in several particular computational CPG models. In particular, analysis of transition mechanisms can be used to predict changes in respiratory phase durations in response to changes in particular external drives, to explain phase invariance of inspiration under hypercapnic conditions, and to explain differences between locomotor CPG rhythms with and without feedback from muscle afferents. More generally, transition mechanisms in CPG rhythms may play key roles in feedback control of CPG outputs.

**Ryan Canolty:** Role of patterns of oscillatory local field potential (LFP) phase coupling in regulating spiking activity

Hebb proposed that cell assemblies  anatomically-dispersed but functionally integrated groups of neurons – are critical for effective perception, cognition, and action. However, evidence for brain mechanisms that coordinate multiple coactive assemblies remains lacking. Neuronal oscillations have been suggested as one possible mechanism for cell assembly coordination. In *Role of patterns of oscillatory local field potential (LFP) phase coupling in regulating spiking activity*, Ryan Canolty presented both experimental evidence and an associated dynamical model that investigate this issue.

Prior studies have shown that spike timing depends upon local field potential (LFP) phase proximal to the cell body, but few studies have examined the dependence of spiking on distal LFP phases in other brain areas far from the neuron, or the influence of LFP-LFP phase coupling between distal areas on spiking. Canolty and colleagues investigated these interactions by recording LFPs and single unit activity using multiple microelectrode arrays in several brain areas, and then used a probabilistic multivariate phase distribution to model the dependence of spike timing on the full pattern of proximal LFP phases, distal LFP phases, and LFP-LFP phase coupling between electrodes.

The results show that spiking activity in single neurons and neuronal ensembles depends on dynamic patterns of oscillatory phase coupling between multiple brain areas, in addition to the effects of proximal LFP phase. Neurons that prefer similar patterns of phase coupling exhibit similar changes in spike rates, while neurons with different preferences show divergent responses  providing a basic mechanism to bind different neurons together into coordinated cell assemblies. Surprisingly, phase-coupling-based rate correlations are independent of inter-neuron distance. Phase-coupling preferences correlate with behavior and neural function, and remain stable over multiple days. These findings suggest that neuronal oscillations enable selective and dynamic control of distributed functional cell assemblies.

**Chris Pack:** Encoding of sensory stimuli by local field potentials in macaque visual cortex

Local field potentials (LFPs) are low-frequency fluctuations in electrical activity that are found throughout the brain. Because they correlate well with electroencephalography and fMRI BOLD signals, LFPs are critical to the study of brain function. In *Encoding of sensory stimuli by local field potentials in macaque visual cortex,* Christopher Pack presented experimental data from primate visual recordings suggesting a surprisingly strong link between the sensory tuning of low-frequency cortical LFPs and afferent inputs, with important implications for the interpretation of imaging studies and for models of cortical function.

**Mark Kramer:** Network oscillations in epilepsy, and beyond

During seizure, the aggregate voltage activity of neural populations often exhibits stereotypical rhythmic patterns, typically dominated by large amplitude voltage oscillations observable at the scalp or cortical surface. These rhythmic activities recorded from separate brain areas often exhibit correlations that also evolve in characteristic ways. In this *Network oscillations in epilepsy, and beyond*, Mark Kramer described correlation patterns observed in invasive voltage recordings from a population of human subjects with epilepsy. He showed that correlations increase at seizure onset and termination compared to pre-seizure intervals, suggesting the surprising result that macroscopic cortical areas decorrelate during the middle intervals of the seizure. Kramer also characterized other network properties during the seizure, including their coalescence and fragmentation. Finally, he applied these analyses to non-seizure recording intervals and to examine the common network structures that emerge.

**Leslie Kay:** Complementary functional and behavioral roles for olfactory beta and gamma oscillations

In *Complementary functional and behavioral roles for olfactory beta and gamma oscillations*, Leslie Kay presented a complementary view on the role of correlated dynamics on coding in olfactory discrimination tasks. Here, the correlations took the role of coordinated, rhythmic spiking across a neural population.

Moving beyond the anatomy or input wiring, Kay discussed a broad range of dynamic processes in the olfactory bulb, the first central and cortical stage of olfactory processing. She showed that olfactory bulb gamma oscillations (40-100 Hz oscillations of the local field potential), representing the precision of the underlying neural population, increase when rats learn to discriminate highly overlapping input patterns in a 2-alternative choice task. The mechanism for producing these oscillations is known to be the reciprocal dendrodendritic synapse between glutamatergic mitral cells and GABAergic granule cells.

Intriguingly, when rats were trained in a similar go/no-go task, gamma oscillations were not enhanced, and beta oscillations ( 20 Hz) instead predominated. These oscillations are not just a different frequency, but they rely on a different network. When centrifugal input to the olfactory bulb is ablated, gamma oscillations increase and beta oscillations disappear. However, the results are not quite as dichotomous as they might seem, because coherence patterns point to an underlying beta oscillation network in the distributed olfactory system in both tasks.

## 2.2 Correlations in specific circuit architectures

**Given that subtle differences in spatiotemporal correlations can have a major impact on encoded information, networks that have even small differences in their architectures may encode stimuli in substantially different ways. A mechanistic theory that connects patterns of neural correlation to network architectures needs to be based on a small family of prototypical circuits. But which should be chosen to best represent signal encoding in the brain? Interaction with experimental neuroscientists such as Alex Reyes (NYU) on cortical circuits and Leonard Maler (U. Ottawa) on sensory circuits that have evolved for different information processing tasks framed many of our discussions of these topics.**

**Valentin Dragoi, Adam Kohn and Andreas Tolias:** Correlated variability in laminar cortical circuits

Valentin Dragoi, Andreas Tolias and Adam Kohn, all researchers working with multi-electrode recordings in primates, gave a joint presentation about the impact of correlated variability in cortical circuits. This is a topic of much current interest. These participants are in the forefront of the field, and each has contributed fundamental results.

All participants agreed that the amount of information encoded by cortical circuits depends critically on the capacity of nearby neurons to exhibit correlations in their responses. Despite the fact that strong trial-by-trial correlated variability in response strength has been reported in many cortical areas, Andreas Tolias suggested that neuronal correlations may be much lower than previously thought. He started with the observation that many cortical areas are organized into functional columns, in which neurons are believed to be densely connected and share common input. Many numerous studies report a high degree of correlated variability between nearby cells. He described the work of his group on the development of chronically implanted multi-tetrode arrays offering unprecedented recording quality to re-examine this question in primary visual cortex of awake macaques. They found that even nearby neurons with similar orientation tuning show virtually no correlated variability. These findings suggest a refinement of current models of cortical microcircuit architecture and function: either adjacent neurons share only a few percent of their inputs or, alternatively, their activity is actively decorrelated.

In response Valentin Dragoi presented his work with laminar probes to revisit the issue of correlated variability in primary visual cortical (V1) circuits. Dragoi found that correlations between neurons depend strongly on local network context - whereas neurons in the input (granular) layer of V1 showed virtually no correlated variability, neurons in the output layers (supragranular and infragranular) exhibited strong response correlations. He showed how to use a linear decoder to demonstrate that, contrary to expectation that the output cortical layers would encode stimulus information most accurately, the input network encodes more information and offers superior discrimination performance compared to the output networks. He noted that laminar dependence of spike count correlations is consistent with recurrent models in which neurons in the middle (granular) layer receive intracortical inputs mainly from nearby cells, whereas neurons in superficial (supragranular) and deep (infragranular) layers receive inputs over larger cortical distances.

Similarly, Adam Kohn reviewed the mounting evidence that suggests that determining how neuronal populations encode information and perform computations will require understanding correlations between neurons, as well as the stimulus and behavioral conditions that modify them. His lab is taking advantage of the recent advent in recording techniques such as multielectrode arrays and two-photon imaging has made it easier to measure correlations, opening the door to detailed exploration of their properties and contributions to cortical processing. He noted a number of participants at the workshop reported discrepant findings, providing a confusing picture about the level and import of correlations in neuronal networks. He reviewed a selection of these studies and presented simulations to explore the influence of several experimental and physiological factors that affect the measurement of correlations. Differences in response strength, the time window over which spikes are counted, internal states, and spike sorting conventions can all dramatically affect measured correlations and systematically bias estimates. He concluded by offering guidelines for measuring and interpreting correlation data.

**Marlene Cohen:** A link between gains, correlations, and behavior

Attention allows observers to focus on a small subset of a cluttered scene and improves perception of attended locations or features. Both spatial and feature attention multiplicatively scale the firing rates of sensory neurons: typically, attending to a location that is close to a neurons spatial receptive field or to a feature that matches its stimulus preference increases sensory responses. We also showed previously that spatial attention tends to decrease correlations between the trial-to-trial fluctuations in the responses of nearby neurons (Cohen and Maunsell, 2009; see also Mitchell et al, 2009). To directly compare the effects of feature and spatial attention on neuronal populations, we recorded simultaneously from dozens of V4 neurons in both hemispheres while animals performed a change detection task in which we varied spatial and feature attention. We found that like spatial attention, feature attention modulates both firing rates and correlations. We found a strong inverse relationship between modulation of rate and correlation for both types of attention: when gains are increased, correlations decrease. While spatial attention increases the firing rates of most neurons, feature attention can either increase or decrease firing rates, depending on the similarity between a neurons tuning and the attended feature. There is an inverse relationship between rate and correlation modulation for cells whose gains decreased as well: feature attention increases the correlations between these neurons. Furthermore, on behavioral trials in which the animal made an error, the correlation structure looks like the correlation structure in the opposite feature and spatial attention condition, suggesting that the animal made an error because attention was misallocated. Together, our results suggest that a single mechanism accounts for the changes in firing rates and spike count correlation caused by any type of attention.

**Bruno Averbeck:** Dopamine, dynamics and information in the basal ganglia

Bruno Averbeck is a researcher at the NIH, and described a problem related to the work of other meeting participants, but of significant importance in medical research: Dopamine depletion in cortical-basal ganglia circuits in Parkinsons disease (PD) grossly disturbs movement and cognition. Classic models relate Parkinsonian dysfunction to changes in firing rates of basal ganglia neurons. Taking both inappropriate firing rates and other dynamics into account, and determining how changes in the properties of these neural circuits that occur during PD impact on information coding, is thus important.

Averbeck described *in vivo* network dynamics in the external globus pallidus (GPe) of rats before and after chronic dopamine depletion. He showed that dopamine depletion leads to decreases in the firing rates of GPe neurons and increases in synchronized network oscillations in the beta frequency (13-30 Hz) band. Using logistic regression models he showed the combined and separate impacts of these factors on network entropy, a measure of the upper bound of information coding capacity. Importantly, changes in these features in dopamine-depleted rats lead to a significant decrease in GPe network entropy. Changes in firing rates have the largest impact on entropy, with changes in synchrony also decreasing entropy at the network level. Changes in autocorrelations tended to off-set these effects as auto-correlations decreased entropy more in the control animals. Thus, it is possible that reduced information coding capacity within basal ganglia networks may contribute to the behavioral deficits accompanying PD.

Complementary to the GPe work his group also examined GPe, STN interactions by analyzing neurons recorded simultaneously from these nucleiin dopamine lesioned and healthy control rats. Both nuclei display a pronounced increase of beta frequency (20 Hz) oscillations in the lesioned state. Additionally, analyses of the information transfer between nuclei show that the transfer was significantly increased in the lesioned state. Furthermore, the temporal profile of the information transfer matches well the known neurochemistry of the nuclei, being inhibitory from the GPe to the STN and excitatory from the STN to the GPe and the dynamics of this interaction match well previously published estimates of the dynamics seen in Parkinsons patients.

Overall, these results showed that data analysis inspired by deep mathematical concepts can be used to provide evidence of specific changes in the functional connectivity between basal ganglia nuclei in the dopamine lesioned state.

**Jeremie Lefebvre:** Driven networks of ON and OFF cells with recurrent feedback

The origin of gamma oscillations in sensory networks continues to attract a lot of attention. Such oscillations are known to occur in the electrosensory system when stimuli have a large spatial extent. These past studies have modeled this phenomenon using a population of ON cells receiving spatio-temporal noise as its input, and with delayed feedback in its network topology. Jeremie Lefebvre presented results on how the presence of ON and OFF cells embedded in such a delayed feedback network influences the genesis of gamma oscillations. He illustrated the responses of neural populations to spatio-temporal forcing, mimicking those found in most sensory systems. ON pyramidal cells received sensory inputs directly, while OFF cells received a mirror image of the stimuli via an interneuron, inverting their response. The connectivity was determined solely by global inhibitory recurrent connections. Using a combination of neural field theory and numerical simulations, he showed that input-induced Andronov-Hopf bifurcations can occur; the stability of oscillations is determined by the spatial features of the input. He also showed how the network can double the frequency of an input in its firing activity. This is a consequence of rectification in the feedback network. He finally showed how adaptation can enhance gamma oscillations in such circuits.

**Alex Pouget and Jeff Beck:** Insights from a simple expression for linear Fisher information in a recurrently connected population of spiking neurons / Neural basis of perceptual basis

Alex Pouget and Jeff Beck presented joint talks. First, in *Insights from a simple expression for linear Fisher information in a recurrently connected population of spiking neurons*, they gave a simple expression for a lower bound of Fisher information for a network of recurrently connected spiking neurons which have been driven to a noise-perturbed steady state. This lower bound is called linear Fisher information, as it corresponds to the Fisher information thatcan be recovered by a locally optimal linear estimator. Unlike recent similar calculations, the approach used here includes the effects of non-linear gain functions and correlated input noise, and yields a surprisingly simple and intuitive expression that allows for substantial insight into the sources of information degradation across successive layers of a neural network. Here, this expression is used to (1) compute the optimal (i.e., information maximizing) firing rate of a neuron, (2) demonstrate why sharpening tuning curves by either thresholding or via the action of recurrent connectivity is generally a bad idea, (3) show how a single cortical expansion is sufficient to instantiate a redundant population code which can propagate across multiple cortical layers with minimal information loss, and (4) show that optimal recurrent connectivity strongly depends upon the covariance structure of the inputs to the network.

Next, these results found application in *Neural basis of perceptual basis*. The motivation was from cognitive neuroscience: extensive training on simple tasks results in large improvements in performance, a form of learning known as perceptual learning. Previous neural models have argued that perceptual learning is the result of sharpening and amplification of tuning curves in early visual areas. However, these models are at odds with the conclusions of psychophysical experiments manipulating external noise, which argue for improved decision making, presumably in later visual areas. Here, Pouget and Beck explore the possibility that perceptual learning for fine orientation discrimination is due to improved probabilistic inference in early visual areas. This mechanism captures both the changes in response properties observed in early visual areas and the changes in performance observed in psychophysical experiments. The modeling also suggests that sharpening and amplification of tuning curves may play only a minor role in improving performance, in comparison to the role played by the reshaping of inter-neuronal correlations.

**Maurice Chacron** Neural variability and contrast coding by correlations

Understanding how populations of neurons encode sensory information is of critical importance. Correlations between the activities of neurons are ubiquitous in the central nervous system and, although their implications for encoding and decoding of sensory information has been the subject of arduous debates, there is a general consensus that their effects can be significant. As such, there is great interest in understanding how correlated activity can be regulated. Recent experimental evidence has shown that correlated activity amongst pyramidal cells within the electrosensory lateral line lobe (ELL) of weakly electric fish can be regulated based on the behavioral context: these cells modulate their correlated activity depending on whether the fish is performing electrolocation or communication tasks without changing the mean firing rate of their response. Moreover, it was shown in the same study that the changes in correlated activity were correlated with changes in bursting dynamics. In this work we explore the role of intrinsic bursting dynamics on the correlated activity of ELL pyramidal neurons. We use a combination of mathematical modeling as well as in vivo and in vitro electrophysiology to show that bursting dynamics can significantly alter the ability of neuronal populations to be correlated by common input. In particular, our model predicts that the ratio of output to input correlations (i.e. the correlation susceptibility) is largely independent of stimulus amplitude when neurons are in the tonic firing model. In contrast, we find that the correlation susceptibility increases with stimulus amplitude when the neurons are in the bursting mode. We then performed in vivo and in vitro experiments to verify this prediction. Our results show that intrinsic dynamics have important consequences on correlated activity and have further revealed a potential coding mechanism for stimulus amplitude through correlated activity.

**Michael Graupner** Correlations in the auditory cortex during spontaneous activity

Spiking correlations between neurons have been found in many regions of the cortex and under multiple experimental conditions. Despite their importance consequences for neural population coding, the origin and the magnitude of such correlations remain a highly debated issue. Potential sources of correlations include shared presynaptic input. However, theoretical investigations have shown that shared inputs do not necessarily lead to correlations. Instead, active decorrelation occurs provided that the neurons are tightly coupled in a balanced configuration of excitation and inhibition. In support of these results, recent experiments measure virtually no correlations. However, those findings are in contrast to a large body of prevalent results suggesting strong correlations. We examine to which extent spontaneously active cortical networks meet the conditions of a balanced, decorrelated activity regime.

To investigate synaptic input, membrane potential and spike-output correlations between pairs of neurons, we perform simultaneous whole-cell recordings from pairs of pyramidal neurons in thalamocortical slices from young mice (P14-18).

We find correlated excitatory as well as inhibitory input to pairs of cortical cells during spontaneous activity. Interestingly, excitatory and inhibitory inputs are anticorrelated leading to cancellation of correlations at intermediate membrane potentials. We furthermore measure weak spike-count correlations between neurons ( 0.01). Together, our results show that nearby ( 100 m) cortical neurons receive correlated synaptic input.

However, spiking correlations are suppressed due to negative correlations between excitatory and inhibitory inputs. Our results suggest that cortical networks are structured to actively suppress correlations and thereby increase their information coding capacities.

## 3 Defining useful metrics for encoding and decoding collective network dynamics

**The standard metrics used to asses encoding of sensory information in spike trains are Fisher and Mutual Information. The former quantifies the accuracy with which sensory stimuli can be estimated from (stochastic) patterns of spikes, and the latter measures the reduction in uncertainty about a stimulus from observations of the response. These metrics can be made mathematically precise yet often assume system optimality, and are not necessarily motivated by the biophysical constraints present in the brain. They also require the specification of the neural 'response', a matter of much debate among experimentalists and theorists. Using mechanistic models of neural response will prompt a principled exploration of these areas, specifically of how correlations shape the neural code. An issue of special focus is models that address emerging large datasets from many simultaneously recorded neurons.**

**Andrea Barreiro:** When are microcircuits well-modeled by pairwise maximum entropy methods?

The conference theme was that collective activity is widespread in the nervous system and has important implications for functionality. When can we represent such activity by lower dimensional models, and how does our ability to do so depend on basic circuit properties such as input statistics, internal dynamics and network connectivity – such as the descriptions in wide use at the conference, where only pairwise spike correlations were considered? In *When are microcircuits well-modeled by pairwise maximum entropy methods?*, Andrea Barreiro took some first steps toward answering this question by studying the ability of maximum entropy models to characterize the spiking activity of networks modeled on retinal circuitry.

She first considered systems of N=3 spiking cells, driven by a common fluctuating input against independent background noise. She probed this circuit over a wide variety of operating regimes and input correlation structures, assessing the efficacy of the PME model by calculating the KL-divergence between the observed and PME distributions. Using a novel visualization method she showed that bimodal inputs generate spiking distributions that break the PME. Barreiro gave an analytical justification of these findings: in the small parameter describing the strength of common inputs to the circuit, $D_{KL}$ is at least an order (often more) smaller for unimodal vs. bimodal inputs. This persists for larger N.

Barreiro then constructed a biophysical model constrained by intracellular recordings of primate parasol RGCs. She exposed a triplet of such cells to stimuli at a wide variety of spatial and temporal scales. Even in the presence of highly correlated inputs, and significant cell-to-cell heterogeneity (induced by blocky spatial patterns of comparable size to the cell receptive fields), spiking outputs are well fit by the PME model. This surprising result explains previous experimental results, and leads to predictions for stimuli and RGC classes that will lead produce departures from PME responses. Preliminary results indicate that the feedforward structure of these circuits is highly significant in achieving the above results; introducing recurrence into this circuit can increase higher-order correlations by a factor of 20.

**Liam Paninski:** Coding and Computation by Neural Ensembles in the Primate Retina

The neural coding problem — deciding which stimuli will cause a given neuron to spike, and with what probability — is a fundamental question in systems neuroscience. We apply statistical modeling methods to analyze data recorded from a complete mosaic of macaque parasol retinal ganglion cells in a small region of visual space. We find that a surprisingly simple model with functional coupling between neurons captures both the stimulus dependence and the detailed spatiotemporal correlation structure of multi-neuronal responses; in addition, ongoing network activity in the retina accounts for a significant portion of the trial-to-trial variability in a neuron's response. We assess the significance of correlated spiking by performing optimal Bayesian decoding of the population spike responses. Finally, we discuss work in progress on the following questions: how much temporal precision is necessary to capture the neural code in the retina? How can we adapt our optimal decoding methods to estimate behaviorally relevant signals such as image velocity? How do we perceive stable images when the retina must contend with the constant motion due to small random eye movements? Finally, what can statistical spike-train analysis methods tell us about the underlying circuitry of the retina?

**Jean-Philippe Thivierge:** The Creative Nature of Neurons: How Heterogeneous Networks Provide a Rich Repertoire of Brain Activity

The brain is a creative organ it never responds to the same sounds and sights in exactly the same way twice. Jean-Philippe Thivierge's work on modeling brain responses across time sheds new light on the origins of this variability. Neurons of the brain are heterogeneous, meaning that they each possess slightly different characteristics that make them unique. According to his new theory, this property gives rise to a rich repertoire of possible brain states, and prevents the brain from getting ?stuck? in certain patterns of activity. Despite this wealth of possible states, neurons can also fine tune their interactions in order to reproduce rhythmic patterns beyond their time of presentation. He argued how this may explain the formation of short-term perceptual memories, as well as persistent rhythms of activity in neuropathological conditions. This presentation elicited a lot of questions about the interplay of fine tuning and network heterogeneity.

**Don Katz:** Modeling the impact of attention on coherent cortical ensembles: decreasing temporal coding variability by increasing noise

In interpreting neural activity, it is often assumed that information available in single-neuron responses is of primary importance: many theoretical models of population function take as their input highly pre-processed characterizations of single-neuron responses (i. e., response magnitudes, collapsed across trials and post-stimulus time), and many discussions of between-neuron correlations center on the concern that any overlap in the information available in each of two single-neuron responses reduces the information available in the population (i. e., "redundancy").

In *Modeling the impact of attention on coherent cortical ensembles: decreasing temporal coding variability by increasing noise*, Don Katz presented an alternative view. His group approaches the population coding of taste from a different angle, restricting our analysis to simultaneously-recorded ensembles of neurons and characterizing activity in single trials, without averaging across within- or between-trial timescales. The data is interpreted the information available in single neurons only in light of a primary population-level characterization, rather than vice-versa, and thus a clearer picture of the true dynamics of the system in action may be appreciated. Specifically, he observes cortical and amygdalar ensembles progressing through a sequence of coherent, nonlinear (attractor-like) firing-rate transitions; the sequences are reliable and stimulus-specific, but the timing of these transitions is highly variable from trial to trial–that is, much of the seeming "correlated noise" in the single-neuron responses reflect important aspects of the population dynamics. We argue that averaging single-neuron activity across time (and collapsing neurons that were collected non-simultaneously into single ensembles) obscures critical aspects of the population dynamics.

**Ila Fiete:** Beyond classical population coding for nearly exact estimation in the brain

The brain represents and transforms external variables to perform computations and achieve goals. Representation and transformation are inherently noisy when performed by neurons. One way to extract a less noisy estimate of the encoded variable is by averaging over large neural populations. Classical population codes, as seen in the sensory and motor peripheries, lead to only modest (polynomial, or  N) improvements in inverse squared error with increasing neuron number (N).

In *Beyond classical population coding for nearly exact estimation in the brain*, Ila Fiete explored an intriguing alternative. She showed that the entorhinal grid cell code for animal location is in a qualitatively different performance class than classical population codes. It allows unprecedented accuracy, enabling nearly exact removal of noise from noisy neural representations, with inverse squared error that improves exponentially ($\sim e^{aN}$ for some $a > 0$) with population size. The noise removal is enabled by the peculiar structure of the grid code, and does not rely on the existence of external cues. Moreover, a simple neural network model, similar to the hippocampus, can decode the grid represenation to take advantage of its error-control properties. This raises the possibility that the grid code is not unique, and that the brain could contain numerous examples of strong error-correcting codes for computing with analog variables.

**Tatyana Sharpee:** Maximally informative irregularities in neural circuits.

In *Maximally informative irregularities in neural circuits* Tatyana Sharpee explored the possibility that irregularities in neural circuits serve a useful computational function. To answer this question she and her colleagues focused on the retina, a well-studied circuit where many aspects of its average organization were previously found to be in good agreement with optimization principles. Previous experimental work has demonstrated the presence of fine scale irregularities in the shapes of individual receptive fields. Sharpee found that, in the presence of lattice irregularities, the irregular receptive field shapes increase the spatial resolution from 60% to 92% of that possible for a perfect lattice. Optimization of receptive field boundaries around their fixed center positions reproduced experimental observations on a neuron-by-neuron basis. These results suggest that lattice irregularities determine the shapes of retinal receptive fields and similar algorithms may improve the performance of the retinal prosthetics where substantial irregularities arise at their interface with the neural tissue. Taken more broadly, the results contribute to the emerging theme that irregularities in the organization of the nervous system are key to achieving its near optimal performance.

**John Beggs:** Information flow in networks of cortical neurons

John Beggs addressed the problem of information flow in networks of cortical neurons. Understanding this question would help us see how the brain integrates information across its different parts – a fundamental problem in neuroscience. Beggs noted that the average pyramidal neuron in cortex makes and receives approximately 7,000 synaptic contacts, suggesting that local cortical networks are connected in a fairly equal manner. The pattern of information flow in such networks, however, is poorly understood and can not be inferred from anatomy alone. Theory indicates that an unequal distribution of flows can actually contribute to network efficiency and robustness.

Beggs' group sought to examine the distribution of information flow in recordings from cortical slice cultures ($n = 6$) and monkey motor cortex ($n = 1$) containing $100 \pm 25$ identified neurons. They used transfer entropy to quantify information flow, as validation tests revealed that this measure could reliably distinguish true from spurious flows in a variety of realistic conditions. Beggs showed that information flow was distributed significantly more unevenly in the networks extracted from the data than in random control networks. This was evident in the distribution of information flow strengths, the distribution of total information flow into and out of each neuron, and in the distribution of connections with significant information flow per neuron. Simulations indicated the observed cortical information flow networks were significantly more efficient in routing signals, could form significantly more combinations among inputs per node, and were significantly more robust than random control networks. This is the first study of information flow in local cortical networks.

Beggs ended with an intriguing conclusion: The highly unequal distribution of information flow among cortical neurons contributes to the efficiency and robustness of information processing in cortex.

**Ruben Moreno Bote:** Weak synchrony in networks with finite input information

Neurons in cortex are correlated, but the functional role of these correlations remains elusive. It has been proposed that neuronal networks work in the so-called asynchronous state where correlations between pairs of neurons become vanishingly small for large networks. In such networks, the dynamics effectively decorrelates neurons firing, a process which is thought to improve coding because the percentage of input information conveyed by the output of the network increases with network size. However, these studies tacitly assume that the input information also increases without bounds with the network size, which is unrealistic. In contrast, we analyzed the dynamics of neuronal networks with finite input information. We find that with finite input information, and dominant inhibition, neuronal networks of integrate-and-fire neurons spontaneously settle down in a state of weak but not vanishingly weak- correlations, despite the dense connectivity, and despite the strong shared noise induced by the finite input information constraint. Moreover, and quite surprisingly, the finite information conveyed by the inputs is completely recoverable from the output spike counts of the network. We also show that excitation-dominated networks generate strong correlations but still preserve the input information in the output spike counts. In other words, whether the network decorrelates or not, there is no change in the amount of information transmitted. This challenges the notion that decorrelation is a universal mechanism for improving the quality of neural code. Moreover, given the relatively small correlation values observed in cortex, we propose that cortical networks are in a weakly synchronous state where inhibition dominates over excitation.

# 4 Linking feedforward and recurrent mechanisms for collective network activity

**Feedforward and recurrent networks both generate correlated activity. However, the mechanisms by which this is achieved are distinct. Real neural circuitry has aspects of both architectures, and it remains to be understood how the development of correlation in mixed feedforward and recurrent networks occurs. A fundamental piece of this puzzle is understanding the genesis of neural correlations in simple pairs of cells with different intrinsic dynamics; the meeting featured several discussions of this unresolved issue as well.**

**Duane Nykamp:** When feedforward intuition deceives: the influence of connectivity motifs on synchronization in recurrent networks

The synchronizability of a network captures how network structure can influence the tendency for a network to synchronize, independent of the dynamical model for each node. I demonstrate a synchronizability analysis that takes advantage of the framework of second order networks, which defines four second order connectivity statistics based on the relative frequency of two-edge network motifs. This analysis allows one to parametrically vary the amount of common input in a recurrent network in order to analyze the intuition from feedforward networks that common input is an important source of correlations and synchrony. In contrast to this intuition, the analysis determines that common input has little influence on synchrony in recurrent networks. Instead, the frequency of two-edge chains in the network plays a critical role, as synchrony increases dramatically with these chains. This dependence of synchrony on chains and not common input holds for a wide variety of neuron models.

**Ashok Kumar:** Shaping magnitude and timescale of correlations with state dependent synaptic input

Spike trains produced by sensory neurons often exhibit correlations that vary depending on neural state and stimulus properties. However, the circuit mechanisms responsible for changing the degree and timescale of pairwise correlated activity remain elusive. In contrast, many well-studied biophysical mechanisms have been shown to modulate single neuron response properties, such as firing rate gain. If these single neuron properties also shape correlations, then mechanisms of single neuron modulation may explain the shifts in correlation observed in sensory systems. We first study this possibility with simplified neuron models receiving varying levels of balanced, conductance-based synaptic input. This input modulates single neuron gain by changing membrane potential variability and conductance, and we show how this mechanism also shapes the timescale of correlation of neuron pairs. Next, we model and analyze data from the electrosensory system of weakly electric fish, in which recruitment of slow inhibitory feedback yields a reduction in single neuron transfer of low-frequency stimuli. We show that this effect also leads to a reduction in long timescale correlations between pairs of electrosensory neurons. These two studies demonstrate that modulatory synaptic input to neuron pairs can differentially shape precise spike time synchrony and average spike rate correlations. The results have consequences for state-dependent processing and propagation of neural activity.

**Robert Rosenbaum:** Using simplified integrate-and-fire models to understand how correlations propagate in neuronal networks

Robert Rosenbaum, a mathematics graduate student at the University of Houston, presented his work on the use of simplified integrate-and-fire models to understand how correlations propagate in neuronal networks. He observed that overlapping afferent populations and correlations between presynaptic spike trains can introduce correlations between the inputs to downstream cells. While in several other talks about the impact of correlations, participants have considered more detailed models, or larger networks, Robert described how simplified models can help us develop an intuitive and mechanistic description of the dominant mechanisms that control how correlations propagate. He showed several new results that proved that the degree to which input correlations are preserved is strongly modulated by cellular dynamics and also by synaptic variability. Both of these factors are frequently ignored in computational studies. He also demonstrated that correlations within an afferent population are significantly amplified by synaptic convergence. This amplification of correlations is the primary mechanism responsible for the synchronization of feedforward chains, a simple observation that did not seem to be widely known.

**Bard Ermentrout:** Correlation transfer by heterogeneous oscillators

Neurons are heterogeneous in many ways and this affects how they respond to inputs. In particular, identical inputs going into heterogeneous neurons will produce different outputs. There are many sources of heterogeneity including different channel distributions, different mean firing rates, and different noisy synaptic inputs. By considering the neurons as oscillators, we can write down equations for the phase of the oscillators as a function of the inputs. Using correlated white noise inputs, it is possible to compute the phases for a pair of heterogenous oscillators. The variance of these and the covariance allows us to compute the output correlation as a function of the input correlation. We derive formulae for the spike-count correlation for short windows (near synchrony) and find that neural oscillators that have a phase-resetting curve which has a nearly zero mean will maximize the transfer of correlation. For long time windows, the flatter is the PRC, the better the transfer of correlation. We combine these results with numerics to complete the picture of correlation transfer in long and short windows. We close by providing analytic approximations for the phase differences between two oscillators that have different frequencies and or differently shaped PRCs.