Connecting Network Architecture and Network Computation

Andrea Barreiro (Southern Methodist University)
Maurice Chacron (McGill University)
Brent Doiron (University of Pittsburgh)
Chris Eliasmith (University of Waterloo)
Krešimir Josić (University of Houston)
Eric Shea-Brown (University of Washington)

December, 6 – December 11, 2015

1 Overview of the Field

The brain solves an astonishing array of computational tasks – from representation, compression, and transformation of highly multivariate and multimodal sensory inputs, to statistical inference based on these inputs, to the expression and control of motor outputs in response. This is accomplished with speed and precision that easily outpaces the most powerful computers and best known algorithms today. Ultimately, the architectures — connectivity and dynamics — of neural circuits are responsible for this success. Understanding the link between architecture and computation is an inescapably mathematical challenge at the core of modern biological science.

Neural circuits achieve their computational power despite operating under severe constraints. Energetic considerations restrict the number of neurons that are available for any given function, together with the number and range of connections among cells. Moreover, communication between each cell incurs a delay. Biological circuits that integrate information from disparate sources need to preserve the relative timing of information arriving along different pathways – a particular challenge given the noise and unreliability of individual synaptic connections in neuronal networks. Finally, the large numbers of inputs that cells receive implies that neuronal networks need to operate in a state of precise balance between excitation and inhibition to prevent runaway synchrony.

2 Recent Developments and Open Problems

A variety of new experimental techniques provide unprecedented clues into how are neuronal networks structured to meet these constraints. In particular, new multielectrode and imaging techniques are revealing the simultaneous activity of neural ensembles — and, in some cases, entire neural populations — on a huge range of scales. At the same time, new methods allow the precise mapping of microcircuit anatomy (including both cellular dynamics and intercellular connectivity structures). In addition, techniques from microstimulation and optogenetics further allow us to manipulate the activity of specific cell types during meaningful behaviors.

Interpreting and optimizing these new empirical methods and the large data sets that they produce requires the development of new mathematics. There is a dual challenge: properly analyzing data with highly nonlinear dependencies among hundreds to thousands of variables, and developing mathematical and computational models that reveal and explain the relationship between network architecture and function. To understand the structure-function relationship – how neurons that display highly nonlinear and stochastic dynamics in isolation interact cooperatively when arranged in a circuit in order to perform behaviorally relevant computations under biological constraints – the research community must overcome major challenges in applied mathematics.

Three branches of the mathematical sciences must be brought together to enable the next set of advances – (1) dynamical systems, especially high-dimensional dynamics, (2) stochastic analysis, especially coupled random processes, and (3) the theory of network computation. Moreover, the space of possible network architectures and dynamics is so vast that working with experimentalists to guide and constrain the mathematical developments will be essential.

The aim of this workshop was to unite and provide a forum to focus the work of a group of international experts on network dynamics, information theory, and biological computation on the mathematical interface between network structure and function.

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 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?

3 Presentation Highlights

The meeting brought together a range of theoretical and computational neuroscientists, and experimental researchers interested in theory to address different aspects of this question. The quality of the talks was high, and – importantly – the presenters heeded the organizer's request to keep the talks accessible. This was particularly important due to the varied backgrounds of the participants. A number of graduate students and postdoctoral fellows also presented their research.

Below we highlight some of the presentations:

Emre Aksay (Weill Cornell Medical College)
Jacob Davidson (University of California Davis)

Network architectures underlying persistent neural activity

Persistent neural activity is important for motor control, short-term memory, and decision making. It is unclear what network processing mechanisms and architectures support this brain dynamic. In this joint presentation an experimental neuroscientist (Aksay), and a computational neuroscientist (Davidson) showed their recent efforts to address this question in the oculomotor integrator, a model system for studying persistent neural activity. They have determined candidate architectures by fitting dynamical network models to population-wide recordings with additional constraints from experiments on cellular morphology, intrinsic excitability, and localized perturbations. They also tested candidate architectures by imaging activity in the dendritic arbor of integrator neurons during persistent firing. These efforts suggest architectures of higher rank than previously assumed. Such architectures may allow persistent activity networks to act as hubs that perform numerous input-output transformations.

John Beggs (Indiana University)

High-degree neurons feed cortical computations

Recent results have shown that functional connectivity among cortical neurons is highly varied, with a small percentage of neurons having many more connections than others. Also, new theoretical work makes it possible to quantify how neurons modify information from the connections they receive. These developments have allowed John Beggs and his lab to investigate how information modification, or computation, depends on the number of connections a neuron receives (in- degree) or sends out (out-degree). John described how his lab uses a high-density 512 electrode array to record spontaneous spiking activity from cortical slice cultures and transfer entropy to construct a network of information flow. They have identified generic computations by the synergy produced wherever two information streams converged, and found that computations did not occur equally in all neurons throughout the networks. Surprisingly, the in-degree of a neuron was not related to the amount of information it computed. Instead, neurons that computed large amounts of information tended to receive connections from high out-degree neurons. To gain insight into these findings, the lab has developed a simple feedforward network model, and found that a degree- modified Hebbian wiring rule best reproduced the pattern of computation and degree correlation results seen in the real data. Interestingly, this rule also maximized signal propagation in the presence of network-wide correlations, suggesting a mechanism by which cortex could deal with common random background input. These are the first results to show that the extent to which a neuron modifies incoming information streams depends on its topological location in the surrounding functional network. Co-authors: Nick Timme and Sunny Nigam

Braden Brinkman (University of Washington)

Crouching tiger, hidden neuron

A major obstacle to understanding population coding in the brain is that neural activity can only be monitored at limited spatial and temporal scales. Inferences about network properties important for coding, such as connectivity between neurons, are sensitive to hidden units: unobserved neurons or other inputs that drive network activity. This problem is important not just for understanding inference from data, but also for which network properties shape spike train statistics as subsampled or pooled signals are transmitted through the brain. Recent computational efforts have been made to fit models to hidden units, but a fundamental theory of the effects of unobserved influences on the statistics of subsampled or pooled network activity remains elusive. Braden Brinkman showed how the methods from statistical physics can be used to develop an analytical framework to begin answering questions about how ground truth properties of neuronal networks are distorted when an experimenter (or downstream neuron) can only observe coarsely resolved activity data. As a specific example, he showed how the coupling filters of a generalized linear model fit to pooled spike train data change as a function of the fraction of spike trains pooled together. Coauthors: Fred Rieke, Eric Shea-Brown, Michael Buice

Nicholas Brunel (University of Chicago)

Statistics of connectivity optimizing information storage in recurrent networks

The rules of information storage in cortical circuits are the subject of ongoing debate. Two scenarios have been proposed by theorists: In the first scenario, specific patterns of activity representing external stimuli become fixed-point attractors of the dynamics of the network. In the second, the network stores sequences of patterns of network activity so that when the first pattern is presented the network retrieves the whole sequence. In both scenarios, the right dynamics are achieved thanks to appropriate changes in network connectivity. Nicolas Brunel described how methods from statistical physics can be used to investigate information storage capacity of such networks, and the statistical properties of network connectivity that optimize information storage (distribution of synaptic weights, probabilities of motifs, degree distributions, etc) in both scenarios. In the final part of the talk Nicolas compared the theoretical results with available data.

Kathleen Cullen (McGill University)

Neural correlates of sensory prediction errors during voluntary self-motion: evidence for internal models in the cerebellum.

The computation of sensory prediction errors is an important theoretical concept in motor control. In this context, the cerebellum is generally considered as the site of a forward model that predicts the expected sensory consequences of self-generated action. Changes in the motor apparatus and/or environment will cause a mismatch between the cerebellum's prediction and the actual resulting sensory stimulation. Thus this mismatch — the sensory prediction error — is thought to be vital for updating both the forward model and motor program during motor learning to ensure that sensory-motor pathways remain calibrated. In addition, through our daily activities, the computation of sensory prediction errors is required to discriminate externally-applied from self-generated inputs. However, direct proof for the existence of this comparison had been lacking.

Kathleen Cullen described how she and her collaborators took advantage of a relatively simple sensory-motor pathway with a well-described organization to gain insight into the computations that drive motor learning. The most medial of the deep cerebellar nuclei (fastigial nucleus), constitutes a major output target of the cerebellar cortex and in turn sends strong descending projections that ensure accurate posture and balance. They carried out a trial-by-trial analysis of these cerebellar neurons during the execution and adaptation of voluntary head movements and found that neuronal sensitivities dynamically tracked the comparison of predictive and feedback signals. When the relationship between the motor command and resultant movement was altered, neurons robustly responded to sensory input as if the movement was externally generated. Neuronal sensitivities then declined with the same time course as the concurrent behavioral learning. These findings demonstrate the output of an elegant computation in which rapid updating of an internal model enables the motor system to sense, and then learn to expect, previously unexpected sensory inputs. This enables both the i) rapid suppression of descending reflexive commands during voluntary movements and ii) rapid updating of motor programs in the face of changes to either the motor apparatus or external environment.

Jeff Dunworth (University of Pittsburgh)

Finite size effects and rare events in balanced cortical networks with plastic synapses

Cortical neuron spiking activity is broadly classified as temporally irregular and asynchronous. Model networks with a balance between large recurrent excitation and inhibition capture these two key features, and are a popular framework relating circuit structure and network dynamics. Balanced networks stabilize the asynchronous state through reciprocal tracking by the inhibitory and excitatory population activity, leading to a cancellation of total current correlations driving cells within the network. While asynchronous network dynamics are often a good approximation of neural activity, in many cortical datasets there are nevertheless brief epochs wherein the network dynamics are transiently synchronized (Buzsáki and Mizuseki (2014), Tan et al. (2014)). Jeff Dunworth described an analysis of paired whole cell voltage-clamp recordings from spontaneously active neurons in mouse auditory cortex slices (Graupner and Reyes (2013)) showing a network where correlated excitation and inhibition effectively cancel, except for intermittent periods when the network shows a macroscopic synchronous event. These data suggest that while the core mechanics of balanced activity are important, we require new theories capturing these brief but powerful periods when balance fails. Traditional balanced networks with linear firing rate dynamics have a single attractor, and fail to exhibit macroscopic synchronous events. Mongillo et al. (2012) showed that balanced networks with short-term synaptic plasticity can depart from strict linear dynamics through the emergence of multiple attractors. Jeff Dunworth and collaborators extended this model by incorporating finite network size, introducing strong nonlinearities in the firing rate dynamics and allowing finite size induced noise to elicit large scale, yet infrequent, synchronous events. They carried out a principled finite size expansion of an associated Markovian birth-death process and identified core requirements for system size and network plasticity to capture the transient synchronous activity observed in experimental data. This model properly mediates between the asynchrony of balanced activity and the tendency for strong recurrence to promote macroscopic population dynamics.

Stefano Fusi (Columbia University)

Computational principles of synaptic plasticity

Memories are stored, retained, and recollected through complex, coupled processes operating on multiple timescales. To understand the computational principles behind these intricate networks of interactions Stefano Fusi and collaborators constructed a broad class of synaptic models that efficiently harnesses biological complexity to preserve numerous memories. In these models the memory capacity scales almost linearly with the number of synapses, which is a substantial improvement over the square root scaling of previous models. This was achieved by combining multiple dynamical processes that initially store memories in fast variables and then progressively transfer them to slower variables. Importantly, the interactions between fast and slow variables are bidirectional. The proposed models are robust to parameter perturbations and can explain several properties of biological memory, including delayed expression of synaptic modifications, metaplasticity, and spacing effects.

Julijana Gjorgjieva (Brandeis University)

Optimal sensory coding by neuronal populations

In many sensory systems the neural signal splits into multiple parallel pathways, suggesting an evolutionary fitness benefit of a very general nature. For example, in the mammalian retina, ~ 20 types of retinal ganglion cells transmit information about the visual scene to the brain. What factors drove the evolution of such an early and elaborate pathway split remains elusive. Gjorgjieva and collaborators test the hypothesis that pathway splitting enables more efficient encoding of sensory stimuli, in the context of a specific prominent instance of sensory splitting: the emergence of ON and OFF pathways that code for stimulus increments and decrements, respectively. They developed a theory of optimal coding for a population of sensory ON and OFF neurons and computed the coding efficiency for different mixtures of ON and OFF cells. They found that optimal ON-OFF ratio in the population can be related to the statistics of natural stimuli, resulting in set of predictions for the optimal response properties of the neurons.

Cameron Harris (University of Washington)

Role and limits of inhibition in an excitatory burst generator

The pre-Bötzinger complex (preBötC) is now recognized as the essential core of respiratory rhythm generation, where it generates the inspiratory phase. Rhythmogenesis occurs through network synchronization. Harris and collaborators use a biophysical model of the entire preBötC to investigate several questions: What is the role of inhibitory cells in the preBötC? How does changing the sparsity of connections and synaptic strengths affect the resulting rhythm? These modeling results were compared to *in vitro* slice experiments in which inhibitory and excitatory synaptic transmission were progressively blocked. Harris found that too much sparsity or inhibition disrupts rhythm generation, yet highly connected networks without inhibition also produce non-biological rhythms. Their slice experiments suggest that the real preBötC lies within the partially synchronized region of network parameter space. As inhibitory neurons are added to the network, some cells fire out-of-phase with the main population rhythm, which offers an explanation for the out-of-phase cells observed in preBötC. However, Harris also reported that it is not possible to produce a two-phase population rhythm in their model, without adding further structure to the network. The preBötC and Bötzinger complexes therefore require structured networks in order to produce alternating inspiratory and expiratory rhythms. Finally, Harris presented preliminary stages of a spin model for oscillator phases which reproduces the qualitative features of the synchronization transition.

Kathryn Hedrick (Southern Methodist University)

Megamap: Flexible representation of a large space embedded with nonspatial information by a hippocampal attractor network

The problem of how the hippocampus encodes both spatial and nonspatial information at the cellular network level remains largely unresolved. Spatial memory is widely modeled through the theoretical framework of attractor networks, but standard computational models can only represent spaces that are much smaller than the natural habitat of an animal. Hedrick and collaborators propose that hippocampal networks are built upon a basic unit they call a *megamap*, or a cognitive attractor map in which place cells are flexibly recombined to represent a large space. Its inherent flexibility gives the megamap a huge representational capacity and enables the hippocampus to simultaneously represent multiple learned memories and naturally carry nonspatial information at no additional cost. On the other hand, the megamap is dynamically stable, as the underlying network of place cells robustly encodes any location in a large environment given a weak or incomplete input signal from the upstream entorhinal cortex. Hedrick's results suggest a general computational strategy by which a hippocampal network enjoys the stability of attractor dynamics without sacrificing the flexibility needed to represent a complex, changing world.

Zachary Kilpatrick (University of Houston and University of Colorado, Boulder)

Learning the volatility of a dynamic environment

Humans and other animals make perceptual decisions based on noisy sensory input. Recent studies focus on ecologically realistic situations in which the correct choice or the informative features of the stimulus change dynamically. Importantly, optimal evidence accumulation in changing environments requires discounting prior evidence at a rate determined by environmental volatility. To explain these observations, Zachary Kilpatrick showed how to extend previous accumulator models of decision making to the case where the correct choice changes at an unknown rate. An ideal observer can optimally infer these transition rates and accumulate evidence to make the best decision. He also discussed a neural implementation for this inference process whereby Hebbian plasticity shapes connectivity between populations representing each choice.

Ashok Litwin-Kumar (Columbia University)

Learning associations with both pure and randomly mixed representations

To make decisions and guide actions based on sensory information, neurons that mediate behavior must learn to respond appropriately to combinations of previously experienced stimuli and/or contexts. Many models of learning assume this is accomplished by a feedforward hierarchy of layers of neurons leading from input to desired output. However, sensory information is often relayed by multiple convergent and divergent pathways, each of which may have different representations of the input. Ashok Litwin-Kumar discussed the ability of output neurons that receive both pure stimulus information and randomly mixed stimulus/context information via an indirect pathway to perform associative learning. He showed that for realistic input-output mappings, the optimal pattern of connectivity is an intermediate one that includes input from both pure and mixed representations converging on the output layer. He also discussed the optimal level of mixing to maximize behavioral performance, finding, surprisingly, that sparse connectivity improves performance compared to the fully connected case. These results shed light on the principles governing learning from random representations, a strategy employed in many areas of the brain.

Artur Luczak (University of Lethbridge)

Neuronal activity packets as basic units of neuronal code

Neurons are active in a coordinated fashion, for example, an onset response to sensory stimuli usually evokes a $\sim 50 - 200$ms long burst of population activity. Recently it has been shown that such 'packets' of neuronal activity are composed of stereotypical sequential spiking patterns. The exact timing and number of spikes within packets convey information about the stimuli. Artur Luczak presented evidence that packets can be a good candidate for basic building blocks or 'the words' of neuronal coding, and can explain the mechanisms underlying multiple recent observations about neuronal coding, such as: multiplexing, LFP phase coding, and provide a possible connection between memory preplay and replay. This presentation summarized and expanded on the opinion paper: Luczak et al. (2015, Nature Rev. Neurosci.; doi:10.1038/nrn4026)

Cheng Ly (Virginia Commonwealth University)

Firing Rate Statistics with Intrinsic and Network Heterogeneity

Heterogeneity of neural attributes has recently gained a lot of attention and is increasing recognized as a crucial feature in neural processing. One outstanding question is how two distinct sources of heterogeneity — synaptic weights at the network level, and intrinsic excitability at the single-cell level — interact and alter neural activity. Cheng Ly used a recurrent spiking neural network model to study how these two forms of heterogeneity lead to different distributions of firing rates. Even in an uncoupled network, intrinsic heterogeneity naturally leads to firing rate heterogeneity. Coupling can lead to amplification or attenuation of firing rate heterogeneity, depending on both the relationship between intrinsic and network heterogeneity and the operating regime of the recurrent network: in particular, whether it is firing asynchronously or rhythmically. To analytically characterize these observations, Ly employed dimension reduction methods and asymptotic analysis to derive compact analytic descriptions of the phenomena. These descriptive formulas show how these two forms of heterogeneity determine the firing rate heterogeneity in various settings.

Dani Marti (École Normale Superieure)

Structured connectivity as a source of slow dynamics in randomly connected networks

Cortical networks exhibit dynamics on a range of timescales. Slow dynamics at the timescale of hundreds of milliseconds to seconds carry information about the recent history of the stimulus, and can therefore act as a substrate for short-term memory. How networks composed of fast units, like neurons, can generate such slow dynamics is still an open question. One possible mechanism is based on positive feedback: in randomly connected networks, the collective timescale can be set arbitrarily long by balancing the intrinsic decay rate of individual neurons with recurrent input. This type of mechanism relies however on fine-tuning the synaptic coupling.

Another possibility is that slow dynamics are induced by structured connectivity between neurons. In fact, the connectivity of cortical networks is not fully random. The simplest and most prominent deviation from randomness found in experimental data is the overrepresentation of bidirectional connections among pyramidal cells. Dani Marti argued that symmetry in the connectivity can act as a robust mechanism for the generation of slow dynamics in networks of fast units.

Using numerical and analytical methods, Dani Marti investigated the dynamics of networks with partially symmetric structure. He considered the two dynamical regimes exhibited by random neural networks: the weak-coupling regime, where the firing activity decays to a single fixed point unless the network is stimulated, and the strong-coupling or chaotic regime, characterized by internally generated fluctuating firing rates. He determined how symmetry modulates the timescale of the noise filtered by the network in the weak-coupling regime, as well as the timescale of the intrinsic rate fluctuations in the chaotic regime. In both cases symmetry increases the characteristic asymptotic decay time of the autocorrelation function. Furthermore, or sufficiently symmetric connections the network operating in the chaotic regime exhibits aging effects, by which the timescale of the rate fluctuations slowly grows as time evolves. Such history-dependent dynamics might constitute a new mechanism for short-term memory storage in random networks.

Michael G Metzen (McGill University)
Volker Hofmann (McGill University)

The role of neural correlations in information coding

The role of correlated neural activity in neural coding remains controversial. The two speakers showed that correlated neural activity can provide information about particular stimulus features independently of single neuron activity using the weakly electric fish, Apteronotus leptorhynchus as an animal model. These fish generate an electric organ discharge (EOD) surrounding their body, the amplitude of which is encoded in the discharge of electroreceptors (P-units) that synapse onto pyramidal neurons in the hindbrain electrosensory lateral line lobe (ELL) that in turn synapse onto neurons within the midbrain torus semicircularis (TS). When two conspecifics come into close proximity, each fish experiences a sinusoidal amplitude modulation (i.e. beat) with a frequency that is equal to the difference between both EOD frequencies. The beat can be further modulated due to movements of the animals, thus creating an envelope. Furthermore, these fish can generate communication signals or chirps (i.e. electrosensory objects) that consist of transient increases in EOD frequency and always occur simultaneously with the beat under natural conditions. The pairwise correlation coefficient but not single neuron spiking activity at the periphery: 1) can reliably be used to predict the stimulus envelope and 2) allows for the emergence of a feature invariant representation of natural communication stimuli that is actually exploited by the electrosensory system. Moreover, information carried by correlated neural activity at the periphery is decoded and further refined in downstream brain areas. This gives rise to similar behavioral responses to stimulus waveforms associated with a given electrosensory object. As such, correlated activity codes for stimulus attributes that are distinct from those coded by firing rate and provide a novel role for neural variability. Furthermore, correlated neural activity is invariant to identity preserving transformations of natural stimuli. This reveals how a sensory system exploits this fact in order to implement the emergence and refinement of invariant neural representations of natural stimuli and how these mediate perception and behavior. Herewith we show that neural correlations can serve as an extra channel of information coding independently of single neuron firing rate. The associated neural circuits are generic and thus likely to be found across systems and species.

Stephan Mihalas (Allen Institute of Brain Sciences)

Cortical circuits implementing optimal cue integration

Neurons in the primary visual cortex (V1) predominantly respond to a patch of the visual input, their classical receptive field. These responses are modulated by the visual input in the surround. This reflects the fact that features in natural scenes do not occur in isolation: lines, surfaces are generally continuous. There is information about a visual patch in its surround. This information is assumed to be passed to a neuron in V1 by neighboring neurons via lateral connections. The relation between visual evoked responses and lateral connectivity has been recently measured in mouse V1. In this study we combine these three topics: natural scene statistics, mouse V1 neuron responses and their connectivity. Stephan Mihalas addressed the question: Given a set of natural scene statistics, what lateral connections would optimally integrate the cues from the classical receptive field with those from the surround?

To do so he assumed a neural code: the firing rate of the neuron maps bijectively to the probability of the feature the neuron is representing to be in the presented image. He next generated a database of features these neurons represent by constructing a parameterized set of models from V1 electrophysiological responses, and used the Berkeley Segmentation Dataset to compute the probabilities of co-occurrences of these features. He computed the relation between probabilities of feature co-occurrences and the synaptic weight which optimally integrates these features. The relation between evoked responses and connectivity which leads to optimal cue integration is qualitatively similar to the measured one, but several additional predictions are made. Finally, he hypothesized that this computation: optimal cue integration is a general property of cortical circuits, and the rules constructed for mouse V1 generalize for other areas and species.

Ruben Moreno-Bote (University Pompeu Fabra)

Causal Inference in Spiking Networks

While the brain uses spiking neurons for communication, theoretical research on brain computations has mostly focused on non-spiking networks. The nature of spike-based algorithms that achieve complex computations, such as object probabilistic inference, is largely unknown. Here Ruben Moreno-Bote demonstrated that a family of high-dimensional quadratic optimization problems with non-negativity constraints can be solved exactly and efficiently by a network of spiking neurons. The network infers the set of most likely causes from an observation using explaining away, which is dynamically implemented by spike-based, tuned inhibition.

Katie Newhall (University of North Carolina-Chapel Hill)

Variability in Network Dynamics

Mathematical models of neuronal network dynamics, such as randomly connected integrate-and-fire model neurons, typically create homogeneous dynamics in the sense that a single neuron in the network is representative of the ensemble behavior, and dynamics in time are statistically repeatable. Katie Newhall will discuss work in progress on experimental data in which neither is true, looking at a statistical method to answer biological questions, and pondering the existence of simple model network motifs capable of producing such variability.

Alex Reyes (New York University)

Homeostatic control of neuronal firing rate and correlation: scaling of synaptic strength with network size

Features of sensory input are represented as the spatiotemporal activities of neuronal population. This network dynamics depends on the balance of excitatory (E) and inhibitory (I) drives to individual neurons. Maintaining balance in the face of continuously changing nervous system is vital for preserving the response properties of neurons and preventing neuropathologies. While homeostatic processes are in place to maintain excitatory level, the conditions for maintaining stable responses are yet unclear. Alex Reyes used a culture preparation to systematically vary the density of the network. Using optogenetic techniques to stimulate individual neurons in the network with high spatial and temporal resolution, he was able to systematically vary the number and correlation of external inputs. He found that the average firing rate and the correlation structure of synaptic inputs are invariant with network size. Finally, he showed how paired recordings could be used to measure the synaptic strengths and connection probability between excitatory (E) and inhibitory (I) neurons to confirm experimentally a long standing theoretical assumption that synaptic strength scales with the number of connections per neuron (N) closer to $N^{-1/2}$ than to N^{-1} .

Robert Rosenbaum (University of Notre Dame)

Correlations and dynamics in spatially extended balanced networks

Balanced networks offer an appealing theoretical framework for studying neural variability since they produce intrinsically noisy dynamics with some statistical features similar to those observed in cortical recordings. However, previous balanced network models face two critical shortcomings. First, they produce extremely weak spike train correlations, whereas cortical circuits exhibit both moderate and weak correlations depending on cortical area, layer and state. Second, balanced networks exhibit simple mean-field dynamics in which firing rates linearly track feedforward input. Cortical networks implement non-linear functions and produce non-trivial dynamics, for example, to produce motor responses. We propose that these shortcoming of balanced networks are overcome by accounting for the distance dependence of connection probabilities observed in cortex.

Robert Rosenbaum generalized the mean-field theory of firing rates, correlations and dynamics in balanced networks to account for distance-dependent connection probabilities. He showed that, under this extension, balanced networks can exhibit either weak or moderate spike train correlations, depending on the spatial profile of connections. Networks that produce moderate correlation magnitudes also produce a signature spatial correlation structure. A careful analysis of in vivo primate data reveals this same correlation structure. Finally, he showed that spatiotemporal firing rate dynamics can emerge spontaneously in spatially extended balanced networks. Principal component analysis reveals that these dynamics are fundamentally high-dimensional and reliable, suggesting a realistic spiking model for the rich dynamics underlying nontrivial neural computations. Taken together our results show that spatially extended balanced networks offer a parsimonious model of cortical circuits.

Woodrow Shew (University of Arkansas)

Functional implications of phase transitions in the cerebral cortex

A long-standing hypothesis at the nexus neuroscience, physics, and network science posits that a network of neurons may be tuned through a phase transition. Originally, this idea was motivated by intriguing analogies between the brain and physical systems which undergo phase transitions including Ising and percolation models. More recently, this idea has graduated from an appealing analogy to an experimentally supported and biophysically important fact. Woodrow Shew reviewed recent experiments and models, which have now established not only that phase transitions can occur in cerebral cortex, but also that neural information processing crucially depends on what phase the cortex is in. Cortical phase can be controlled by myriad biophysical mechanisms including tuning the balance of excitation and inhibition and adaptation to sensory input. Importantly, multiple aspects of information processing, such as sensory dynamic range and discrimination are optimized when the network operates nearby (but not exactly at) the critical point of a phase transition. These studies suggest that by operating in the vicinity of criticality the cerebral cortex may be tuned to accommodate changing information processing needs depending on behavioral context.

Aubrey Thompson (Carnegie Mellon University)

Relating spontaneous dynamics and stimulus coding in competitive networks

Understanding the relation between spontaneously active and stimulus evoked cortical dynamics is a recent challenge in systems neuroscience. Recordings across several cortices show highly variable spike trains during spontaneous conditions, and that this variability is promptly reduced when a stimulus drives an evoked response (Churchland, Yu, et al., Nat Neuro, 2010). Aubrey Thompson showed how networks of spiking neuron models with clustered excitatory architecture capture this key feature of cortical dynamics (Litwin-Kumar and Doiron, Nat. Neuro, 2012). In particular, clusters show stochastic transitions between periods of low and high firing rates, providing a mechanism for slow cortical variability that is operative in spontaneous states. She expanded on past work and explored a simple Markov neural model with clustered architecture, where spontaneous and evoked stochastic dynamics can be examined more carefully. She modeled the activity of each cluster in the network as a birth-death Markov process, with positive self feedback and inhibitory cluster-cluster competition. Her Markov model allows a calculation of the expected transition times between low and high activity states, yielding an estimate of the invariant density of cluster activity. Using this theory, she explored how the strength of inhibitory connections between the clusters sets the maximum likelihood for the number of active clusters in the network during spontaneous conditions. This work relates two disparate aspects of cortical computation—lateral inhibition and stimulus coding.

Jochen Triesch (Frankfurt Institute for Advanced Studies)

Wheres the noise? Key features of spontaneous activity and neural variability arise through learning in a deterministic network

Even in the absence of sensory stimulation the brain is spontaneously active. This background noise seems to be the dominant cause of the notoriously high trial-to-trial variability of neural recordings. Recent experimental observations have extended our knowledge of trial-to-trial variability and spontaneous activity in several directions: 1. Trial-to-trial variability systematically decreases following the onset of a sensory stimulus or the start of a motor act. 2. Spontaneous activity states in sensory cortex outline the region of evoked sensory responses. 3. Across development, spontaneous activity aligns itself with typical evoked activity patterns. 4. The spontaneous brain activity prior to the presentation of an ambiguous stimulus predicts how the stimulus will be interpreted. At present it is unclear how these observations relate to each other and how they arise in cortical circuits.

Jochen Triesch demonstrated that all of these phenomena can be accounted for by a deterministic self-organizing recurrent neural network model (SORN), which learns a predictive model of its sensory environment. The SORN comprises recurrently coupled populations of excitatory and inhibitory threshold units and learns via a combination of spike-timing dependent plasticity (STDP) and homeostatic plasticity mechanisms. Similar to balanced network architectures, units in the network show irregular activity and variable responses to inputs. Additionally, however, the SORN exhibits sequence learning abilities matching recent findings from visual cortex and the networks spontaneous activity reproduces the experimental findings mentioned above. Intriguingly, the networks behaviour is reminiscent of sampling-based probabilistic inference, suggesting that correlates of sampling-based inference can develop from the interaction of STDP and homeostasis in deterministic networks. He concluded that key observations on spontaneous brain activity and the variability of neural responses can be accounted for by a simple deterministic recurrent neural network which learns a predictive model of its sensory environment via a combination of generic neural plasticity mechanisms.

Aaron Russel Voelker (University of Waterloo)

Computing with temporal representations using recurrently connected populations of spiking neurons

The modeling of neural systems often involves representing the temporal structure of a dynamic stimulus. Aaron Voelker extended the methods of the Neural Engineering Framework (NEF) to generate recurrently connected populations of spiking neurons that compute functions across the history of a time-varying signal, in a biologically plausible neural network. To demonstrate the method, he proposed a novel construction to approximate a pure delay, and use that approximation to build a network that represents a finite history (sliding window) of its input. Specifically, he solved for the state-space representation of a pure time-delay filter using Pade-approximants, and then map this system onto the dynamics of a recurrently connected population. This construction is robust to noisy inputs over a range of frequencies, and can be used with a variety of neuron models including: leaky integrate-and-fire, rectified linear, and Izhikevich neurons. Furthermore, he extended the approach to handle various models of the post-synaptic current (PSC), and characterize the effects of the PSC model on overall dynamics. Finally, he showed that each delay may be modulated by an external input to scale the spacing of the sliding window on-the-fly.

Joel Zylberberg (University of Colorado School of Medicine)

Correlated stochastic resonance

Even when repeatedly presented with the same stimulus, sensory neurons show high levels of inter-trial variability. Similarly high levels of variability are observed throughout the brain, leading us to wonder how variability affects the function of neural circuits.

One the one hand, prior work on stochastic resonance (SR) has shown that random fluctuations can enhance information transmission by nonlinear circuit elements like neurons. Specifically, the thresholding inherent in spike generation means that much of the information contained within the membrane potential can fail to propagate downstream. Random membrane potential fluctuations soften spike thresholds, allowing more information to survive the spike-generation process. This phenomenon reflects a tradeoff between the positive effects of threshold-softening, and the negative effects of corrupting signals by noise.

While membrane potential fluctuations are often correlated between neurons in vivo, the role of this collective behavior in SR is largely unknown. Concurrently to the SR studies, other work investigated the impact of correlations on signal encoding by noisy non-spiking populations. For these non-spiking models, coding performance is highest when the noise is absent altogether: the noise is always a hindrance to the population codes. Consequently, those studies cannot reveal conditions under which collective variability enhances information coding. Despite these limitations, the prior studies of non-spiking models show that depending on the patterns of inter-neural correlation correlations can mitigate corruption of signals by noise. Combining ideas about correlations, and about SR, Joel Zylberberg showed that correlated membrane potential fluctuations can soften neural spiking thresholds without substantially corrupting the underlying signals with noise, thereby significantly enhancing spiking neural information coding.

4 Scientific Progress Made

Understanding the mechanisms by which the nervous system represents and processes information is a fundamental challenge for mathematical neuroscience. The goal of this meeting was to bring together experimental and theoretical neuroscientists to discuss the questions that will drive research in mathematical and computational neuroscience for years to come. Large initiatives, like the BRAIN Initiative in the US, and the Human Brain Project in Europe are driving the development of new experimental tools to study neural circuits, and large scale computational models of the brain. However, all participants agreed that an understanding of the resulting data and simulation results will require the development of new mathematical, and statistical approaches.The participants presented a range of such approaches.

One common theme that emerged was the necessity to extend the mathematical tools that have been developed to describe populations of statistically independent neurons. New experimental approaches have revealed that the joint activity of neural populations can carry information, and that neural computations are carried out collectively by many cells. A number of presentations pointe to how correlations, synchrony, collective oscillations and state transitions in population activity are essential to understanding brain function. Mathematical methods have been essential for this understanding, and for the development of computational model that will emulate the brain. It also became clear that, as of yet, no overarching theory or approach seemed to be sufficient to achieve these goals. Participants thus presented a range of techniques and approaches.

5 Outcome of the Meeting

The meeting succeeded in bringing together a number of leaders in related areas of computational and experimental neuroscience. The participants discussed different approaches to common central questions in neuroscience research: How neural activity subtends neural computation, and how both emerge from the patterns of connections and interaction in neural tissue.

The week was spent seeking bridges among mathematical disciplines, computational and experimental approaches. The ample time that we allowed for discussions, and the excellent facility at BIRS, resulted in discussions that continued far after the end of each talk. Many meeting participants reported that they have made new connections, learned about new techniques and ideas, and received valuable feedback and comments on their research.

A number of graduate students and postdoctoral fellows participated in the meeting. Nearly all of these participants also gave talks during the meeting. 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.