THE UNIVERSITY OF CHICAGO

DESCRIBING VARIABILITY IN NEOCORTICAL MICROCIRCUIT DYNAMICS WITH MULTINEURONAL SPIKE PATTERNS AND FUNCTIONAL NETWORKS

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES

AND THE PRITZKER SCHOOL OF MEDICINE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY

JOSEPH BERTRAND DECHERY

CHICAGO, ILLINOIS

AUGUST 2018

DEDICATION

To my parents, whose unending support gave me the confidence and resilience I needed.

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

The most important operations of the brain, from vision and audition to complex movements and decision making, necessitate coordinated activity in populations of neurons. Within an area of neocortex, nearby neurons form complex networks of connections that help shape the activity patterns we observe. Since neurons are inextricably embedded within their local population, we must consider neural populations as a whole if we are to understand how the information in the outside world is interpreted and represented in the brain. In this thesis, I present two complementary approaches to study how groups of neurons cooperatively generate patterns of activity. Together, the results demonstrate how trial-to-trial variability of population activity can be explained by functional relationships within groups of neurons. The first approach identifies small ensembles of neurons that reliably spike in sequence to investigate spike propagation in emergent, spontaneous activity. By resolving spatiotemporal spike patterns into sequences, I describe single trial variability and substantiate evidence for neocortical assembly phase sequences. Secondly, I construct a functional graph from correlated activity between pairs of neurons to holistically represent shared variability in visually-evoked population dynamics. These graphs displayed a specific pattern of functional connections underlying accurate predictions of trial-to-trial variability, illustrating a signature of informative correlations in neocortical networks. These analyses delineate the ways in which ensembles of neurons coordinate their activity to shape population dynamics. Extending models of neural activity from single cells to the networks they form have been difficult. By strengthening the lexicon used to capture variability and describe population activity on single trials, we can better investigate the biological sources and behavioral consequences of variability.

CHAPTER I: GENERAL INTRODUCTION

Thesis Outline

Spatiotemporal spike patterns define the behavior of populations of neurons. It is thought that the activity patterns in neocortex are responsible for representing sensory information, controlling complex motor commands, and guiding decisions. It remains a challenge both experimentally and analytically, however, to identify the specific features of population dynamics that represent this information. Beyond representation, the brain also needs to perform computations on incoming and internal information to appropriately direct behavior in a dynamically changing world. New experimental technologies now allow for simultaneous recording of hundreds of nearby neurons, but the descriptions for how these populations behave are still lacking. Importantly, behavior takes place in real time and so an understanding of spatiotemporal spike patterns must ultimately extend beyond trial-averaged responses to capture variability and describe single trial responses.

Due to the multi-disciplined nature of this work, I begin by providing context and discussing relevant background that motivates research into describing variability in neocortical microcircuits. This begins with a discussion of how to represent activity generated in the brain, followed by anatomical background that motivates a microcircuit perspective. The structure of neocortical microcircuits motivates and influences how they behave, and I summarize results to that end, as well as the technology underlying the innovative experiments allowing for investigation of neuronal populations. In Chapter II, I analyze spontaneous activity *in vitro* in order to observe how microcircuit activity propagates while isolated from top-down or bottom-

up influences. I develop a statistical method to identify sequences of spikes in small ensembles of neurons that occur more frequently than one would expect from the population-wide spatiotemporal spiking progression. These sequences densely manifest on single trials and provide experimental evidence for Hebbian assembly phase sequences in neocortical microcircuits. With *in vivo* recordings on visually-evoked activity, Chapter III extends this approach in order to study how microcircuits behave while representing stimulus information. Single trial cofluctuations across the population are represented as a functional network, and I find that a specific triplet motif is predictive of my ability to predict a neuron's activity from ongoing population activity. This finding illustrates how applied graph theory can bridge the scale between pairs and ensembles of neurons and the microcircuit. I synthesize and interpret the previous chapters in the discussion.

From Single-Cell Activity to Network Dynamics in Neocortex

Since the discovery of the action potential, neuroscientists have considered these discrete spikes in membrane voltage as the carriers of information in the brain. The neural correlates of perception and behavior are patterns of spikes emitted by populations of neurons. However, the computational complexity of discrete spiking patterns among all the neurons in the brain is unfathomable, and is further complicated by the biological complexity of the electrochemical synapse and dendritic computation. The study of microcircuit dynamics will not be tractable until we find a simplifying model of spiking activity.

Neuroscientists have yet to agree on a consistent representation of spiking in a single neuron. Spike trains in a single trial can be convolved with a Gaussian to obtain an estimate of the instantaneous firing rate at any given moment (Lehky 2010). Within a cluster of spikes, however, this process can wash out information about the relative timing of spikes leaving only rate information. To maintain spike-time information, time can be binned at a sufficient resolution to obtain a binary sequence with each bin containing a spike or not. These approaches have led to the dichotomy between a rate code versus a temporal code, with experimentalists and theorists on both sides of the debate. This may not be a true dichotomy, however. After all, to obtain a time-varying firing rate, one must choose the width of the Gaussian (Lehky 2010). A small enough width results in a firing rate trace approximated by the binary spike train. Conversely, larger time bins will result in multiple spikes per bin and can be interpreted as a step-function representation of instantaneous firing rate, rather than a continuous representation. The relevant spiking timescale is likely contextual, as both approaches have led to significant results. From the dissection of early visual processing in retina, lateral geniculate nucleus, and V1, receptive fields and tuning properties are classically defined by firing rates alone (Kandel 2012). More contemporary analyses of motor control by M1 microcircuits use instantaneous firing rates of neurons to relate joint kinematics (Vargas-Irwin et al. 2010; Saleh, Takahashi, and Hatsopoulos 2012). Auditory cortex, on the other hand, is capable of emitting temporally precise spiking, and recordings in nonhuman primates show that information about complex auditory stimuli begins to decrease with time bins longer than roughly 5ms (Kayser, Logothetis, and Panzeri 2010). Similarly, somatosensory cortex in rats exhibits activity in response to whisker stimulation that is encoded at a timescale of about 20ms (Jadhav, Wolfe, and Feldman 2009). It appears that a

strong determinant of the relevant timescale is the area of cortex in question (Runyan et al. 2017).

Regardless of the timescale used to represent spiking, a consistent feature of neural responses is its variability in response to identical sensory stimuli (Shadlen and Newsome 1998). This variability is often compared to a Poisson process (Tomko and Crapper 1974; Softky and Koch 1993; Goris, Movshon, and Simoncelli 2014), though this model is frequently only a point of comparison. Non-Poisson spiking behavior is common (Baddeley et al. 1997; Barbieri et al. 2001), due in part to biophysical properties such as burstiness or refratory periods. Futhermore, balanced networks exhibit chaotic dynamics (van Vreeswijk and Sompolinsky 1996), meaning variability is unavoidable, even under perfectly controlled experimental conditions. Recently, cortical state, or similarly, the ongoing activity in nearby neurons, has been identified as a major source of the variability. Spatiotemporal activity patterns in cat V1 measured by voltage sensitive dyes were capable of explaining variability in stimulus-driven responses (Arieli et al. 1996). A model of response strength incorporating a global term shared across all neurons was also very successful at predicting spike count variability (Lin et al. 2015).

The study of response variability is a promising direction toward a causal understanding of behavioral and perceptual variability. Even with the computational power of a neocortex, mammalian behavior is imperfect. Furthermore, organisms need to react in real-time. Theories of how information is encoded in microcircuit dynamics eventually need to extend beyond a trial-averaged perspective to a single-trial perspective. A set of analytic techniques have been applied

and adapted to microcircuit activity (Churchland et al. 2007). A common theme to these techniques, including but not limited to principal components analysis (Briggman, Abarbanel, and Kristan 2005), locally-linear embedding (Stopfer, Jayaraman, and Laurent 2003), and Gaussian-process factor analysis (Byron et al. 2009), is an attempt to identify a reduced-dimension subspace (equivalently latent factors or manifold) within which to embed microcircuit activity patterns (Cunningham and Yu 2014). Dimensionality-reduction techniques remain powerful models that capture variability in the high dimensionality of microcircuit activity and allow decoding of stimuli, but it is difficult to interpret how another brain area would respond to a single-trial trajectory in the reduced-dimensional subspace. A representation of population dynamics that maintains neuron identity or spiking activity will more interpretable when placing this code upstream or downstream of other brain areas. In this thesis, I identify a specific multi-neuronal spike pattern in Chapter II. In Chapter III, I represent similarity of pairwise activity patterns within the microcircuit in a graph theoretic framework. Though this methodology abstracts away the spiking activity, it maintains the identities of individual neurons and critically allows for relating single-neuron response properties like tuning to drifting gratings to their embedding and position in the graph.

The Structure of Neocortical Microcircuits

The foundations of the neocortical microcircuit date back to the influential research of Ramon y Cajal, who correctly identified the discrete nature of neurons and helped to develop the neural theory of neuroanatomy. A functional organizing principle in neocortex was proposed decades later, informed by the work of Vernon Mountcastle on the response properties in cat visual

cortex (1957). By traversing an electrode perpendicular to the pial surface, Mountcastle observed neurons with shared receptive fields. The first ideas of neocortical microcircuits began to take shape, with Mountcastle ultimately hypothesizing that cortical columns acted as the fundamental computational unit of the brain. Stated more directly, the microcircuit hypothesis posits that neurons connect in a specific organization to form a microcircuit repeated across cortex, and this microcircuit is fundamentally responsible for the computations performed in neocortex. Though the strict definition of the microcircuit may vary over time and between research groups, this underlying hypothesis is foundational to this thesis.

The structure of neocortical circuitry has proven to be complex, yet it displays several guiding principles. The first is the laminar organization of neocortex, with different densities and neuronal morphologies distributed between layers. This organization is set in action during development, with neuronal specification, migration, and connectivity driven by the interplay of gene expression and regulation (Molyneaux et al. 2007). More recent evidence in mouse barrel cortex shows that lamina development can be disrupted by blocking thalamo-cortical neurotransmission, identifying an activity-dependent mechanism for laminar development (Li et al. 2013). Through physiological and anatomical experiments, this organization results in a long list of connectivity patterns between cell types in different lamina, between cortical regions, and from cortex to subcortical areas (Harris and Shepherd 2015). Of particular interest is the inter-laminar connections which would define a singular cortical column microcircuit. These include the canonical loop of excitatory neurons from layer (L) 4 excitatory neurons primarily forming synapses onto superficial L2/3 neurons, projecting down to L5, then to L6, finally closing the

loop back to L4 (Gilbert and Wiesel 1983), though these connections are a drastic simplification. The second key element to microcircuit organization is the incorporation of inhibitory neurons that consistent of several different morphologies and serve different functions. They can broadly be classified by the expression of one of three different genes: parvalbumin, somatostatin, or vasoactive intestinal peptide (VIP) (Rudy et al. 2011). All three receive connections from local excitatory neurons, and VIP neurons inhibit both other inhibitory subtypes. Somatostatin and parvalbumin neurons differentially inhibit excitatory neurons by targeting their dendrites and soma, respectively. Finally, somatostatin disinhibit the circuit by also connecting to parvalbumin neurons (Xu et al. 2013; Pfeffer et al. 2013). Connection specificity between neurons can be found based on a myriad of factors—morphology, laminar location, gene expression, physiology—pointing to a well-defined, yet complex neocortical microcircuit.

Despite identifying organization principles of neocortical microcircuitry, not every neuron connects with every other possible neighbor. The connectivity patterns are only tendencies and are probabilistic, so a complementary approach has been studying the statistics of connectivity. Connection probability between excitatory and inhibitory neurons have been measured by painstaking experiments using whole-cell patch recordings to simultaneously measure direct synaptic connections *in vitro.* Nearby pyramidal cells only connect to each other with roughly 20% probability in L4 mouse barrel cortex (Lefort et al. 2009). In mouse auditory cortex, paired patch clamp recordings have measured the connectivity between excitatory and inhibitory neurons to be much higher, around 50% probability (Oswald et al. 2009). A consistent feature of synaptic connectivity statistics is its dependence on the somatic distance between neurons

(Oswald et al. 2009; Perin, Berger, and Markram 2011). This spatial decay in connectivity likelihood suggests a fixed space constant or size of the microcircuit, though the decay is difficult to accurately measure given the low connection probability, and it may differ between excitatory and inhibitory connections. Beyond the spatial organization, there are topological principles determining microcircuit structure. Though the canonical microcircuit consists mostly of asymmetric connections (see above), reciprocal connections in pyramidal neurons are more frequent that would occur if connections were independent (Song et al. 2005). This same paper shows that triplets of pyramidal neurons have biased connectivity toward certain motifs. Excitatory neurons also exhibit clustered connectivity patterns. That is, two neurons are more likely to be connected the more neighbors they share (Perin, Berger, and Markram 2011). Indirect measurement of connectivity via an inferred causal impact on activity, 'effective connectivity,' has revealed one topological feature of connectivity present among the entire microcircuit—the rich club network (Nigam et al. 2016)—and has since been corroborated in a reconstructed network of rat somatosensory cortex (Gal et al. 2017). This organization contains a small subset of neurons that are densely interconnected with each other and with the remaining network. Neocortical microcircuits also possess the property of small-worldness (Perin, Berger, and Markram 2011). This property is defined by the confluence of strong clustering in combination with short average path length (Watts and Strogatz 1998) which, in the case of neural networks, results in efficient signal processing properties (Lago-Fernández et al. 2000). These statistical and topological descriptions of neocortical microcircuitry substantiate the concept of a structural neocortical microcircuit.

Challenges remain, however, in finding consensus on the definition of the microcircuit. One contributing factor is evidence for differences in the cytoarchitecture across neocortical areas – famously primary motor cortex lacking a granular layer. Not all differences are genetically determined; some occur during development and are shaped by thalamocortical activity from the periphery. The distinction between V1 and higher order visual areas is dependent on intact thalamocortical drive to layer 4 (Chou et al. 2013), and the structure of rodent barrel cortex requires intact sensory drive from whisker stimulation (Erzurumlu and Gaspar 2012). These two papers exemplify the fact that sensory cortical areas are not only functionally, but anatomically shaped by the nature of the inputs they receive. Nonetheless, there is little evidence for major inter-areal differences in the underlying microcircuitry outside of layer 4 (Harris and Shepherd 2015). Canonical microcircuit structure remains compatible with these differences, provided that its function is generalizable to the requirements of different cortical areas. Another challenge posed to the microcircuit hypothesis could be difference across species. While the total volume, surface area, and cortical thickness may differ drastically across mammals (Barton and Harvey 2000), a consistent microcircuit structure may still exist at a small enough spatial scale. The density of neurons in cortex varies across species of primates (Herculano-Houzel et al. 2008), which would result in a microcircuit that differs either in spatial scale, number of neurons, or density across species.

Another difficulty in precisely defining the neocortical microcircuit is due to the tremendous recurrence in connectivity intra- and inter-cortically. Direct and indirect loops occur throughout cortex, as evidenced by the famous illustration of hierarchical connectivity in macaque visual

brain areas (Essen, Anderson, and Felleman 1992). While it is clear that neurons are discrete cellular units, a microcircuit cannot be discretized, and it will always have some degree of connectivity to nearby and distant microcircuits. In the context of a cortical column, these are commonly referred to as horizontal connections, with vertical connections occurring within the column. Though estimates of the proportion of horizontal to vertical connections vary, a recent paper reconstructed labeled neurons in cat area 17 and estimated the proportion of excitatory connections within an iso-orientation column that originate outside the column to be 92% (Stepanyants et al. 2009). The question then becomes 'Where does one draw the boundary between the end of one circuit and the beginning of the next.' Using the vertical column in visual cortex as the definitive representation of the microcircuit has received criticism (Olshausen and Field 2006; Rakic 2008; Boucsein et al. 2011), but the study of neocortex from a microcircuit perspective remains sound. Resolving the structure and function of neocortical microcircuits will drive important neuroscientific insights and could help inform future research into artificial intelligence and machine learning.

Functional Consequences of Network Architecture

No element of a biological system is completely described by its structure or anatomy alone. But understanding the structure can help to understand its function, and together, the structure and function give the full picture. Structure-function relationships play a large part in any area of biological research, from organismal biology to molecular biochemistry or biophysics. Understanding these relationships provide important context for comparative biology, and they have direct impact on potential engineering or medical intervention applications. In

neuroscience, the structure of neocortical microcircuits will inform how they behave and in turn, how they process information. In other words, the underlying connections within a population of neurons shapes and constrains the patterns of activity that are generated (Luczak and MacLean 2012). Structure-function relationships in microcircuits are difficult to identify due to the complexity of the system and its constituent element, the neuron (Horton and Adams 2005). This complexity compounds the difficulty in causally identifying a feature of microcircuit architecture that impacts its function. Nonetheless, researchers have made theoretical and experimental progress in understanding how the structure of the microcircuit influences its dynamics and its downstream impact on perception and behavior.

Theory of neural network dynamics often begins with a highly simplified model and advances toward more and more realistic models informed by biology. The seminal work of Wilson and Cowan (1973) exemplify this approach by developing a mean-field approximation to the dynamics of neural populations. Understanding population dynamics at the resolution of single neurons is possible from dynamical systems that model neuron subthreshold and spiking behavior. Considerable effort has been made to find accurate models of single neuron spiking behavior that has low computational complexity and is tunable to specific features of spiking including adaptation, burstiness, and refractoriness (Izhikevich 2003; Brette and Gerstner 2005), as opposed to the biophysical Hodgkin-Huxley model (Hodgkin and Huxley 1952) and its multi-compartmental extensions (Lindsay, Lindsay, and Rosenberg 2005). Despite the exponential improvement of computational power, biologically realistic simulations of networks still require large-scale collaborations and state of the art supercomputers (Markram et al. 2015).

Neocortical microcircuits exhibit asynchronous, irregular spiking patterns (Ecker et al. 2010), which can be replicated in models of recurrent neural networks with balanced excitation and inhibition (van Vreeswijk and Sompolinsky 1996; Brunel 2000). The balance of excitatory and inhibitory inputs in cortical networks has a strong experimental evidence (Shu, Hasenstaub, and McCormick 2003; Haider et al. 2006; Xue, Atallah, and Scanziani 2014). As a consequence of irregular firing patterns, the neural spiking becomes a sparse, efficient code with benefits for information propagation (Kumar, Rotter, and Aertsen 2008; Zhou and Yu 2018). Incorporating features of spatial connectivity in cortical microcircuits into network models has shown these theoretical predictions to be consistent in more biologically realistic models (Litwin-Kumar and Doiron 2014; Rosenbaum et al. 2017). Beyond acting as a counterpart to excitation and balancing runaway spiking, inhibition along with synaptic plasticity allows for microcircuits to implement an approximation to expectation maximization, a powerful inference technique (Nessler et al. 2013). Attempts to simulate microcircuits as realistically as possible have promisingly reproduced an array of experimental findings, but have yet to yield strong theoretical predictions (Markram et al. 2015). The structure of microcircuits has been extensively studied, but similar to population activity patterns, the complexity has caused difficulty in identifying which aspects are impactful and which are inconsequential.

Experimental Foundations of Microcircuit Investigation

Scientific understanding is born from the interplay between theory and experiment. Modern neuroscience follows a continuous cycle of experimental design and data collection, leading to

interpretation of the data and formation of models and theory explaining the observations. This informs further experimentation for confirmation, extension, and new investigation. Neither experiment nor theory can advance uninterrupted while the other is stalled. Experimental investigation of intact neural systems requires precision to measure the small fluctuations in electric potential within densely-packed neurons in the delicate biologic tissue. Furthermore, perhaps more-so than other biologic systems, neuroscientific recording requires careful ethical considerations. A generalizable theory of how the neocortical microcircuit represents sensory information and the motor control underlying behavior is currently lacking, in part because the ability to record from the intact microcircuit is a relatively recent development.

New experimental techniques probing neural function have historically led to large advances in neuroscientific understanding. Beginning over half a century ago, the electrophysiological behavior of neurons in the form of action potentials was revealed with dual electrodes recordings (Hodgkin, Huxley, and Katz 1952). Later, carefully crafted glass micropipette electrodes and recording hardware led to patch-clamp experiments with wide-reaching applications from ion channel dynamics (Hamill et al. 1981) to the multi-cell recordings that began revealing the behavior of the cortical microcircuit (Song et al. 2005). From a top-down perspective, non-invasive recordings of aggregate activity using blood-flow dependent magnetic resonance are solidifying the neural correlates of behavior in humans (Fellows 2004). Multineuronal recordings are necessary to bridge our understanding of single-cell activity and large, cortical areal function.

The investigation of neocortical microcircuits has only recently been possible, partly due to the development and refinement of two-photon scanning microscopy beginning in 1990 (Denk, Strickler, and Webb 1990). The measurement of neural activity is critically dependent on calcium-sensitive fluorophores to be present in the neuron cytoplasm, relating the spiking-dependent calcium concentration to the measured fluorescence (Eilers et al. 2001). Genetic engineering now allows for constitutively-expressed fluorescent proteins in mice (Dana et al. 2014) and makes microcircuit investigations much more accessible to research groups. Focused laser light has a spatial resolution necessary to measure activity in single-cells and even subcellular compartments. However, changes in fluorescence are caused by changes in calcium concentration – an indirect measurement of activity. Future development of voltage-sensitive fluorophores (Kuhn, Denk, and Bruno 2008) may be necessary for these optical techniques to achieve temporal resolution matching electrophysiological recordings. Population recordings with two-photon scanning microscopy is further limited in its temporal resolution, since the fluorescence can only be measured from one neuron at a time and the excitation laser must be sequentially directed to across the microscope field of view. A common approach to increase temporal resolution is to minimize the scan time (Sadovsky et al. 2011; Schuck, Annecchino, and Schultz 2014).

Beyond temporal and spatial resolution, two-photon techniques confer other advantages for the design of investigatory experimentation. Two-photon absorption occurs with photons of a higher wavelength, allowing for imaging in deep cortical layers, past several hundred micrometers of light-scattering tissue. Using standard techniques, two-photon scanning microscopy has been

applied *in vivo* to image all six layers of cortex (Tischbirek et al. 2015). Finally, and critical to the investigation of microcircuit dynamics, is the ability to record from spatially proximal neurons in a local population, making two-photon microscopy advantageous over multi-electrode arrays. Though recording from a single imaging plane is common and restricts data to one cortical layer, techniques using acousto-optic deflectors (Reddy et al. 2008) or temporal focusing of excitation light (Schrödel et al. 2013; Pégard et al. 2017) allow for recordings in a 3d volume. Chapter III includes *in vivo* recordings of over three hundred neurons in an imaging plane of layer 2/3 mouse V1 excitatory neurons in a field of view diameter of less than 800μm. *In vivo* recordings of dozens of neurons in small fields of view are now common (Carrillo-Reid et al. 2015; Ecker et al. 2010; Rikhye and Sur 2015; Harvey, Coen, and Tank 2012; Montijn et al. 2016). Development of this minimally-invasive technique was necessary to record network dynamics in the intact animal, leading to development of theories of microcircuit function.

What We Can Learn from Spatiotemporal Spike Patterns

Until neuroscientists crack the neural code, a valuable approach to understanding information representation in spike trains is one that is assumption-free, or one that assumes as little as possible. Information theory, originally developed for signal processing and communication systems (Shannon 1948), can elegantly measure informational content in spike trains with minimal assumptions about how that information is represented (Strong et al. 1998). In general, these methods have been useful in showing under what conditions stimuli are faithfully represented by spike trains (e.g. Lazar and Pnevmatikakis 2008). In applications to microcircuitry, a directed metric of information propagation, transfer entropy, has revealed

computational principles in microcircuit activity (Neymotin et al. 2011; Nigam et al. 2016). Maximum entropy models similarly have foundations in information theory, and have shown that activity patterns in retinal populations can be captured with pairwise interactions alone (Schneidman et al. 2006). Higher-order correlations (Ohiorhenuan et al. 2010), or spike history (Truccolo, Hochberg, and Donoghue 2010), however, are necessary to explain microcircuit dynamics in neocortex. Apart from an information theoretical framework, generalized linear models, inspired by receptive field properties, attempt to parameterize spiking probability by a set of linear filters applied to recent stimulus, spiking history, and other neurons, and incorporating known non-linearities in spiking activity (Truccolo et al. 2005). By studying the resulting filters, these models have revealed the structure of correlated activity within neural populations—a hallmark of neocortical dynamics. Correlations within a population can have diverse and sometimes surprising effects on stimulus encoding (Abbott and Dayan 1999; Josić et al. 2009; Zylberberg 2017). The sources and consequences of correlated variability can help dissect stimulus encoding properties (Pillow et al. 2008). This theoretical background can be categorized into two different approaches to studying population dynamics that are mirrored in the following two chapters. The first is to remain agnostic about what spike patterns represent in search of generalizable computing capabilities. Alternatively, stimulus-evoked activity can be anchored to features of an external stimulus, directly revealing how spike patterns encode information.

References

Abbott, L. F., and Peter Dayan. 1999. "The Effect of Correlated Variability on the Accuracy of a Population Code." *Neural Computation* 11 (1): 91–101.

Arieli, Amos, Alexander Sterkin, Amiram Grinvald, and A. D. Aertsen. 1996. “Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses.” *Science* 273 (5283): 1868–1871.

Baddeley, Roland, Larry F. Abbott, Michael CA Booth, Frank Sengpiel, Tobe Freeman, Edward A. Wakeman, and Edmund T. Rolls. 1997. “Responses of Neurons in Primary and Inferior Temporal Visual Cortices to Natural Scenes.” *Proceedings of the Royal Society of London B: Biological Sciences* 264 (1389): 1775–1783.

Barbieri, Riccardo, Michael C. Quirk, Loren M. Frank, Matthew A. Wilson, and Emery N. Brown. 2001. “Construction and Analysis of Non-Poisson Stimulus-Response Models of Neural Spiking Activity.” *Journal of Neuroscience Methods* 105 (1): 25–37.

Barton, Robert A., and Paul H. Harvey. 2000. “Mosaic Evolution of Brain Structure in Mammals.” *Nature* 405 (6790): 1055.

Boucsein, Clemens, Martin Nawrot, Philipp Schnepel, and Ad Aertsen. 2011. “Beyond the Cortical Column: Abundance and Physiology of Horizontal Connections Imply a Strong Role for Inputs from the Surround.” *Frontiers in Neuroscience* 5.

Brette, Romain, and Wulfram Gerstner. 2005. “Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity.” *Journal of Neurophysiology* 94 (5): 3637–42.

Briggman, Kevin L., Henry DI Abarbanel, and William B. Kristan. 2005. “Optical Imaging of Neuronal Populations during Decision-Making.” *Science* 307 (5711): 896–901.

Brunel, Nicolas. 2000. “Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons.” *Journal of Computational Neuroscience* 8 (3): 183–208.

Byron, M. Yu, John P. Cunningham, Gopal Santhanam, Stephen I. Ryu, Krishna V. Shenoy, and Maneesh Sahani. 2009. “Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity.” In *Advances in Neural Information Processing Systems*, 1881–1888.

Carrillo-Reid, Luis, Jae-eun Kang Miller, Jordan P. Hamm, Jesse Jackson, and Rafael Yuste. 2015. “Endogenous Sequential Cortical Activity Evoked by Visual Stimuli.” *Journal of Neuroscience* 35 (23): 8813–8828.

Chou, Shen-Ju, Zoila Babot, Axel Leingärtner, Michele Studer, Yasushi Nakagawa, and Dennis D. M. O’Leary. 2013. “Geniculocortical Input Drives Genetic Distinctions between Primary and Higher-Order Visual Areas.” *Science (New York, N.Y.)* 340 (6137): 1239–42.

Churchland, Mark M, Byron M Yu, Maneesh Sahani, and Krishna V Shenoy. 2007. “Techniques for Extracting Single-Trial Activity Patterns from Large-Scale Neural Recordings.” *Current Opinion in Neurobiology*, Neuronal and glial cell biology / New technologies, 17 (5): 609–18.

Cunningham, John P., and Byron M. Yu. 2014. “Dimensionality Reduction for Large-Scale Neural Recordings.” *Nature Neuroscience* 17 (11): 1500–1509.

Dana, Hod, Tsai-Wen Chen, Amy Hu, Brenda C. Shields, Caiying Guo, Loren L. Looger, Douglas S. Kim, and Karel Svoboda. 2014. “Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging in Vivo.” *PloS One* 9 (9): e108697.

Denk, W., J. H. Strickler, and W. W. Webb. 1990. “Two-Photon Laser Scanning Fluorescence Microscopy.” *Science* 248 (4951): 73–76.

Ecker, Alexander S., Philipp Berens, Georgios A. Keliris, Matthias Bethge, Nikos K. Logothetis, and Andreas S. Tolias. 2010. “Decorrelated Neuronal Firing in Cortical Microcircuits.” *Science* 327 (5965): 584–87.

Eilers, Jens, Tim D. Plant, Nima Marandi, and Arthur Konnerth. 2001. “GABA-Mediated Ca2+ Signalling in Developing Rat Cerebellar Purkinje Neurones.” *The Journal of Physiology* 536 (2): 429–437.

Erzurumlu, Reha S., and Patricia Gaspar. 2012. “Development and Critical Period Plasticity of the Barrel Cortex.” *The European Journal of Neuroscience* 35 (10): 1540–53.

Essen, DC Van, C. H. Anderson, and D. J. Felleman. 1992. “Information Processing in the Primate Visual System: An Integrated Systems Perspective.” *Science* 255 (5043): 419–23.

Fellows, Lesley K. 2004. “The Cognitive Neuroscience of Human Decision Making: A Review and Conceptual Framework.” *Behavioral and Cognitive Neuroscience Reviews* 3 (3): 159–172.

Gal, Eyal, Michael London, Amir Globerson, Srikanth Ramaswamy, Michael W. Reimann, Eilif Muller, Henry Markram, and Idan Segev. 2017. “Rich Cell-Type-Specific Network Topology in Neocortical Microcircuitry.” *Nature Neuroscience* 20 (7): 1004–13.

Gilbert, C. D., and T. N. Wiesel. 1983. “Functional Organization of the Visual Cortex.” *Progress in Brain Research* 58: 209–18.

Goris, Robbe LT, J. Anthony Movshon, and Eero P. Simoncelli. 2014. “Partitioning Neuronal Variability.” *Nature Neuroscience* 17 (6): 858.

Haider, Bilal, Alvaro Duque, Andrea R. Hasenstaub, and David A. McCormick. 2006. “Neocortical Network Activity in Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition.” *Journal of Neuroscience* 26 (17): 4535–4545.

Hamill, Owen P., A. Marty, Erwin Neher, Bert Sakmann, and F. J. Sigworth. 1981. “Improved Patch-Clamp Techniques for High-Resolution Current Recording from Cells and Cell-Free Membrane Patches.” *Pflügers Archiv* 391 (2): 85–100.

Harris, Kenneth D., and Gordon MG Shepherd. 2015. “The Neocortical Circuit: Themes and Variations.” *Nature Neuroscience* 18 (2): 170.

Harvey, Christopher D., Philip Coen, and David W. Tank. 2012. “Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task.” *Nature* 484 (7392): 62.

Herculano-Houzel, S., C. E. Collins, P. Wong, J. H. Kaas, and R. Lent. 2008. “The Basic Nonuniformity of the Cerebral Cortex.” *Proceedings of the National Academy of Sciences* 105 (34): 12593–98.

Hodgkin, A. L., and A. F. Huxley. 1952. “A Quantitative Description of Membrane Current and Its Application to Conduction and Excitation in Nerve.” *The Journal of Physiology* 117 (4): 500–544.

Hodgkin, A. L., A. F. Huxley, and B. Katz. 1952. “Measurement of Current-Voltage Relations in the Membrane of the Giant Axon of Loligo.” *The Journal of Physiology* 116 (4): 424–48.

Horton, Jonathan C, and Daniel L Adams. 2005. “The Cortical Column: A Structure without a Function.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 360 (1456): 837–62.

Izhikevich, E. M. 2003. “Simple Model of Spiking Neurons.” *IEEE Transactions on Neural Networks* 14 (6): 1569–72.

Jadhav, Shantanu P., Jason Wolfe, and Daniel E. Feldman. 2009. “Sparse Temporal Coding of Elementary Tactile Features during Active Whisker Sensation.” *Nature Neuroscience* 12 (6): 792–800.

Josić, Krešimir, Eric Shea-Brown, Brent Doiron, and Jaime de la Rocha. 2009. “Stimulus-Dependent Correlations and Population Codes.” *Neural Computation* 21 (10): 2774–2804.

Kandel, Eric R. 2012. *Principles of Neural Science, Fifth Edition*. Edited by James H. Schwartz, Thomas M. Jessell, Steven A. Siegelbaum, and A. J. Hudspeth. 5th edition. New York: McGraw-Hill Education / Medical.

Kayser, Christoph, Nikos K. Logothetis, and Stefano Panzeri. 2010. “Millisecond Encoding Precision of Auditory Cortex Neurons.” *Proceedings of the National Academy of Sciences of the United States of America* 107 (39): 16976–81.

Kuhn, B., W. Denk, and R. M. Bruno. 2008. “In Vivo Two-Photon Voltage-Sensitive Dye Imaging Reveals Top-down Control of Cortical Layers 1 and 2 during Wakefulness.” *Proceedings of the National Academy of Sciences* 105 (21): 7588–93.

Kumar, Arvind, Stefan Rotter, and Ad Aertsen. 2008. “Conditions for Propagating Synchronous Spiking and Asynchronous Firing Rates in a Cortical Network Model.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 28 (20): 5268–80.

Lago-Fernández, L. F., R. Huerta, F. Corbacho, and J. A. Sigüenza. 2000. “Fast Response and Temporal Coherent Oscillations in Small-World Networks.” *Physical Review Letters* 84 (12): 2758–61.

Lazar, Aurel A., and Eftychios A. Pnevmatikakis. 2008. “Faithful Representation of Stimuli with a Population of Integrate-and-Fire Neurons.” *Neural Computation* 20 (11): 2715–44.

Lefort, Sandrine, Christian Tomm, J.-C. Floyd Sarria, and Carl C. H. Petersen. 2009. “The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex.” *Neuron* 61 (2): 301–16.

Lehky, Sidney R. 2010. “Decoding Poisson Spike Trains by Gaussian Filtering.” *Neural Computation* 22 (5): 1245–71.

Li, Hong, Sofia Fertuzinhos, Ethan Mohns, Thomas S. Hnasko, Matthijs Verhage, Robert Edwards, Nenad Sestan, and Michael C. Crair. 2013. “Laminar and Columnar Development of Barrel Cortex Relies on Thalamocortical Neurotransmission.” *Neuron* 79 (5): 970–86.

Lin, I.-Chun, Michael Okun, Matteo Carandini, and Kenneth D. Harris. 2015. “The Nature of Shared Cortical Variability.” *Neuron* 87 (3): 644–656.

Lindsay, A. E., K. A. Lindsay, and J. R. Rosenberg. 2005. “Increased Computational Accuracy in Multi-Compartmental Cable Models by a Novel Approach for Precise Point Process Localization.” *Journal of Computational Neuroscience* 19 (1): 21–38.

Litwin-Kumar, Ashok, and Brent Doiron. 2014. “Formation and Maintenance of Neuronal Assemblies through Synaptic Plasticity.” *Nature Communications* 5: 5319.

Luczak, Artur, and Jason N. MacLean. 2012. “Default Activity Patterns at the Neocortical Microcircuit Level.” *Frontiers in Integrative Neuroscience* 6: 30.

Markram, Henry, Eilif Muller, Srikanth Ramaswamy, Michael W. Reimann, Marwan Abdellah, Carlos Aguado Sanchez, Anastasia Ailamaki, et al. 2015. “Reconstruction and Simulation of Neocortical Microcircuitry.” *Cell* 163 (2): 456–92.

Molyneaux, Bradley J., Paola Arlotta, Joao R. L. Menezes, and Jeffrey D. Macklis. 2007. "Neuronal Subtype Specification in the Cerebral Cortex." *Nature Reviews Neuroscience* 8 (6): 427–37.

Montijn, Jorrit S., Guido T. Meijer, Carien S. Lansink, and Cyriel MA Pennartz. 2016. "Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective." *Cell Reports* 16 (9): 2486–2498.

Mountcastle, Vernon B. 1957. "Modality and Topographic Properties of Single Neurons of Cat's Somatic Sensory Cortex." *Journal of Neurophysiology* 20 (4): 408–34.

Nessler, Bernhard, Michael Pfeiffer, Lars Buesing, and Wolfgang Maass. 2013. "Bayesian Computation Emerges in Generic Cortical Microcircuits through Spike-Timing-Dependent Plasticity." *PLOS Computational Biology* 9 (4): e1003037.

Neymotin, Samuel A., Kimberle M. Jacobs, André A. Fenton, and William W. Lytton. 2011. "Synaptic Information Transfer in Computer Models of Neocortical Columns." *Journal of Computational Neuroscience* 30 (1): 69–84.

Nigam, Sunny, Masanori Shimono, Shinya Ito, Fang-Chin Yeh, Nicholas Timme, Maxym Myroshnychenko, Christopher C. Lapish, et al. 2016. "Rich-Club Organization in Effective Connectivity among Cortical Neurons." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 36 (3): 670–84.

Ohiorhenuan, Ifije E., Ferenc Mechler, Keith P. Purpura, Anita M. Schmid, Qin Hu, and Jonathan D. Victor. 2010. "Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks." *Nature* 466 (7306): 617.

Olshausen, Bruno A., and David J. Field. 2006. "What Is the Other 85 Percent of V1 Doing." *23 Problems in Systems Neuroscience*, 182–221.

Oswald, Anne-Marie M., Brent Doiron, John Rinzel, and Alex D. Reyes. 2009. "Spatial Profile and Differential Recruitment of GABAB Modulate Oscillatory Activity in Auditory Cortex." *Journal of Neuroscience* 29 (33): 10321–10334.

Pégard, Nicolas C., Alan R. Mardinly, Ian Antón Oldenburg, Savitha Sridharan, Laura Waller, and Hillel Adesnik. 2017. "Three-Dimensional Scanless Holographic Optogenetics with Temporal Focusing (3D-SHOT)." *Nature Communications* 8 (1): 1228.

Perin, Rodrigo, Thomas K. Berger, and Henry Markram. 2011. "A Synaptic Organizing Principle for Cortical Neuronal Groups." *Proceedings of the National Academy of Sciences* 108 (13): 5419–5424.

Pfeffer, Carsten K., Mingshan Xue, Miao He, Z. Josh Huang, and Massimo Scanziani. 2013. "Inhibition of Inhibition in Visual Cortex: The Logic of Connections between Molecularly Distinct Interneurons." *Nature Neuroscience* 16 (8): 1068–76.

Pillow, Jonathan W., Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M. Litke, E. J. Chichilnisky, and Eero P. Simoncelli. 2008. "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population." *Nature* 454 (7207): 995.

Rakic, Pasko. 2008. "Confusing Cortical Columns." *Proceedings of the National Academy of Sciences* 105 (34): 12099–100.

Reddy, Gaddum Duemani, Keith Kelleher, Rudy Fink, and Peter Saggau. 2008. "Three-Dimensional Random Access Multiphoton Microscopy for Functional Imaging of Neuronal Activity." *Nature Neuroscience* 11 (6): 713.

Rikhye, Rajeev V., and Mriganka Sur. 2015. "Spatial Correlations in Natural Scenes Modulate Response Reliability in Mouse Visual Cortex." *Journal of Neuroscience* 35 (43): 14661–14680.

Rosenbaum, Robert, Matthew A. Smith, Adam Kohn, Jonathan E. Rubin, and Brent Doiron. 2017. "The Spatial Structure of Correlated Neuronal Variability." *Nature Neuroscience* 20 (1): 107.

Rudy, Bernardo, Gordon Fishell, SooHyun Lee, and Jens Hjerling-Leffler. 2011. "Three Groups of Interneurons Account for Nearly 100% of Neocortical GABAergic Neurons." *Developmental Neurobiology* 71 (1): 45–61.

Runyan, Caroline A., Eugenio Piasini, Stefano Panzeri, and Christopher D. Harvey. 2017. "Distinct Timescales of Population Coding across Cortex." *Nature* 548 (7665): 92–96.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. 2011. "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging." *Journal of Neurophysiology* 106 (3): 1591–1598.

Saleh, Maryam, Kazutaka Takahashi, and Nicholas G. Hatsopoulos. 2012. "Encoding of Coordinated Reach and Grasp Trajectories in Primary Motor Cortex." *Journal of Neuroscience* 32 (4): 1220–32.

Schneidman, Elad, Michael J. Berry II, Ronen Segev, and William Bialek. 2006. "Weak Pairwise Correlations Imply Strongly Correlated Network States in a Neural Population." *Nature* 440 (7087): 1007.

Schrödel, Tina, Robert Prevedel, Karin Aumayr, Manuel Zimmer, and Alipasha Vaziri. 2013. "Brain-Wide 3D Imaging of Neuronal Activity in Caenorhabditis Elegans with Sculpted Light." *Nature Methods* 10 (10): 1013–20.

Schuck, Renaud, Luca A. Annecchino, and Simon R. Schultz. 2014. “Scaling up Multiphoton Neural Scanning: The SSA Algorithm.” *Conference Proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference* 2014: 2837–40.

Shadlen, Michael N., and William T. Newsome. 1998. “The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding.” *Journal of Neuroscience* 18 (10): 3870–3896.

Shannon, C. E. 1948. “A Mathematical Theory of Communication.” *Bell System Technical Journal* 27: 379–423.

Shu, Yousheng, Andrea Hasenstaub, and David A. McCormick. 2003. “Turning on and off Recurrent Balanced Cortical Activity.” *Nature* 423 (6937): 288–93.

Softky, W. R., and C. Koch. 1993. “The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 13 (1): 334–50.

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B. Chklovskii. 2005. “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits.” *PLOS Biology* 3 (3): e68.

Stepanyants, Armen, Luis M. Martinez, Alex S. Ferecskó, and Zoltán F. Kisvárday. 2009. “The Fractions of Short- and Long-Range Connections in the Visual Cortex.” *Proceedings of the National Academy of Sciences of the United States of America* 106 (9): 3555–60.

Stopfer, Mark, Vivek Jayaraman, and Gilles Laurent. 2003. “Intensity versus Identity Coding in an Olfactory System.” *Neuron* 39 (6): 991–1004.

Strong, Steven P., Roland Koberle, Rob R. de Ruyter van Steveninck, and William Bialek. 1998. “Entropy and Information in Neural Spike Trains.” *Physical Review Letters* 80 (1): 197.

Tischbirek, Carsten, Antje Birkner, Hongbo Jia, Bert Sakmann, and Arthur Konnerth. 2015. “Deep Two-Photon Brain Imaging with a Red-Shifted Fluorometric Ca2+ Indicator.” *Proceedings of the National Academy of Sciences of the United States of America* 112 (36): 11377–82.

Tomko, G. J., and D. R. Crapper. 1974. “Neuronal Variability: Non-Stationary Responses to Identical Visual Stimuli.” *Brain Research* 79 (3): 405–18.

Truccolo, Wilson, Uri T. Eden, Matthew R. Fellows, John P. Donoghue, and Emery N. Brown. 2005. “A Point Process Framework for Relating Neural Spiking Activity to Spiking History,

Neural Ensemble, and Extrinsic Covariate Effects." *Journal of Neurophysiology* 93 (2): 1074–89.

Truccolo, Wilson, Leigh R. Hochberg, and John P. Donoghue. 2010. "Collective Dynamics in Human and Monkey Sensorimotor Cortex: Predicting Single Neuron Spikes." *Nature Neuroscience* 13 (1): 105–11.

Vargas-Irwin, Carlos E., Gregory Shakhnarovich, Payman Yadollahpour, John M. K. Mislow, Michael J. Black, and John P. Donoghue. 2010. "Decoding Complete Reach and Grasp Actions from Local Primary Motor Cortex Populations." *Journal of Neuroscience* 30 (29): 9659–69.

Vreeswijk, C. van, and H. Sompolinsky. 1996. "Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity." *Science (New York, N.Y.)* 274 (5293): 1724–26.

Watts, Duncan J., and Steven H. Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature* 393 (6684): 440.

Wilson, H. R., and J. D. Cowan. 1973. "A Mathematical Theory of the Functional Dynamics of Cortical and Thalamic Nervous Tissue." *Kybernetik* 13 (2): 55–80.

Xu, Han, Hyo-Young Jeong, Robin Tremblay, and Bernardo Rudy. 2013. "Neocortical Somatostatin-Expressing GABAergic Interneurons Disinhibit the Thalamorecipient Layer 4." *Neuron* 77 (1): 155–67.

Xue, Mingshan, Bassam V. Atallah, and Massimo Scanziani. 2014. "Equalizing Excitation–inhibition Ratios across Visual Cortical Neurons." *Nature* 511 (7511): 596–600.

Zhou, Shanglin, and Yuguo Yu. 2018. "Synaptic E-I Balance Underlies Efficient Neural Coding." *Frontiers in Neuroscience* 12: 46.

Zylberberg, Joel. 2017. "Untuned but Not Irrelevant: A Role for Untuned Neurons in Sensory Information Coding." *BioRxiv*, 134379.

CHAPTER II

EMERGENT CORTICAL CIRCUIT DYNAMICS CONTAIN DENSE, INTERWOVEN ENSEMBLES OF SPIKE SEQUENCES

Abstract

Temporal codes are theoretically powerful encoding schemes, but their precise form in neocortex remains unknown in part because of the large number of possible codes and the difficulty in disambiguating informative spikes from statistical noise. A biologically plausible and computationally powerful temporal coding scheme is the Hebbian assembly phase sequence (APS), which predicts reliable propagation of spikes between functionally related assemblies of neurons. Here, we sought to measure the inherent capacity of neocortical networks to produce reliable sequences of spikes, as would be predicted by an APS code. To record microcircuit activity, the scale at which computation is implemented, we used two-photon calcium imaging to densely sample spontaneous activity in murine neocortical networks *ex vivo*. We show that the population spike histogram is sufficient to produce a spatiotemporal progression of activity across the population. To more comprehensively evaluate the capacity for sequential spiking that cannot be explained by the overall population spiking, we identify statistically-significant spike sequences. We found a large repertoire of sequence spikes that collectively comprise the majority of spiking in the circuit. Sequences manifest probabilistically and share neuron membership, resulting in unique ensembles of interwoven sequences characterizing individual spatiotemporal progressions of activity. Distillation of population dynamics into its constituent sequences provides a way to capture trial-to-trial variability, and may prove to be a powerful

decoding substrate *in vivo*. Informed by these data, we suggest that the Hebbian APS be reformulated as interwoven sequences with flexible assembly membership due to shared, overlapping neurons.

Introduction

Despite reliable perception and behavior, a hallmark of neural activity is a relatively high degree of trial-to-trial variability (Vogels, Spileers, and Orban 1989; Heggelund and Albus 1978; Arieli et al. 1996). It is not yet known to what degree this variability is related to cortical function or is simply uninformative biological noise. Identifying a complete picture of the statistics describing population dynamics is necessary to distinguish between noisy and informative spiking and in turn understand cortical function and computation. A general feature of neuronal population dynamics is recurring spatiotemporal patterns of neuronal activity, occurring in a wide range of preparations and under a variety of conditions (Johnson, Goel, and Buonomano 2010; Beggs and Plenz 2004; Peters, Chen, and Komiyama 2014; Bathellier, Ushakova, and Rumpel 2012; Villa et al. 1999). Spatiotemporal patterns have been found to correlate with perception, sensory input, and motor behavior (Harvey, Coen, and Tank 2012; Luczak, Barthó, and Harris 2009; Churchland et al. 2012), but we still lack the necessary tools to fully capture and describe complex population spike patterns. Noisy spatiotemporal patterns can be summarized as trajectories using dimensionality reduction techniques, facilitating accurate decoding and interpretability (Byron et al. 2009; Briggman, Abarbanel, and Kristan 2005; Broome, Jayaraman, and Laurent 2006). However, neocortical computation is implemented at the level of spiking neurons, making the mechanisms underlying these reduced-dimension trajectories difficult to

resolve. Analytic methods that retain detail relating neurons and their spikes to the information they carry will facilitate an understanding of the network mechanisms supporting meaningful dynamics.

Temporally structured population activity is an expedient mechanism for information representation, particularly for time-varying stimuli. Temporal spike patterns provide a powerful computational substrate (Crowe, Averbeck, and Chafee 2010; Villa et al. 1999; Rabinovich, Huerta, and Laurent 2008; Wehr and Laurent 1996), but the extent to which local cortical networks can support reliable spike propagation is an open question (Diesmann, Gewaltig, and Aertsen 1999). Neural networks often exhibit a relatively consistent advancement of activity, summarized as a mean global progression (MGP) of activity (Fig II-1E; Harvey, Coen, and Tank 2012; Luczak et al. 2007; Luczak, McNaughton, and Harris 2015; Pastalkova et al. 2008), suggesting neocortical networks have the capacity for reliable spike propagation. The MGP, revealed after sorting neurons by their mean spike time, is consistent with a Hebbian assembly phase sequence (APS) in which functionally related assemblies of neurons propagate activity through the network (Hebb 1961). Cell assemblies facilitate the propagation of spiking, providing an important bridge between the monosynaptic building blocks of a cortical network and trajectories within population dynamics. While spatiotemporal patterns are a robust feature of cortical network activity, neurons are noisy, and consequently, MGPs generally exhibit large trail-to-trial variability (Lin et al. 2015; Luczak, Barthó, and Harris 2009). The structure of this variability is understudied. Here we sought to identify subpopulations of the network exhibiting reliable spike propagation due to shared spike-time variability.

In order to study reliable spike propagation, we employed two-photon calcium imaging to record spontaneous population activity in large neocortical networks *ex vivo.* This preparation allows us to densely record the local cortical network and observe activity as it propagates undisturbed by top-down modulation or bottom-up sensory drive. The spatiotemporal activity patterns resulted in an MGP that was qualitatively similar to those observed *in vivo.* We show that in these networks, the population post-event time histogram (PETH) alone can lead to an MGP. To evaluate the capacity of these networks for structured spike timing that could not be explained by the PETH, we identified small ensembles of neurons exhibiting sequential spiking. The large repertoire of spike sequences collectively comprise the majority of all spiking in the network. Viewing network activity from the perspective of these spike sequences helped to capture trial-to-trial variability. Sequences densely manifested within the population dynamics and strongly overlapped in time and in neuronal membership, leading to interwoven ensembles of sequences that characterized each population event. Importantly, the occurrence of any given sequence is not deterministic, and the full ensemble of sequences in any trial is unique. From the perspective of an APS, our results suggest that rather than a concatenated, feedforward group of assemblies, the dynamics are characterized by interwoven assemblies. Structured spatiotemporal activity patterns can encode information and compute (Rabinovich, Huerta, and Laurent 2008; Durstewitz and Deco 2008; Buonomano and Maass 2009), suggesting that spike sequences may underlie aspects of cortical function, including sensory representations.

Materials and Methods

Experimental preparation

All procedures were performed in accordance and approved by the Institutional Animal Care and Use Committee at the University of Chicago. Data analyzed here was previously reported in Sadovsky and MacLean (2013). C57BL/6 mice of either sex, P14-17, were anesthetized by intraperitoneal ketamine-xylazine injection. After confirming anesthetic level, mice were rapidly decapitated and had their brains removed and placed in ice-cold cutting solution (containing in mM: 3 KCl, 26 NaHCO3, 1NaH2PO4, 0.5 CaCl2, 3.5 MgSO4, 25 dextrose, 123 sucrose) bubbled with carbogen (95% O2, 5% CO2). Coronal and thalamocortical slices of primary auditory cortex (described in Cruikshank, Rose, and Metherate 2002), 500um thick, were cut using a vibratome (VT1000S; Leica). These slices were then incubated in 35°C oxygenated artificial CSF (containing in mM: 123 NaCl, 3 KCl, 26 NaHCO3, 1 NaH2PO4, 2 CaCl2, 6 MgSO4, 25 dextrose) for 35-40 minutes. Slices were then transferred into a small Petri dish containing 2ml ACSF with an aliquot of 50g Fura-2AM (Invitrogen), 2L Pluronic F-127 (Invitrogen), and 13L DMSO (Sadovsky and MacLean 2013).

Data acquisition

Data was acquired with slices submerged in standard ACSF (containing in mM: 123 NaCl, 3 KCl, 26 NaHCO3, 1 NaH2PO4, 2 CaCl2, 2 MgSO4, and 25 dextrose) which was continuously aerated with carbogen (95% O2, 5% CO2). Primary auditory cortex was identified visually by internal capsule landmarks under brightfield, and two-photon scanning microscopy was performed using the Heuristically Optimal Path Scanning technique (Sadovsky et al. 2011)

resulting in imaging frames lasting 78±14.5ms. Following data acquisition, action potentials were inferred from fluorescence traces using a modified fast non-negative inference algorithm (Vogelstein et al. 2010). Datasets with fewer than four identified spontaneous populations events, and neurons active in fewer than two events were excluded from analysis.

Experimental Design and Statistical Analysis

Calcium imaging was performed on 11 mice (p14-17; both sexes). Data from all experiments were included in all analyses. Coronal (n=5) and thalamocortical (n=6) slices were pooled as no statistical differences were found between datasets, as described previously (Sadovsky and MacLean 2013). All statistical tests were performed with MATLAB (MathWorks). Unless otherwise noted, data are shown as mean±std. Significance of correlation coefficients were computed using a Student's t statistic. A difference in the means of Figure II-6C was confirmed using a two-way ANOVA (spike category and animal identity). All other tests assumed nonparametric distributions and used Mann-Whitney U.

Normalized count

Normalized lagged-spiking counts were computed for every pair of neurons as

$$\frac{\sum_{n=1}^{N}\sum_{t=1}^{T-1} x_i(t)\cdot x_j(t+1)}{\sum_{n=1}^{N}\sum_{t=1}^{T} x_i(t)}$$

Where $x_i(t)$ is the binary spike train of neuron *i,* for each event *n* and each time frame *t*. These lagged spike counts quantify the probability of cell *j* spiking after cell *i*. The matrix of normalized-counts, between pairs of neurons, served to identify potential sequences before testing for significance (see Fig II-2A).

Candidate sequence construction

Candidate sequences are constructed from the normalized lagged-spike count matrix. Each sequence length is defined beforehand. We tested 1000 candidate sequences for sequence lengths of 3 through 8. For each sequence length, pairs of neurons were deterministically drawn from the highest values in the normalized-count matrix. The first pair was drawn starting from the highest value within the entire matrix. A third neuron was added by choosing the neuron with the highest value among all the (2 , i) values, excluding the neuron in position 1. Iteratively, neuron n was chosen by the largest (n-1 , i) values, excluding neurons 1,…,n-1. This guarantees that no neuron repeats within a single candidate sequence. The procedure ends once the predetermined sequence length is reached. To maximize the average normalized count within the candidate sequences, the next candidate was created starting from position k, where the maximum values among all (k , i) is greater than the maximum value among all other starting positions. In this way, the next candidate can branch off the previous candidate, and they would share neurons 1 through k.

When generating random sequences, normalized counts were ignored, but neurons were again not allowed to repeat within individual candidate sequences. After generating surrogate datasets of inhomogeneous Poisson spiking or trial-shuffled spike trains, a new normalized count matrix was computed, and the identical candidate generation method was applied.

Sequence score

Each candidate sequence was characterized by its sequence score (*S*), roughly equal to the fraction of spiking in sequence relative to all spiking in the sequence neurons. First, the binary spike rasters of the candidate sequence are transformed by a term frequency-inverse document frequency (*TF-IDF*) weighting (Carrillo-Reid et al. 2015). The binary spike then becomes the product of its term-frequency weighting and its inverse-document-frequency weighting:

$$W_{c,e}(t) = w_{c,e}^{tf}(t) \cdot w_{c,e}^{idf}(t)$$

With term-frequency weighting

$$w_{c,e}^{tf}(t) = log\left(\frac{f_e}{1 + n_{c,e}}\right) \cdot X_{c,e}(t)$$

And the document-frequency weighting

$$w_{c,e}^{idf}(t) = log\left(1 + \frac{1}{E}\sum_{e=1}^{E} H\left(n_{c,e}\right)\right) \cdot X_{c,e}(t)$$

Where $X_{c,e}(t)$ is the binary spike train of cell *c* in event *e*, f_e is the number of frames in event *e,* $n_{c,e}$ is the number of spikes in cell *c* in event *e*, *E* is the total number of events, and *H(x)* is the Heaviside function, with the convention H(0) = 0. Term-frequency weighting decreases the weight of a spike when a given neuron spikes multiple times within an event (i.e. it may not participate in sequences at a unique time-point). Inverse-document weighting increases the weight of all spikes in an event when the neuron is active in many different events (i.e. it is more likely to be reliably active in spike sequences).

To compute the sequence score, we first convolve each $W_{c,e}(t)$ with a 5 frame Hamming window. Then, the sequence score (S) and sequence onset time (T_e) is defined as

$$T_e = \underset{1 \le t \le f_e}{argmax}\left(\sum_{c=1}^{L} W_{c,e}(t + c - 1)\right)$$

$$S = \frac{1}{E}\sum_{e=1}^{E} s_e$$

$$s_e = \frac{\sum_{c=1}^{L} W_{c,e}(T_e + c - 1)}{\sum_t \sum_{c=1}^{L} W_{c,e}(t)}$$

Where L is the sequence length. These equations define the sequence score as the mean fraction of weighted spikes occurring within the identified sequence window. The sequence window is offset by one frame for each subsequent cell in the sequence, and its location is simply that which gives the maximum within-event sequence score.

Surrogate spike trains are then constructed for each candidate sequence. The number of spikes per event is maintained for each cell. In each event, the spike times of all candidate cells are pooled to estimate the multi-unit spike-time histogram. New spikes are randomly drawn from this distribution for each cell until it has the same number of spikes in the original data. This is done independently for each event, until a complete set of new spike trains is generated. This procedure is repeated one thousand times, leading to one thousand surrogate sequence scores. This distribution of scores represents how sequential we can expect this set of neurons to behave given their coactivity and spike timing. Each candidate sequence is accepted as significant if its own sequence score is greater than 99% of the surrogate sequence scores.

MGP analysis

The MGP is computed by sorting the neurons by their mean spike time across all trials. The percent of spiking within the MGP was calculated by the number of spikes within a spike window centered at each neuron's mean spike time. All spikes from all trials were convolved with a 60ms Gaussian window to generate a time-varying firing rate. To define whether a spike

occurred within the MGP or not, an MGP window was defined as the nearest local minimum before and after each neuron's location in the MGP (i.e. its mean spike time). Surrogate MGPs were computed from inhomogeneous Poisson populations for each dataset 40 times, one example shown in Fig II-1E.

Sequence analysis

Only first spike times were included to compute the mean and variance of spike times. Cells that were active in fewer than four events were excluded. Spikes were classified as sequence spikes or non-sequence spikes. According to the cell identities, ordering, and event-specific sequence window of each sequence, we expect to see spikes occurring at specific time points in specific neurons. Thus, a sequence spike is a spike that satisfies this condition (i.e. it spikes in the sequence window). Each spike was compared against the expected spike timing of every sequence the given neuron participated in. If the spike matched the expected spike time of any sequences, it was classified as a sequence spike; if it matched none, it was a non-sequence spike. All other analyses are described in Results.

Results

Neocortical dynamics exhibit structured progression of neuronal activity

In order to understand the intrinsic properties of network activity generated by local circuitry in neocortex, we chose an *ex vivo* preparation (Fig II-1A). This approach allowed us to study the statistical features of intrinsic dynamics that are exclusively produced by local circuit connectivity, excluding long-range top down and bottom up influences. Population activity in

this preparation emerged spontaneously as discrete events of coordinated network activity resulting in elevated membrane potential across the population and spiking in a subset of the population (Cossart, Aronov, and Yuste 2003; Sadovsky and MacLean 2013; Kruskal, Li, and MacLean 2013). This activity was brief (1.28±0.46 seconds) and was separated by long periods of quiescence, making each period of activity a distinct observation of spike propagation through the local network (Fig II-1B). Using two-photon calcium imaging, we recorded 9.4±4.6 of these events from large populations (595±101 neurons) sampling activity in the population at 12.8±2.0 Hz per slice preparation (n=11). We used a fast nonnegative deconvolution algorithm (Sadovsky and MacLean 2013; Vogelstein et al. 2010) to infer spikes from calcium transients in each neuron. Across the imaged neuronal population, spiking was relatively sparse, encompassing 63±25% of the population with active neurons spiking 1.5±0.8 times per event (Fig II-1B). On average, the population activity followed a common progression: activity in a small number of neurons rapidly recruiting activity in other neurons reaching a maximum population firing rate after 512.5±232.2ms (Fig II-1C). Afterwards, spiking gradually decayed until the end of the event. The mean number of active cells per frame was well fit by a difference of exponentials (R^2 = 0.94). Sorting neurons by their mean spike time revealed mean global progression (MGP) of activity (Fig II-1E), consistent with many reports (Harvey, Coen, and Tank 2012; Luczak, Barthó, and Harris 2009; Luczak, McNaughton, and Harris 2015; Churchland et al. 2012). Alongside this mean progression of activity was a high degree of trial-to-trial variance. At the level of individual cells, variance was correlated with the mean first spike time, resulting in an accumulation of variance throughout the event, as previously reported (Luczak, McNaughton, and Harris 2015). The few neurons that initiated activity had more reliable latencies, while

neurons that spiked later did so more variably ($r=0.51$, $p=2.2*10^{-131}$ Student's t-test; Fig II-1D). We further found that the population PETH was sufficient to replicate the MGP. We generated surrogate spike trains with a population of Poisson processes with inhomogeneous firing rates equal to the PETH. By conserving the number of spikes in each neuron and the mean PETH, the distribution of mean spike times was replicated and in turn generated a qualitatively similar MGP (Fig II-1E). The ordering of the MGP was different, but the percent of spiking within the MGP was similar between the data and the Poisson spiking (data: 47.0±10.8%; Poisson 50.8±7.9%; $p=0.148$ Mann-Whitney U) suggesting that the MGP is largely the result of inhomogeneous firing rates across a specific population. There is substantial evidence that MGPs correspond to specific sensory inputs, motor outputs and behavioral choice, but the high variability of the MGP presents specific challenges from a mechanistic account and an information theoretic perspective. As a result, we sought to determine whether we could identify sequential spiking patterns beyond that described by the global mean.

Identifying reliable spike patterns in global assembly dynamics

To study whether population activity supports reliable propagation of spiking beyond the MGP, we sought to identify smaller ensembles of neurons with sequential spike times across trials (Gansel and Singer 2012; Gourévitch and Eggermont 2010; Reyes-Puerta et al. 2014; Abeles and Gerstein 1988). Testing all potential sequences of moderate length in populations of this size is not computational tractable; therefore, we only tested groups that were most likely to generate sequential spiking (Fig II-2A). Since we were searching for sequential spikes between neurons, we generated potential, 'candidate,' sequences using a normalized lagged-spike count matrix and

then tested each candidate sequence for statistical significance. These normalized counts are better conceptualized as a consistent propagation of spikes between two neurons (Villa et al. 1999), rather than simultaneous correlated fluctuations. We systematically identified candidate sequences using the normalized count matrix and extended group size by linking one pair of strongly-correlated neurons to another (Fig II-2A; for more details, see Materials and Methods). All potential sequences were deterministically constructed from the largest elements in the normalized count matrix. We did not allow neurons to repeat in any single sequence to ensure that each neuron had a unique position in the sequence. All candidate sequences were then independently tested to determine whether they exhibited statistically significant sequential spiking.

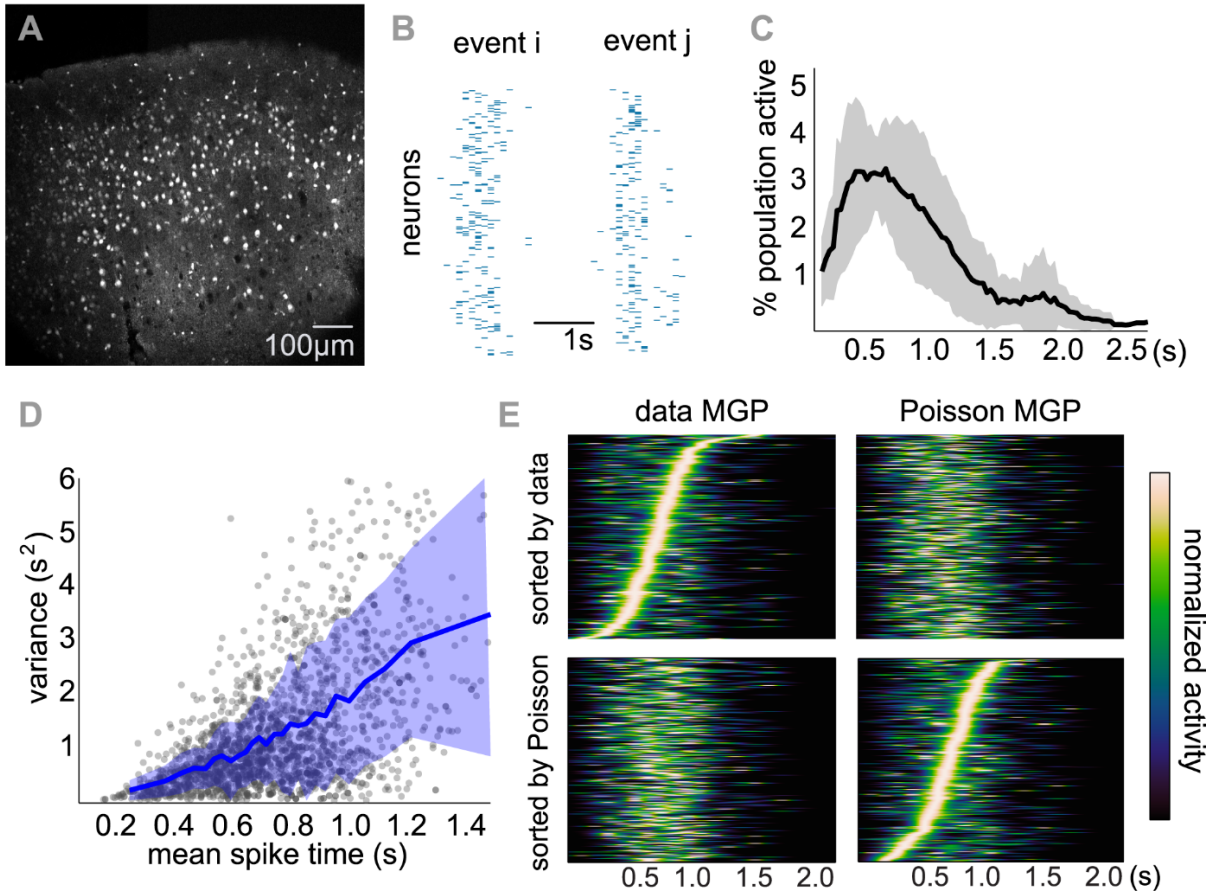

Figure II-1: Structured, emergent dynamics in neocortical populations
A) Representative field of view showing imaged neurons spanning all cortical lamina. B) Two example spontaneous events of activity. C) PETH of population showing a fast rise in activity and slower decay (mean±std). D) Accumulation of variance across the population. Each circle is a neuron (all datasets included; n=11). For each neuron, spikes were pooled across all events aligned to the start of each event. Overlaid running mean±sem in blue. E) Mean global progression is a product of the population PETH; representative MGP from data and MGP generated from an inhomogeneous Poisson population. MGPs in each case have unique ordering but similar structure. Normalized activity of neurons sorted by their mean spike time; spikes convolved with a 60 ms Gaussian window for visualization.

To determine whether each group of neurons showed reliable, sequential spiking, we used a novel metric to quantify a group of neurons' sequential firing (Fig II-2). First, spike trains are weighted to highlight spikes more likely to participate in sequences; the metric is then computed

as the fraction of weighted spiking occurring in sequence. Binary spike matrices are weighted according to a *term frequency - inverse document frequency* (TF-IDF) function (Carrillo-Reid et al. 2015). This weighting procedure makes sparse, binary spike trains more amenable to a continuous, rather than discrete, metric of reliability. Briefly, the weighting of a spike was increased when a neuron had few spikes per event and when a neuron spiked in many different events. This served to more strongly weight neurons that were reliably active across events and that spiked at a unique time point within trials.

We then computed a sequence score for each candidate sequence, measured as the in-sequence weight divided by the total weight in the event (i.e. the fraction of sequence spiking compared to all spiking). In-sequence weight was determined by a sequential time window characterizing the expected spike pattern (Fig II-2A), positioned to maximize the sequence score. To allow small jitter in the sequences, TF-IDF weights were convolved with a 5-frame Hamming window so that presence in the window is not all-or-none. Thus, a neuron not spiking was less costly than a spike out of sequence. Due to the temporal resolution of calcium imaging in these experiments, these sequences cannot be simply the manifestation of a chain of monosynaptic connections (Chambers and MacLean 2015), but rather represent network-dependent propagation of spikes (Chambers and MacLean 2016).

To establish statistical significance, we then compared the sequence score to a distribution of scores using surrogate data. To construct surrogate data, we drew random spike times from the spike time distribution of the candidate cells while maintaining the number of spikes in each cell

in each event. Each spike train of the surrogate data has the same distribution of spike times across trials as its corresponding neuron, maintaining its position in the overall MGP. The surrogate data contains the spike features necessary for a population-wide MGP; consequently, differences between real data and surrogate data were due to specific spike-timing relationships between the candidate neurons themselves. We constructed 1000 surrogate spike trains for each neuron in the candidate sequence and computed a surrogate sequence score for each set of spike trains. If the true sequence score was greater than 99% of the surrogate sequence scores (i.e. $p < 0.01$), we rejected the null hypothesis that the degree of sequential spiking (i.e. sequence score) occurred by chance. Our approach rejected the vast majority of candidate sequences (85.2%±10.1%). Despite this stringent statistical test, we identified many significant sequences after 1000 candidates from each sequence length to generate a suitable set for quantitative analysis (69.7±69.5 sequences of length 3, 132.8±121.3 length 4, 182.2±126.1 length 5, 202.2±108.3 length 6, 177.1±92.5 length 7, 123.0±72.3 length 8; Fig II-2C). We did not include any sequences of lengths greater than 8 because the candidate sequences had a rejection rate of 95.0±6.1%, which was not likely to provide enough statistical power for quantitative analysis. This result suggests a clear limit in the capacity of the local network to support reliable spontaneous sequences on its own. In addition, the total duration of an event acted as an upper bound by limiting the likelihood that we would find a sequence of length nine or longer. Although the significance of a sequence was dependent both on the sequence score and its surrogate distribution, we found that sequences with higher sequence scores had lower p-values and were thus more likely to be significant ($r=-0.43$, $p=0$ Student's t-test).

The sizeable number of sequences with large (6-8) numbers of neurons indicated that there is a high degree of shared variability among large groups of neurons in the population, despite the fact that single-cell spike times are highly variable after even a few frames. Illustrative examples of a sequence of each length in a single event show qualitative results of the identification analysis (Fig II-2B). Due to the design of the identification method, missing spikes and misaligned spikes were common in individual events. When we pooled all spikes from all events (aligned to the occurrence of each sequence), we found a high degree of precision within the neurons with relatively low alignment errors and longer sequences tended to incur more errors than shorter sequences (Fig II-2B).

We performed three control tests for the identification procedure. First, we generated surrogate data in a similar manner to the statistical test, but instead used the entire population raster as the spike time distribution, rather than only using spikes from the candidate sequence neurons (see Materials and Methods). These data are identical to those used in Figure II-1E. The correlation structure in these surrogate spike trains is wholly generated from the population post event time histogram (PETH), removing all influences of intrinsic reliability, co-variability, and higher-order correlations. Thus, one should expect no significant sequential structure from these data. Indeed, the acceptance rate for this control was near the p-value threshold of the statistical test (0.4±0.8%; Fig II-2C, right). As a second control, we created surrogate data by shuffling each neuron's spike trains across trials. This control will destroy sequences that strongly vary in their onset time, but sequences with a reliable onset latency will remain. Indeed, this second class of surrogate data contained fewer significant sequences than the data (147.8±106.8 data; 62.2±52.6

shuffle; Fig II-2C middle-left). Informed by this control, we estimated that 42% of sequences had a consistent onset latency. Finally, we tested the efficacy of the candidate construction method by randomly linking neurons to build candidate sequences. This method was extremely computationally inefficient at identifying sequences (Fig II-2C, middle-right). Constructing candidate sequences by chaining normalized count relationships together lead to a 43-fold increase in the probability of finding a significant sequence. Because of the combinatorial nature of sequences, leveraging the normalized count matrix to generate candidate sequences was necessary for practical tractability.

Due to the deterministic sequence construction and the need for explicit sequence length specification, we found that many of the shortest sequences were themselves subsets of the longer sequences (Fig II-3C). To determine whether long sequences were only comprised of shorter significant sequences we also tested every possible subsequence from the original significant sequences. We found that some, but not all, subsequences were significant (37±23.4%). Longer subsequences were more likely to be significant than shorter subsequences (10.8±8.8% length 3, 24.2±15.5% length 4, 37.9±16.8% length 5, 51.3±17.1% length 6, 61.0±16.6% length 7). Broadly, however, significant sequential spiking within one given group of neurons does not guarantee sequential spiking with a similar group. This suggests that higher-order correlations are responsible for the presence of longer sequences (Chambers and MacLean 2016), and these sequences are not simply composed of chains of pairwise correlations. The abundance of identified sequences motivated a closer analysis of their manifestation in network dynamics.

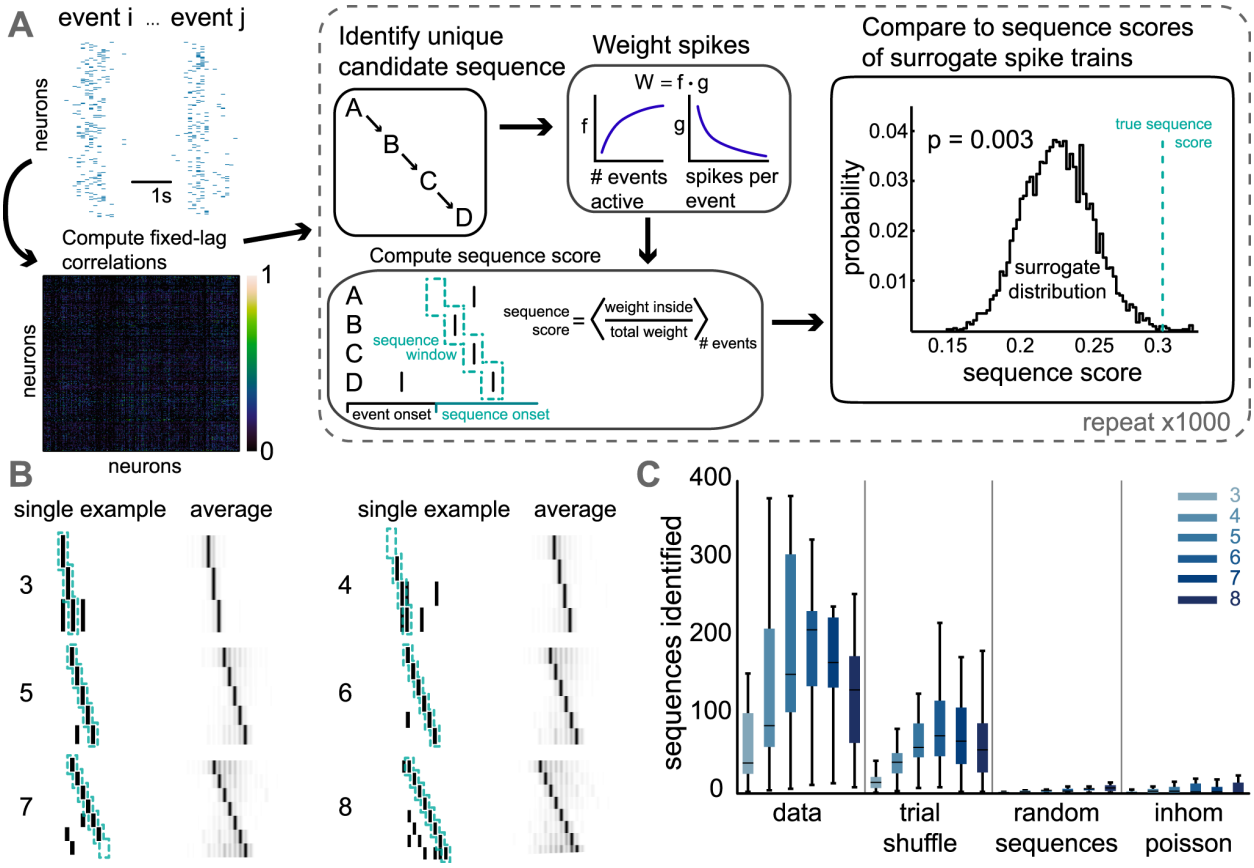

Figure II-2: Identifying spike sequences

A) Illustration of sequence identification method (details in Materials and Methods). Candidate sequences are constructed, and a sequence score is computed from weighted spike trains. The sequence is accepted if score is significantly larger than scores from surrogate spike trains. B) Spike sequences of each length. Single example columns: spike trains from one event (sequence spikes outlined in green); average columns: normalized spike weight across all sequences and all events showing strong sequential structure. C) Number of sequences identified across all datasets (data); in trial-shuffled datasets (trial shuffle); in surrogate spikes replicating population PETH (shuffle); sequences identified with randomly sorted candidate sequences (random). Boxes show median and interquartile range. Whiskers extend to $2.7 \cdot \sigma \cdot \mathrm{IQR}$.

Neuronal participation within sequences

Neurons were allowed to have membership in multiple sequences, and among the significant sequences, we found a long-tailed distribution of sequence membership across neurons. Most neurons participated in a few sequences, and a few neurons participated in a large proportion of all sequences. Surprisingly, the neurons with highest sequence membership were also those with the latest average sequence position (Fig II-3A). Taken together, this indicated that on average sequences begin in a diversity of neurons, and as sequences grew longer, they pool into fewer and fewer neurons. The broad distribution of sequence membership, spanning two orders of magnitude emphasizes the amount of convergence at the end of long sequences. This is consistent with the accumulation of spike-time variance throughout events, as additional variability would reduce the group of possible neurons to occur at the end of long sequences. With this in mind, we tested the hypothesis that long sequences are converging onto low variance cells. We needed to control for the accumulation of variance, since neurons at the end of long sequences (i.e. participating in many sequences) are likely to be active later in the events when variance is higher. We separated the distribution of mean spike times into 39 quantiles and computed the z-scored variance of each neuron according to its mean-spike time, eliminating the differences in mean variance and heteroscedasticity (Fig II-3B, inset). We found that neurons with higher sequence membership have lower variance in spike times on average (Fig II-3B). Neurons with more precise spike-times preferentially participated in many sequences. This finding suggests that in order to participate in many sequences, it is beneficial to be locked to the population onset, as sequences with variable onset latencies would show higher spike-time variability.

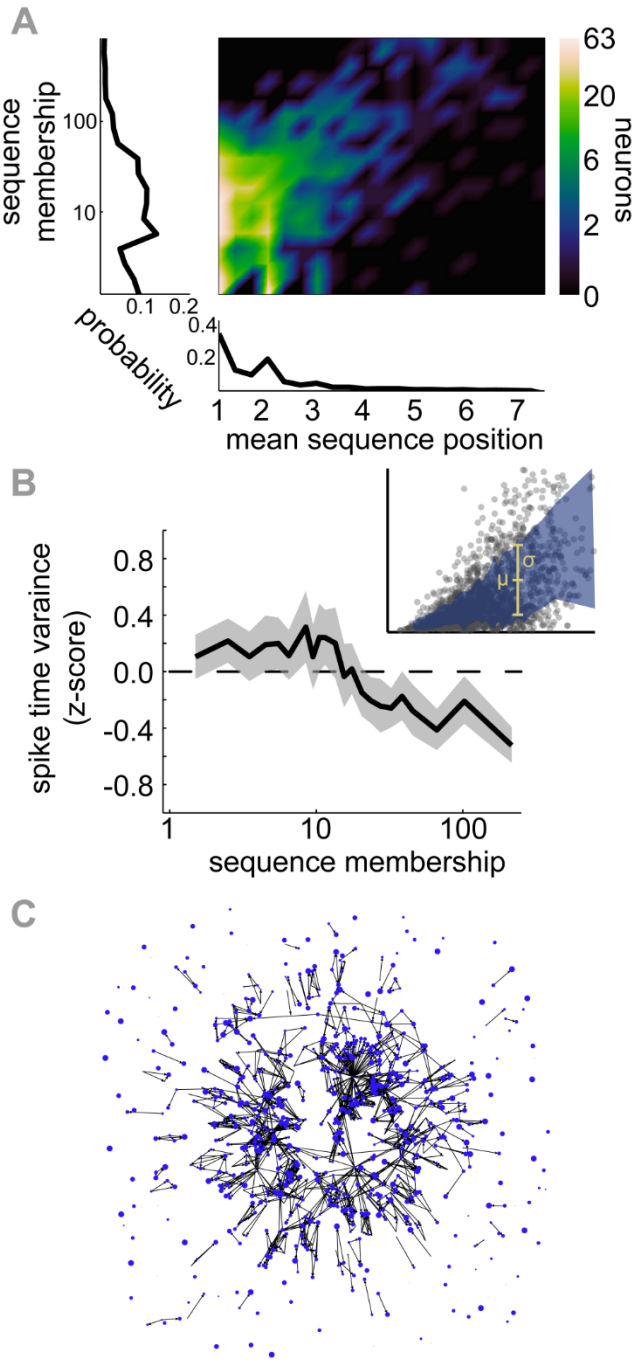

Figure II-3: Sequences converge onto high membership neurons with low variance spiking

A) Sequence membership of a neuron vs mean sequence position (note log scale on ordinate axis). Majority of neurons initiate few sequences, and a minority participate promiscuously at the end of sequences, leading to broad convergence onto a small subset of neurons. Marginal distributions left & bottom. (B) Promiscuous neurons have more precise spike times, relative to event onset. Each neuron variance was z-scored according to its mean-spike time quantile (see inset). Mean variance shown with bootstrapped 95% confidence intervals. (C) Illustrative example of shared neurons between sequences. Directed graph showing sequences that share every neuron (unordered) with another sequence. Each node (n=855) is a sequence (size corresponds to sequence length); Edges lead from one sequence to another in which the target contains all neurons of the source.

Population capacity for spike-time reliability places an upper bound on long, reliable sequences

The accumulation of variance led us to expect sequences to occur near the beginning of events, when spike-time precision was highest. We found that on average longer sequences tended to start later than short sequences relative to the onset of activity, but the onset of shorter sequences was more widely distributed than long sequences (Fig II-4A). The consistently delayed onset of long sequences occurred alongside the peak of population activity (see Fig II-1C), possibly suggesting that a requisite level of excitability is necessary for the initiation of stable spike propagation. Onset times of all sequences showed a similar accumulation of variance to individual cells, and this relationship was independent of sequence length (linear fit slope 3.25 for single cells (see Fig II-1D), 3.40±0.23 across sequence lengths). Thus, sequence onsets reflected the overall population PETH.

While every sequence in this analysis is statistically significant, there is a distribution of sequence scores among significant sequences. We next asked how the reliability of a sequence was impacted by the accumulation of variance by analyzing the relationship between sequence score and mean onset time. We found a strong relationship between sequence score and onset time for long sequences, but short sequences showed no such relationship. For short sequences, reliability could be generated by co-variability within the small group of neurons, and could occur throughout the full duration of the event despite the global accumulation of variability. Conversely, reliability in long sequences was highly dependent on its onset relative to the

population (Fig II-4B). Intrinsic co-variability was not sufficient to support a high degree of reliability for such long sequences.

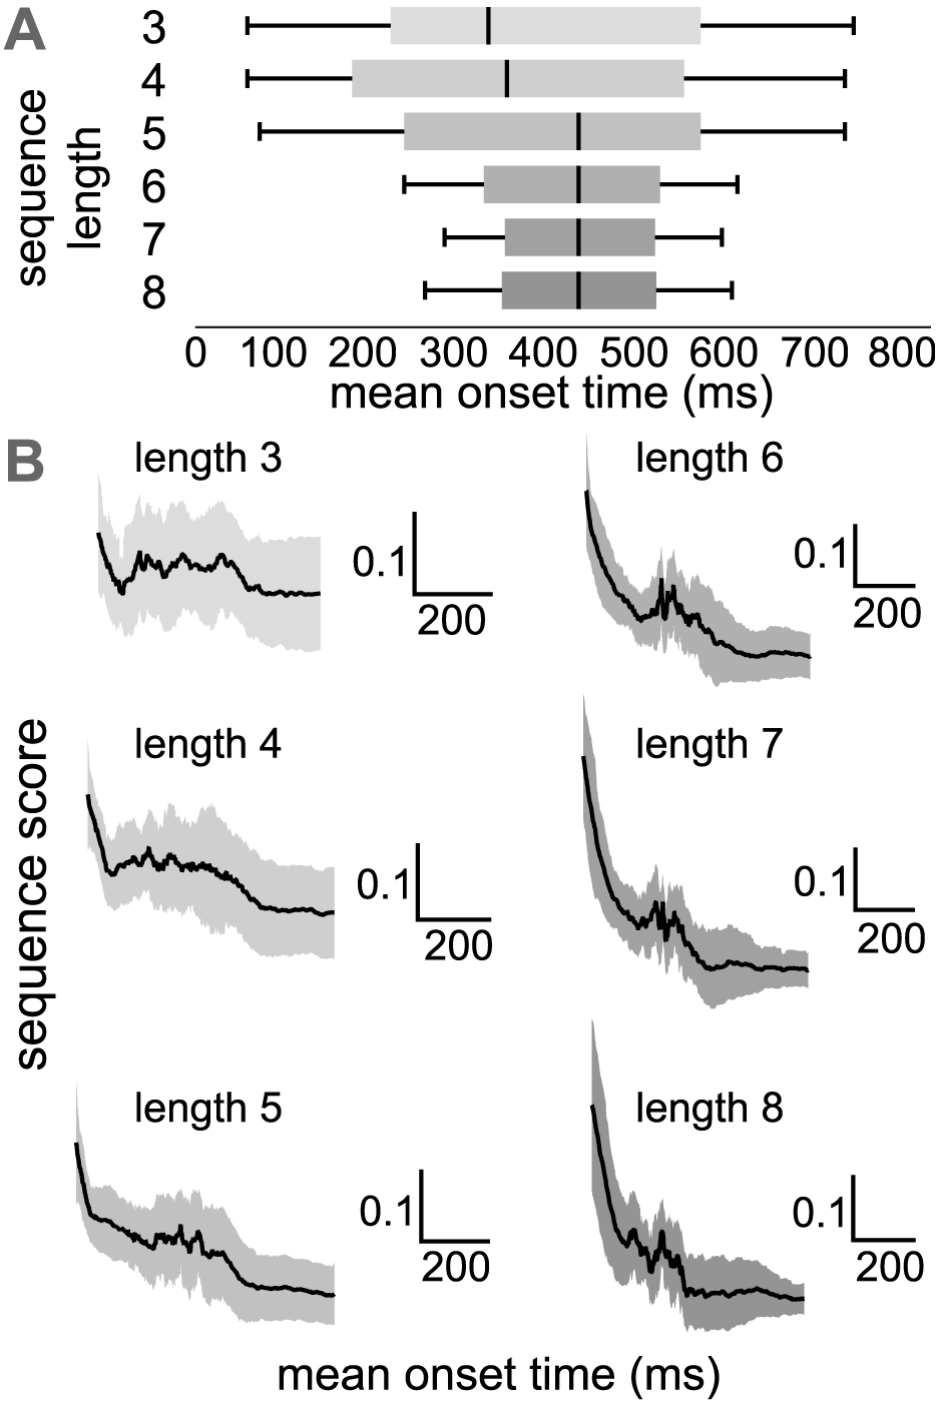

Figure II-4: Long sequences are bounded by the duration and variance of population activity
A) Distributions of onset times for each sequence length. Short sequences occur throughout MGP; long sequences are temporally coupled to peak activity (see Fig II-1C). Only sequences in which 75% percent of neurons were active in an event included. Box and whiskers as in Figure II-2C. B) Running mean±std of sequence score against sequence onset time. Short sequences robustly reliable over time; longer sequences lose reliability in conjunction with accumulation of variability (Fig II-1D).

Sequence composition of global dynamics

We next sought to decompose instances of population activity into spikes captured by sequences, and spikes unexplained by sequences and consequently categorized each spike as a sequence spike or a non-sequence spike. For each neuron, we compared its spike times to all expected spike-times given its sequence membership. If the neuron spiked at the correct time for any sequence, it was categorized as a sequence spike; if the spike belonged to none of its sequences it was categorized as a non-sequence spike.

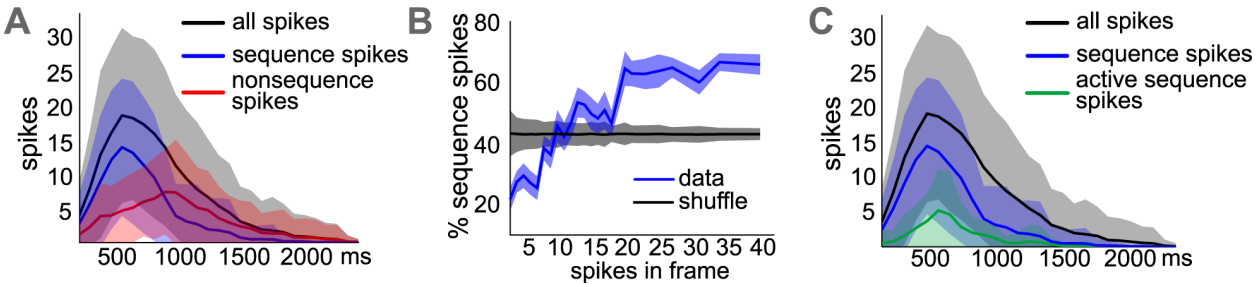

Figure II-5: Sequences comprise a large fraction of population activity
A) Total population spiking, and spiking categorized as either sequence or non-sequence spiking (mean+/-std). B) Sequences track population activity; more spikes belong to sequences when more neurons are spiking (mean±sem). Uniform probability of sequence spikes was generated by shuffling spikes in each event. C) Active sequences, within any one event, are less common, yet show similar evolution in time as the average of events (mean±std).

On average, 55.1% of all spikes belonged to an identified and statistically significant sequence. The progression of sequence spikes, non-sequence spikes, and total spikes over time are shown in Figure II-5A. On average, the number of sequence spikes is highly correlated with the total number of spikes. Non-sequence spikes were fewer in number and tended to occur later. This is presumably because of increased spike variance later in events and a bias for sequences to start early in an event (see Fig II-4A). We also found that the more coincident spikes in any frame, the higher fraction of them were sequence spikes ($r=0.37$, $p=3.49*10^{-42}$ Student's t-test; Fig II-5B).

In other words, increased excitability within the population contained more sequences, while the sparse activity was composed of very little sequential structure. The moderate correlation coefficient and the concave appearance in Figure II-5B suggests that this is predominately a nonlinear relationship. We next restricted our analyses further by only considering sequences that fully manifested during an event (i.e. every cell in the sequence produced a sequence spike). Complete sequences were deemed 'active,' while all other sequences were 'inactive.' Generally, the presence of active sequences, according to this definition, was uncommon (27.0±26.3% length 3, 14.5±17.8% length 4, 7.2±10.8% length 5, 3.5±6.6% length 6, 1.4±4.1% length 7, 0.3±1.0% length 8). Unsurprisingly, active shorter sequences on average occurred more often than active longer sequences. Though only a small portion of the population activity was explained by active sequences, its progression was closely tied to the full population (Fig II-5C).

Single trial manifestations of active sequences

We next studied how active sequences were coordinated during events of population activity. We found distinct but overlapping sequences defined population spiking in single trials (see examples in Fig II-6A). As expected by the distribution of sequence membership, the sequences were highly interwoven with many shared neurons among active sequences. To quantify the degree of shared sequences between events, we computed the Jaccardian index between active sequence spikes in all pairs of events. We found that the median sequence overlap between pairs of events was low (33.7%), and many pairs of events shared nearly no sequences spikes (Fig II-6B). This suggests that while all trials share the same MGP and repertoire of possible sequences,

they are differentiated by their unique ensembles of active sequences. We then computed the absolute deviation from a neuron's mean spike time (i.e. its position in the MGP) for all of its sequence spikes and non-sequence spikes separately. We found that sequence spikes were more similar to the MGP than non-sequence spikes (Fig II-6C) consistent with our result that sequence spikes tracked population activity (see Fig II-5A).

We then tested whether the convergence of all sequences onto a small subset of high membership neurons was also present within single instantiations of active sequences. First, analyzing all sequences aligned to their first neuron (as in Fig II-3A), we computed the number of unique neurons at each sequence position, normalized by the expected number of unique neurons. This expectation was calculated assuming a uniform probability of observing each neuron spiking in sequence. The ratio of these two quantities told us how often neurons repeat in sequences as compared to expectations set by chance. Consistent with Figure II-3A, we found relatively diverse neurons at the beginning of sequences, but high redundancy later in sequences. Next, we exclusively analyzed active sequences within individual events, computing the number of unique cells participating in each frame, normalized by the expected number of unique neurons assuming a uniform distribution. This resulted in a comparable measure of how often neurons repeat in active sequences. We found a constant ratio of unique neurons throughout population activity (Fig II-6D). Due to the specific subsets of active sequences and the distribution of onset times, the overall convergence seen among all sequences was suppressed. Among the overlapping subsets of active sequences, the reliable propagation of spiking was balanced between convergence onto shared neurons and divergence into distinct neurons. Thus,

sequences were highly interwoven and co-occurred throughout the duration of population activity. While on average sequences pooled into a small subset of individual neurons, on single trials this was not the case. Interwoven sequences during single events were drawn from a large repertoire of possible sequences and demonstrated that variance about the global mean is not noise, but the manifestation of reliable spike times in the population.

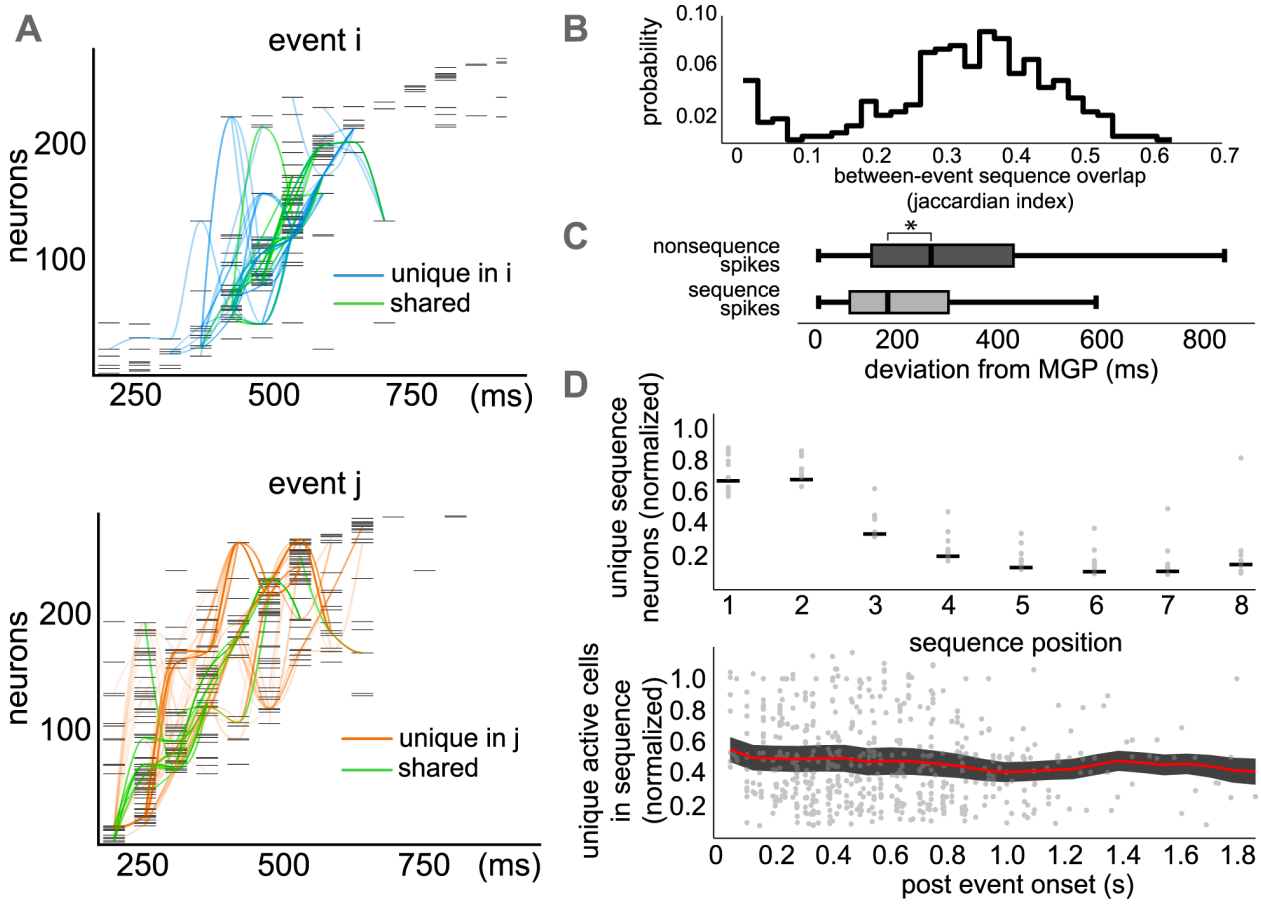

Figure II-6: Single trails of MGP are composed of unique subsets of interwoven sequences
A) Two example rasters with overlaid active sequences. Neurons sorted according to MGP; events exhibit unique (left: blue; right: orange) and overlapping (green) sequences. B) Distribution of shared sequence spikes between event pairs. Low overall overlap, with many events sharing nearly no sequences. C) Sequence spikes occur closer to expected MGP position

($p=1.1*10^{-3}$ two-way ANOVA). Box and whiskers as in Figure II-2C. D) All sequences show strong pooling at the end of long sequences (black bar is mean, grey points for each dataset); active sequences show balanced convergence and divergence throughout event (mean+/-sem; grey points show each time point across all events).

Discussion

Consistent trajectories of population activity have been observed in many cortical areas, are qualitatively similar *in vitro* and *in vivo* (Sadovsky and MacLean 2013), and have been shown to correlate with behavioral choice (Harvey, Coen, and Tank 2012), motor output (Churchland et al. 2012), and sensory inputs (Luczak, Barthó, and Harris 2009). Our analyses provide a more detailed quantitative description of population spike propagation. Across the population, sorting neurons by their mean spike latency reveals the MGP. However, the population PETH is sufficient to generate an MGP in these data indicating that a more thorough description of activity is necessary. The most prevalent theoretical framework underlying spatiotemporal transmission of spikes is the Hebbian assembly phase sequence (APS). While an APS may be implemented in different forms, a common interpretation of Hebb's seminal work is that distinct, non-overlapping assemblies form a propagating sequence of activity, with one assembly recruiting the next, resulting in a temporal progression that reflects a common function (Fig II-7A). An assembly composed of well-defined groups of functionally related neurons is consistent across many disciplines (Litwin-Kumar and Doiron 2014; O'Neill et al. 2008; Pastalkova et al. 2008; Carrillo-Reid et al. 2015; Levy et al. 2001; Wehr and Laurent 1996). Until now, however, there has been little beyond this qualitative description of how the phase sequences manifest within local cortical populations comprised of interconnected neurons. Our data and analytical approach motivate an alternate formulation of the Hebbian APS. Naively, neurons with similar

onset latencies might be clustered into distinct assemblies. However, due to shared variability within subpopulations of neurons, our results suggest that the fundamental unit of the APS is the temporal spike sequence, rather than the active group of neurons at one specific time. As a consequence of sequences highly overlapping in time and in neuron membership, assemblies are then defined by co-occurring sequences. Different ensembles of sequences manifest in single trials of population activity, resulting in variable sequence co-activity and assemblies. Identifying the repertoire of sequences allowed for description of unique sequence ensembles, and helped to explain the large variation about the MGP. Each event is uniquely described by its ensemble of sequences, suggesting that the APS is a propagation of spikes through an interwoven subset of sequences among a common repertoire. As a consequence of sequences with variable onset time and probabilistic sequence activation, different assemblies manifest in each trial, but the MGP and the repertoire of sequences is consistent. We conclude that the MGP is the result of averaging across these interwoven micro-assemblies, rather than a single, concatenated chain of assemblies (Fig II-7B).

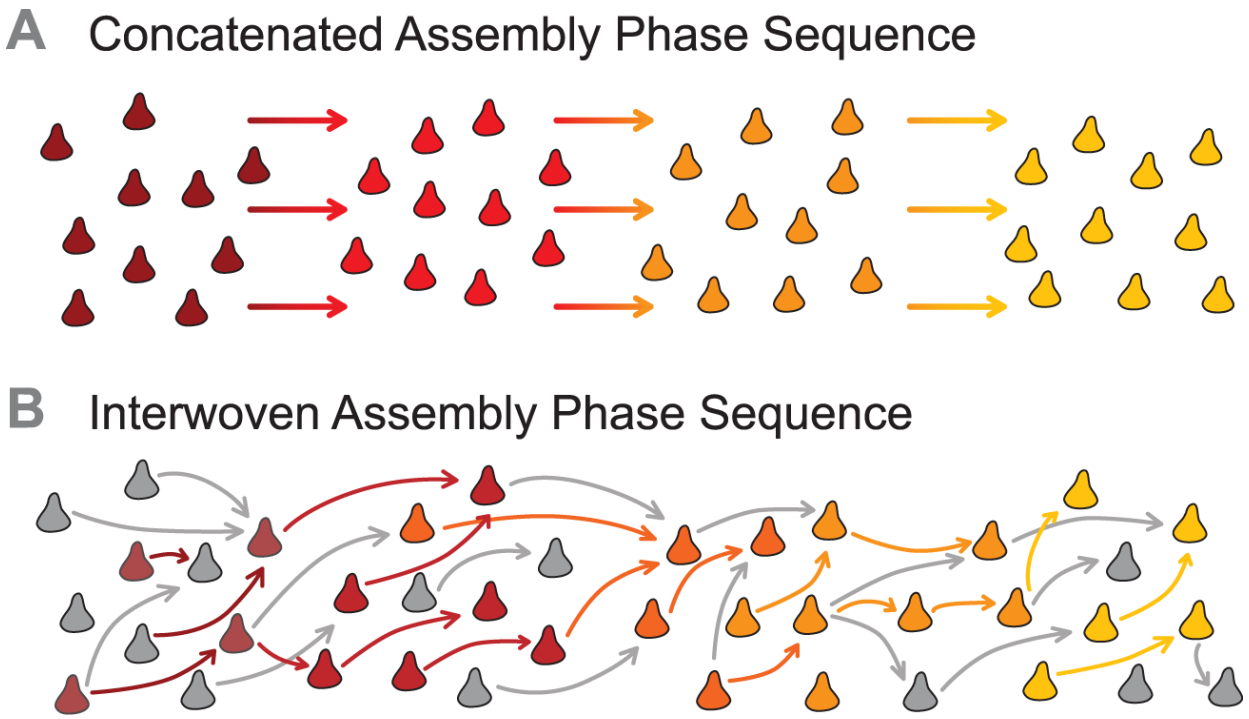

Figure II-7: Sequence composition of MGP suggests an interwoven assembly phase sequence framework
A) Distinct, non-overlapping assemblies (neuron color) underlies propagation of activity. Traditionally, the hypothesis is interpreted as dynamics being driven by a concatenated, feedforward assembly phase sequence. Assemblies are defined by the group of coactive neurons in a time window and noise intrinsic to each assembly leads to trial-to-trial variability. B) In our suggested revision, based on our data analysis, neurons belong to many assemblies, and probabilistic activation of sequences underlies propagation of activity. Neuron color denotes single-trial assembly membership, arrows denote underlying sequences; grey neurons and arrows are inactive in this trial. Assemblies are defined by the ensemble of active sequences.

Previous attempts to quantify temporally structured spiking in local populations of neurons have had mixed results (Abeles and Gerstein 1988; Baker and Lemon 2000; Ikegaya et al. 2004; Roxin, Hakim, and Brunel 2008; Mokeichev et al. 2007; Humphries 2011; Torre et al. 2016). The combinatorial nature of spatiotemporal spike sequences poses substantial difficulty in identifying spike patterns that cannot be explained by simpler hypotheses. One can attempt to calculate the contributions from every individual neuron and subsequence, then estimate the

divergence from this expectation in the data (Schneidman et al. 2003; Martignon et al. 2000). These approaches require large amounts of data which increases non-linearly as pattern length increases. A more common approach, taken here, is to compare patterns in the data to surrogate data. Features built into the surrogate data delineate the null hypothesis, and patterns unique to the data are taken as evidence against the null hypothesis (Lee and Wilson 2004; Grun 2009; Harrison, Amarasingham, and Truccolo 2015). Our analysis here diverges from many other works in three main ways. First, we do not consider synfire-like patterns with strict, millisecond precision (Baker and Lemon 2000; Ikegaya et al. 2004; Lee and Wilson 2004), as our sampling rate and identification method allow for spike jitter. While many sequence-identification methods allow for spike jitter (Gansel and Singer 2012; Torre et al. 2016), they also require pattern completion. The second difference is that the sequences here may have spikes missing in individual trials. Finally, other identification methods have taken an agnostic approach to the specific spike-patterns that could manifest in neural networks (Abeles and Gerstein 1988; Tatsuno, Lipa, and McNaughton 2006; Gansel and Singer 2012); this work exclusively searched for single-frame lagged spike sequences which have a clear correspondence with APS's and help to better understand how subpopulations contributed to the population-wide MGP. While this assumption greatly restricts the number of possible patterns, it also limits the scope of this work to a specific form of a Hebbian APS. Furthermore, we did not allow neurons to repeat in sequences. From the perspective of an MGP, a neuron should have a specific time of activation, but highly recurrent networks may produce sequences that reverberate back to an earlier neuron (Hebb 1961). Due to the data acquisition rate and the assumptions built into the method, the sequences identified here are not a complete sampling of spike sequences in the data. Rather,

they represent the reliable spike sequences underlying the population MGP. The identification method presented here could be extended to allow for variable lags and repeated neurons in sequences, but other methods are better suited for assumption-free analyses of spatiotemporal spike patterns.

Calcium imaging allows recording of neuronal populations with single cell resolution spanning large fields of view (~1000um). Different data acquisition techniques, such as high-density multi-electrode arrays (Berdondini et al. 2009), may be capable of capturing aspects of Hebbian assembly dynamics not observed here, and will likely be necessary to reveal generalizable features of spike sequences beyond those measured in these data. Future explorations of sequences propagating at different intervals will help to further delineate the capacity of cortical networks to produce reliable spatiotemporal patterns of activity.

Trajectories of activity in neural networks has attracted attention from theorists and experimentalists alike, as they are capable of encoding information and performing computations (Crowe, Averbeck, and Chafee 2010; Klampfl and Maass 2013; Harvey, Coen, and Tank 2012). Other computational frameworks have shown that time-varying trajectories of activity in neural populations are a powerful computational substrate (Buonomano and Maass 2009; Rabinovich, Huerta, and Laurent 2008; Durstewitz and Deco 2008). Experimentally, single-trial deviations of the MGP have been hypothesized to encode stimulus identity (Luczak, McNaughton, and Harris 2015). Our results show that the population activity deviations are captured by unique subsets of active sequences, demonstrating that local circuitry is able to generate sequences which in turn

have the capacity to be stimulus-specific. The large repertoire of sequences and their interwoven propagation could even allow for simultaneous encoding of external variables, internal cortical state, and their interactions similar to a coupled hidden Markov model. Synaptic plasticity rules are capable of embedding interwoven spike sequences in a neural network (Izhikevich 2006), and these sequences show a natural capacity to encode stimulus information (Paugam-Moisy, Martinez, and Bengio 2008). Our analyses summarize complex population dynamics while maintaining individual spikes in specific neurons, leading to a major advantage in interpretation over other methods of representing population activity. Since each sequence is composed of specific neurons, there is a direct biological interpretation of the information represented in sequence activity. In contrast, dimensionality reduction techniques are powerful at harnessing information within population activity, but they project activity onto latent dimensions making it difficult to attribute the activity of specific neurons to the low dimensional projection (Mante et al. 2013; Briggman, Abarbanel, and Kristan 2005; Byron et al. 2009; Broome, Jayaraman, and Laurent 2006). The ability to understand how a specific computation or stimulus representation is implemented by local circuitry necessitates an understanding of each neuron's contribution. In addition, the temporal coding framework of reliably propagating spikes has a clear relationship with the underlying biological substrate. The modulation of activity via synaptic transmission occurs with a delay; therefore, spike propagation has a more direct relationship with network activity than simultaneous modulation of spiking. Neocortex produces both highly variable (Heggelund and Albus 1978; Vogels, Spileers, and Orban 1989), and highly precise (Reich et al. 1997; Wehr and Zador 2003; Bair and Koch 1996) activity. By maintaining spike-time

information and capturing moments of reliable activity, the capacity for information representation is increased (Kayser et al. 2009).

Motivated by observations that neural population activity exhibits a mean global progression, we present a detailed description of the spike sequences underlying this population-wide phenomenon. We decomposed the individual observations of population activity into smaller sequences of spikes that recur across events. Most neurons only participated in a few sequences and tended to initiate sequences. This diversity of initiating neurons converged onto a small subset of neurons that participated in many sequences. Across the population, spiking in sequences was tightly linked to the population firing rate, and each event was comprised of a specific subset of active sequences. The specific subset of active sequences was highly interwoven, sharing neurons and spike timing. Our analyses capture much of the trial-to-trial variability in the spontaneous activity, and suggest that sequences probabilistically converge and diverge. The full set of identified sequences represents the dynamical repertoire of the population, but population activity manifests in the network through unique subsets of these sequences. The probabilistic trajectories of sequences emphasize that neurons belong to multiple, interwoven assemblies, rather than a feedforward cascade of spike propagation. This method of representing population activity is closely tied to the structure and function of neocortical networks and reveals temporally structured activity with the capacity to aid in information representation.

References

Abeles, Moshe, and George L. Gerstein. 1988. "Detecting Spatiotemporal Firing Patterns among Simultaneously Recorded Single Neurons." *Journal of Neurophysiology* 60 (3): 909–924.

Arieli, Amos, Alexander Sterkin, Amiram Grinvald, and A. D. Aertsen. 1996. "Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses." *Science* 273 (5283): 1868–1871.

Bair, Wyeth, and Christof Koch. 1996. "Temporal Precision of Spike Trains in Extrastriate Cortex of the Behaving Macaque Monkey." *Neural Computation* 8 (6): 1185–1202.

Baker, Stuart N., and Roger N. Lemon. 2000. "Precise Spatiotemporal Repeating Patterns in Monkey Primary and Supplementary Motor Areas Occur at Chance Levels." *Journal of Neurophysiology* 84 (4): 1770–1780.

Bathellier, Brice, Lyubov Ushakova, and Simon Rumpel. 2012. "Discrete Neocortical Dynamics Predict Behavioral Categorization of Sounds." *Neuron* 76 (2): 435–449.

Beggs, John M., and Dietmar Plenz. 2004. "Neuronal Avalanches Are Diverse and Precise Activity Patterns That Are Stable for Many Hours in Cortical Slice Cultures." *Journal of Neuroscience* 24 (22): 5216–5229.

Berdondini, Luca, Kilian Imfeld, Alessandro Maccione, Mariateresa Tedesco, Simon Neukom, Milena Koudelka-Hep, and Sergio Martinoia. 2009. "Active Pixel Sensor Array for High Spatio-Temporal Resolution Electrophysiological Recordings from Single Cell to Large Scale Neuronal Networks." *Lab on a Chip* 9 (18): 2644–2651.

Briggman, Kevin L., Henry DI Abarbanel, and William B. Kristan. 2005. "Optical Imaging of Neuronal Populations during Decision-Making." *Science* 307 (5711): 896–901.

Broome, Bede M., Vivek Jayaraman, and Gilles Laurent. 2006. "Encoding and Decoding of Overlapping Odor Sequences." *Neuron* 51 (4): 467–82.

Buonomano, Dean V., and Wolfgang Maass. 2009. "State-Dependent Computations: Spatiotemporal Processing in Cortical Networks." *Nature Reviews Neuroscience* 10 (2): 113.

Byron, M. Yu, John P. Cunningham, Gopal Santhanam, Stephen I. Ryu, Krishna V. Shenoy, and Maneesh Sahani. 2009. "Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity." In *Advances in Neural Information Processing Systems*, 1881–1888.

Carrillo-Reid, Luis, Jae-eun Kang Miller, Jordan P. Hamm, Jesse Jackson, and Rafael Yuste. 2015. “Endogenous Sequential Cortical Activity Evoked by Visual Stimuli.” *Journal of Neuroscience* 35 (23): 8813–8828.

Chambers, Brendan, and Jason N. MacLean. 2015. “Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints.” *Journal of Neurophysiology* 114 (3): 1837–1849.

Chambers, Brendan, and Jason N. MacLean. 2016. “Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks.” *PLoS Computational Biology* 12 (8): e1005078.

Churchland, Mark M., John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. 2012. “Neural Population Dynamics during Reaching.” *Nature* 487 (7405): 51.

Cossart, Rosa, Dmitriy Aronov, and Rafael Yuste. 2003. “Attractor Dynamics of Network UP States in the Neocortex.” *Nature* 423 (6937): 283.

Crowe, David A., Bruno B. Averbeck, and Matthew V. Chafee. 2010. “Rapid Sequences of Population Activity Patterns Dynamically Encode Task-Critical Spatial Information in Parietal Cortex.” *Journal of Neuroscience* 30 (35): 11640–11653.

Cruikshank, Scott J., Heather J. Rose, and Raju Metherate. 2002. “Auditory Thalamocortical Synaptic Transmission in Vitro.” *Journal of Neurophysiology* 87 (1): 361–384.

Diesmann, Markus, Marc-Oliver Gewaltig, and Ad Aertsen. 1999. “Stable Propagation of Synchronous Spiking in Cortical Neural Networks.” *Nature* 402 (6761): 529.

Durstewitz, Daniel, and Gustavo Deco. 2008. “Computational Significance of Transient Dynamics in Cortical Networks.” *European Journal of Neuroscience* 27 (1): 217–227.

Gansel, Kai S., and Wolf Singer. 2012. “Detecting Multineuronal Temporal Patterns in Parallel Spike Trains.” *Frontiers in Neuroinformatics* 6: 18.

Gourévitch, Boris, and Jos J. Eggermont. 2010. “Maximum Decoding Abilities of Temporal Patterns and Synchronized Firings: Application to Auditory Neurons Responding to Click Trains and Amplitude Modulated White Noise.” *Journal of Computational Neuroscience* 29 (1–2): 253–277.

Grun, Sonja. 2009. “Data-Driven Significance Estimation for Precise Spike Correlation.” *Journal of Neurophysiology* 101 (3): 1126–1140.

Harrison, Matthew T., Asohan Amarasingham, and Wilson Truccolo. 2015. “Spatiotemporal Conditional Inference and Hypothesis Tests for Neural Ensemble Spiking Precision.” *Neural Computation* 27 (1): 104–150.

Harvey, Christopher D., Philip Coen, and David W. Tank. 2012. “Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task.” *Nature* 484 (7392): 62.

Hebb, D. O. 1961. *Organization of Behavior. New York: Science Editions*. Inc.
Heggelund, P., and K. Albus. 1978. “Response Variability and Orientation Discrimination of Single Cells in Striate Cortex of Cat.” *Experimental Brain Research* 32 (2): 197–211.

Humphries, Mark D. 2011. “Spike-Train Communities: Finding Groups of Similar Spike Trains.” *Journal of Neuroscience* 31 (6): 2321–2336.

Ikegaya, Yuji, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste. 2004. “Synfire Chains and Cortical Songs: Temporal Modules of Cortical Activity.” *Science* 304 (5670): 559–564.

Izhikevich, Eugene M. 2006. “Polychronization: Computation with Spikes.” *Neural Computation* 18 (2): 245–282.

Johnson, Hope A., Anubhuthi Goel, and Dean V. Buonomano. 2010. “Neural Dynamics of in Vitro Cortical Networks Reflects Experienced Temporal Patterns.” *Nature Neuroscience* 13 (8): 917.

Kayser, Christoph, Marcelo A. Montemurro, Nikos K. Logothetis, and Stefano Panzeri. 2009. “Spike-Phase Coding Boosts and Stabilizes Information Carried by Spatial and Temporal Spike Patterns.” *Neuron* 61 (4): 597–608.

Klampfl, Stefan, and Wolfgang Maass. 2013. “Emergence of Dynamic Memory Traces in Cortical Microcircuit Models through STDP.” *Journal of Neuroscience* 33 (28): 11515–11529.

Kruskal, Peter B., Lucy Li, and Jason N. MacLean. 2013. “Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex.” *Nature Communications* 4: 2574.

Lee, Albert K., and Matthew A. Wilson. 2004. “A Combinatorial Method for Analyzing Sequential Firing Patterns Involving an Arbitrary Number of Neurons Based on Relative Time Order.” *Journal of Neurophysiology* 92 (4): 2555–2573.

Levy, Nir, David Horn, Isaac Meilijson, and Eytan Ruppin. 2001. “Distributed Synchrony in a Cell Assembly of Spiking Neurons.” *Neural Networks* 14 (6–7): 815–824.

Lin, I.-Chun, Michael Okun, Matteo Carandini, and Kenneth D. Harris. 2015. “The Nature of Shared Cortical Variability.” *Neuron* 87 (3): 644–656.

Litwin-Kumar, Ashok, and Brent Doiron. 2014. “Formation and Maintenance of Neuronal Assemblies through Synaptic Plasticity.” *Nature Communications* 5: 5319.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. 2009. “Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations.” *Neuron* 62 (3): 413–425.

Luczak, Artur, Peter Barthó, Stephan L. Marguet, György Buzsáki, and Kenneth D. Harris. 2007. “Sequential Structure of Neocortical Spontaneous Activity in Vivo.” *Proceedings of the National Academy of Sciences* 104 (1): 347–352.

Luczak, Artur, Bruce L. McNaughton, and Kenneth D. Harris. 2015. “Packet-Based Communication in the Cortex.” *Nature Reviews Neuroscience* 16 (12): 745.

Mante, Valerio, David Sussillo, Krishna V. Shenoy, and William T. Newsome. 2013. “Context-Dependent Computation by Recurrent Dynamics in Prefrontal Cortex.” *Nature* 503 (7474): 78.

Martignon, Laura, Gustavo Deco, Kathryn Laskey, Mathew Diamond, Winrich Freiwald, and Eilon Vaadia. 2000. “Neural Coding: Higher-Order Temporal Patterns in the Neurostatistics of Cell Assemblies.” *Neural Computation* 12 (11): 2621–2653.

Mokeichev, Alik, Michael Okun, Omri Barak, Yonatan Katz, Ohad Ben-Shahar, and Ilan Lampl. 2007. “Stochastic Emergence of Repeating Cortical Motifs in Spontaneous Membrane Potential Fluctuations in Vivo.” *Neuron* 53 (3): 413–425.

O’Neill, Joseph, Timothy J. Senior, Kevin Allen, John R. Huxter, and Jozsef Csicsvari. 2008. “Reactivation of Experience-Dependent Cell Assembly Patterns in the Hippocampus.” *Nature Neuroscience* 11 (2): 209.

Pastalkova, Eva, Vladimir Itskov, Asohan Amarasingham, and György Buzsáki. 2008. “Internally Generated Cell Assembly Sequences in the Rat Hippocampus.” *Science* 321 (5894): 1322–1327.

Paugam-Moisy, Hélene, Régis Martinez, and Samy Bengio. 2008. “Delay Learning and Polychronization for Reservoir Computing.” *Neurocomputing* 71 (7–9): 1143–1158.

Peters, Andrew J., Simon X. Chen, and Takaki Komiyama. 2014. “Emergence of Reproducible Spatiotemporal Activity during Motor Learning.” *Nature* 510 (7504): 263.

Rabinovich, Misha, Ramon Huerta, and Gilles Laurent. 2008. “Transient Dynamics for Neural Processing.” *Science*, 48–50.

Reich, Daniel S., Jonathan D. Victor, Bruce W. Knight, Tsuyoshi Ozaki, and Ehud Kaplan. 1997. "Response Variability and Timing Precision of Neuronal Spike Trains in Vivo." *Journal of Neurophysiology* 77 (5): 2836–2841.

Reyes-Puerta, Vicente, Jyh-Jang Sun, Suam Kim, Werner Kilb, and Heiko J. Luhmann. 2014. "Laminar and Columnar Structure of Sensory-Evoked Multineuronal Spike Sequences in Adult Rat Barrel Cortex in Vivo." *Cerebral Cortex* 25 (8): 2001–2021.

Roxin, Alex, Vincent Hakim, and Nicolas Brunel. 2008. "The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons." *Journal of Neuroscience* 28 (42): 10734–10745.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. 2011. "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging." *Journal of Neurophysiology* 106 (3): 1591–1598.

Sadovsky, Alexander J., and Jason N. MacLean. 2013. "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry." *Journal of Neuroscience* 33 (35): 14048–14060.

Schneidman, Elad, Susanne Still, Michael J. Berry, and William Bialek. 2003. "Network Information and Connected Correlations." *Physical Review Letters* 91 (23): 238701.

Tatsuno, Masami, Peter Lipa, and Bruce L. McNaughton. 2006. "Methodological Considerations on the Use of Template Matching to Study Long-Lasting Memory Trace Replay." *Journal of Neuroscience* 26 (42): 10727–10742.

Torre, Emiliano, Carlos Canova, Michael Denker, George Gerstein, Moritz Helias, and Sonja Grün. 2016. "ASSET: Analysis of Sequences of Synchronous Events in Massively Parallel Spike Trains." *PLoS Computational Biology* 12 (7): e1004939.

Villa, Alessandro EP, Igor V. Tetko, Brian Hyland, and Abdellatif Najem. 1999. "Spatiotemporal Activity Patterns of Rat Cortical Neurons Predict Responses in a Conditioned Task." *Proceedings of the National Academy of Sciences* 96 (3): 1106–1111.

Vogels, Rufin, Werner Spileers, and Guy A. Orban. 1989. "The Response Variability of Striate Cortical Neurons in the Behaving Monkey." *Experimental Brain Research* 77 (2): 432–436.

Vogelstein, Joshua T., Adam M. Packer, Timothy A. Machado, Tanya Sippy, Baktash Babadi, Rafael Yuste, and Liam Paninski. 2010. "Fast Nonnegative Deconvolution for Spike Train Inference from Population Calcium Imaging." *Journal of Neurophysiology* 104 (6): 3691–3704.

Wehr, Michael, and Gilles Laurent. 1996. "Odour Encoding by Temporal Sequences of Firing in Oscillating Neural Assemblies." *Nature* 384 (6605): 162.

Wehr, Michael, and Anthony M. Zador. 2003. “Balanced Inhibition Underlies Tuning and Sharpens Spike Timing in Auditory Cortex.” *Nature* 426 (6965): 442.

CHAPTER III

FUNCTIONAL TRIPLET MOTIFS UNDERLIE ACCURATE PREDICTIONS OF SINGLE-TRIAL RESPONSES IN POPULATIONS OF TUNED AND UNTUNED V1 NEURONS

Abstract

Visual stimuli evoke activity in visual cortical neuronal populations. Neuronal activity can be selectively modulated by particular visual stimulus parameters, such as the direction of a moving bar of light, resulting in well-defined trial averaged tuning properties. However, given any single stimulus parameter, a large number of neurons in visual cortex remain unmodulated, and the role of this untuned population is not well understood. Here, we use two-photon calcium imaging to record, in an unbiased manner, from large populations of layer 2/3 excitatory neurons in mouse primary visual cortex to describe co-varying activity on single trials in neuronal populations consisting of both tuned and untuned neurons. Specifically, we summarize pairwise covariability with an asymmetric partial correlation coefficient, allowing us to analyze the resultant population correlation structure, or functional network, with graph theory. Using the graph neighbors of a neuron, we find that the local population, including both tuned and untuned neurons, are able to predict individual neuron activity on a moment to moment basis, while also recapitulating tuning properties of tuned neurons. Variance explained in total population activity scales with the number of neurons imaged, suggesting larger sample sizes are required to fully capture local network interactions. We also find that a specific functional triplet motif in the graph results in the best predictions, suggesting a signature of informative correlations in these populations.

Variance explained in total population activity scales with the number of neurons imaged, suggesting larger sample sizes are required to fully capture local network interactions. In summary, we show that unbiased sampling of the local population can explain single trial response variability as well as trial-averaged tuning properties in V1, and the ability to predict responses is tied to the occurrence of a functional triplet motif.

Introduction

In the visual system, decades of research have probed stimulus parameters that evoke responses in single neurons and these responses have been generally trial-averaged (Hubel and Wiesel 1962). These response properties have revealed principles of functional organization in primary visual cortex (V1), such as orientation columns (Lund, Angelucci, and Bressloff 2003), and canonical computations, such as divisive normalization (Heeger 1992). However, responses are variable across trials (Shadlen and Newsome 1998), making the relationship between perceptual stability and neuronal single-trial stimulus representations unclear (Tolhurst, Movshon, and Dean 1983). The fluctuations of response strength are not independent across neurons, and this shared variability impacts population level representations of visual stimuli (Seriès, Latham, and Pouget 2004; Zohary, Shadlen, and Newsome 1994). Neurons are highly interconnected and connection likelihood is biased toward spatially proximal neurons (Perin, Berger, and Markram 2011), suggesting that trial-to-trial response variability may be in part the manifestation of the state of the surrounding neuronal population (Meshulam et al. 2017; Buonomano and Maass 2009). Pairwise interactions within a population can shape information representation (Moreno-Bote et al. 2014; Averbeck, Latham, and Pouget 2006; Ecker et al. 2011) and can be regulated by top-

down influences (Cohen and Maunsell 2009). Therefore, comprehensive descriptions of stimulus representations in primary sensory cortex require a network perspective. Here, we used two-photon imaging to record from large populations of L2/3 excitatory neurons in mouse V1 to study effects of local population activity on trial-to-trial variability.

Understanding the sources and consequences of response variability is necessary to extend theories of sensory computation from the average case to single trials. Perception and behavior take place in real time, after all, so variable responses must be taken into account to understand stimulus representations in cortex. Shared variability in neural responses is commonly quantified by the set of pairwise correlations between neurons, and the structure of these correlations can have constructive or destructive effects on stimulus encoding in populations of neurons (Hu, Zylberberg, and Shea-Brown 2014; Montijn et al. 2016), highlighting the importance of its characterization. Moreover, complex patterns of population activity in retina can be captured by considering only neuron firing rates and pairwise correlations (Schneidman et al. 2006). Covariability can also be shaped by cognitive properties such as attention in order to improve perceptual acuity (Cohen and Maunsell 2009). Whether trial-to-trial variability is harnessed to improve the fidelity of sensory representations, or accounted for when decoding from noisy signals, the properties of response variability have a large impact on neural function.

Research on the correlation structure of population activity is still incomplete, however, and can be meaningfully expanded by incorporating a more comprehensive sampling of the network (Olshausen and Field 2006). V1 populations consist of neurons whose activity is not modulated

by, or is untuned to, a given stimulus. It is still an open question how this subpopulation contributes to neuron correlations. Two-photon imaging results in a relatively unbiased sampling of spatially proximal neurons including both tuned and untuned subpopulations. Neurons unrelated to a behavioral task can help predict activity in neighboring neurons in hippocampal CA1 (Meshulam et al. 2017). In V1, it has been shown that untuned neurons can help to decode the orientation of drifting gratings (Zylberberg 2017). We investigate how co-fluctuations in the activity within tuned and untuned neurons interacts with responses to drifting grating stimuli.

We characterize population activity and correlations between tuned and untuned subpopulations in order to understand the relationship between single-cell response properties and recurrent network dynamics. Traditional noise correlation analyses study covariability independent from stimulus-driven activity (Cohen and Kohn 2011). However, untuned neurons have no stimulus modulation, so we use an analogous partial-correlation based method that additionally accounts for population-wide covariability. Fluctuations common across a local population are an important determinant of single-trial responses in mouse V1 (Okun et al. 2015; Lin et al. 2015). Capturing this additional variable allows us to study the correlations in the entire unbiased sampling of the population. Additionally, the partial correlation matrices are asymmetric and relatively sparse, and can thus be represented as a weighted, directed graph. Graph theory analysis is used to summarize structure in complex networks and can resolve emergent properties resulting from pairwise relationships. Connectivity patterns, or motifs, in graphs can be characterized within small groups of neurons (Chambers and MacLean 2016) or across an entire population (Nigam et al. 2016; Watts and Strogatz 1998). Motifs have proved to be impactful for

understanding complex biological processes including transcription networks (Alon 2007) and spike propagation (Chambers and MacLean 2016). More generally, motifs patterns impact information representation in complex systems (Timme et al. 2014; Yu et al. 2011) and have increasingly been a subject of interest among neuroscientific disciplines.

From the graph neighbors of a given neuron, we can accurately predict activity on single trials using a simple, linear model. The local population contains information sufficient to predict trial-to-trial variability and recapitulates average tuning properties. Furthermore, neurons that are well-modeled by the activity of their neighbors have specific signatures of functional connection motifs. Across the entire graph, the most predictive motifs are also the most prevalent, suggesting that this structure is responsible for the overall quality of reconstruction observed. The triplet motif that facilitates predictions of neural activity may have a broad impact on information representation in graphs. Notably, total variance explained in a field of view scales with the number of neurons imaged, suggesting larger sample sizes are required to fully capture local network interactions.

Results

Response properties of V1 populations

To study interactions and response variability in local, cortical populations (<800μm diameter imaging plane), we imaged L2/3 excitatory neurons (72–347 neurons; 25–33 Hz; n=8 animals; 23 distinct fields of view; Fig III-1A) in mouse V1 during presentation of drifting gratings (Fig III-1B, C). Square-wave gratings at 12 directions were presented in pseudo-random order for 5

seconds each, interleaved with 3 seconds of mean-luminance matched grey screen. We designed grating stimuli with slightly longer durations than many studies (Montijn et al. 2016; Durand et al. 2016; Reimer et al. 2014) for two reasons: first, to allow for the slow decay of the calcium indicator to fall to baseline in order to remove any confounds from the previous grating, and second, to study the sustained response in the population, rather than a transient response to stimulus onset (McLelland et al. 2010). Mice were awake and allowed to freely run on a linear treadmill. The majority of neurons showed significantly increased activity to one or more gratings over grey screen (3023/4535). Of the responsive subpopulation, most neurons were significantly tuned to orientation or direction (2073/3023; 540/3023 respectively). Neuron tuning was measured by fitting an asymmetric circular Gaussian tuning curve to the trial-averaged mean fluorescence in each grating direction (Fig III-1D). These numbers of tuned and untuned neurons are in line with other population studies of awake, mouse V1 (Montijn et al. 2016; Reimer et al. 2014). In subsequent analyses, we pooled all visually responsive neurons without significant direction or orientation tuning into a class of 'untuned' neurons, differentiating two distinct subpopulations in V1 by their responsivity to drifting gratings.

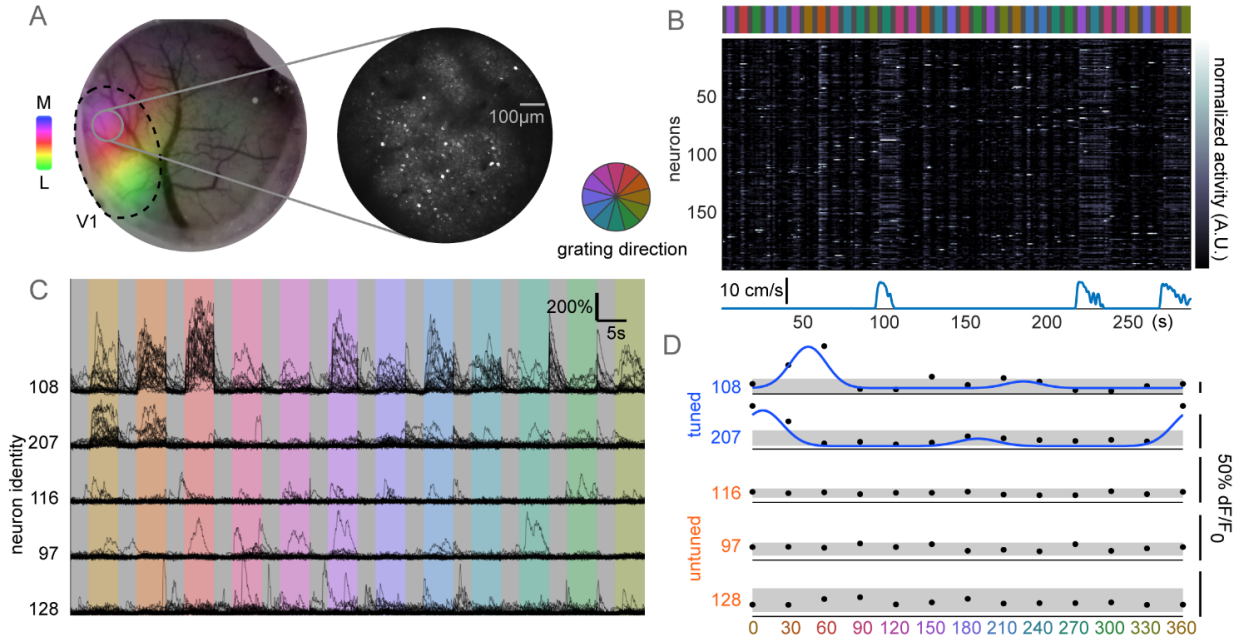

Figure III-1. Categorizing response properties in V1 populations
(A) Left, anecdotal cranial window, V1 retinotopy mapped by intrinsic signal overlaid (colormap shows visual field azimuth, lateral to medial). Right, two-photon field of view L2/3 population from grey region on left. (B) One block of stimulus presentation (top; grey interleaved gratings) with simultaneous imaging data showing neuron normalized fluorescence over time (middle). Bouts of running (bottom) and gratings show coordinated increased activity across neurons. (C) Five example neurons' activity, black traces represent one trial, organized by grating direction. Grey period preceding gratings are kept together, with discontinuities at end of gratings. Top two neurons show increased activity to specific directions, others are untuned with substantial trial-to-trial variability. (D) Trial-averaged activity reveals classic tuning properties in a subpopulation of neurons. Tuned neurons fit with asymmetric circular Gaussian tuning curve (blue), untuned neurons show no increased activity relative to grey periods (grey bar mean+/-std average response).

Responses are highly variable across trials

Single trial responses to gratings showed a high degree of variability, even in strongly tuned neurons, manifesting as occasional strong responses to null-directions and weak or absent responses to preferred directions (Fig III-1C, III-2A). Tuning curves described average response strength of tuned neurons well (0.70+/-0.20 R^2). Responses across single trials in tuned neurons,

however, were not well-described by their tuning curves. The mean fluorescence for each direction only explained a small fraction of the total trial-to-trial variance which we calculated by subtracting mean fluorescence in each direction from fluorescence in each trial (Fig III-2B). This finding matches earlier results in awake, mouse V1 (Montijn et al. 2016). We computed the distribution of response strength (mean fluorescence across a grey or grating presentation) within single trials, z-scoring to account for neurons with different activity levels. Single trial response distributions were skewed, with most responses weaker than the mean (Fig III-2C). Tuned neurons showed slightly stronger responses during gratings, as compared to untuned neurons which had nearly identical response distributions in grey and grating trials. Tuned response distributions were strongly overlapping, however, consistent with the hypothesis that individual neuron activity is not solely driven by tuning properties.

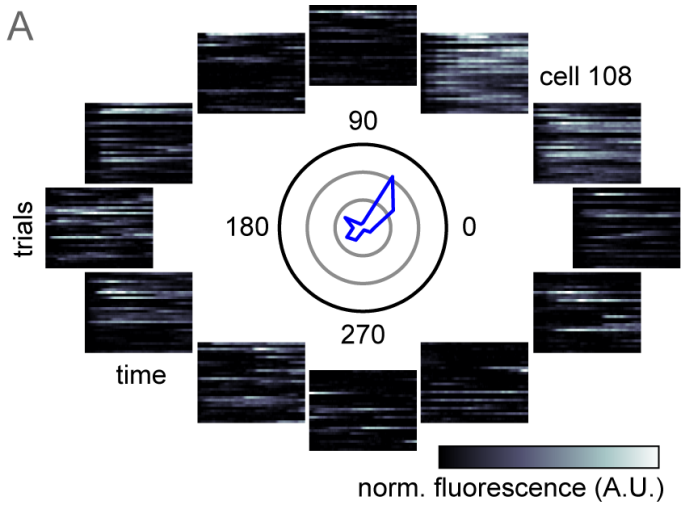

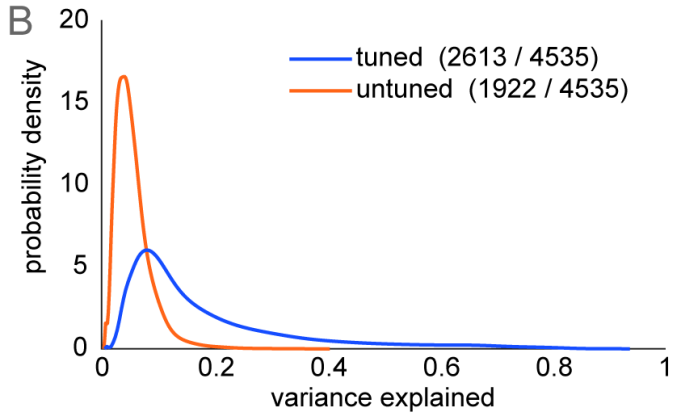

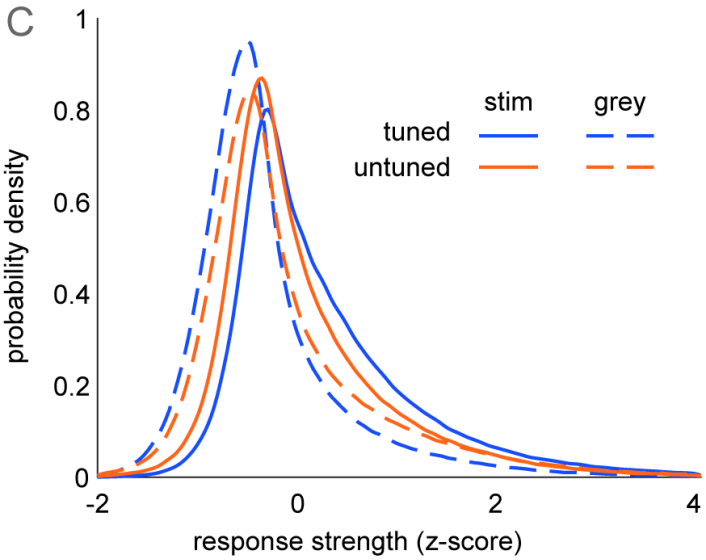

Figure III-2. Single trial response variability

(A) Single trial responses of representative tuned neuron. Mean tuning curve in polar coordinates shown in center (rings=100% dF/F_0). Each heatmap positioned by grating direction shows normalized time-varying activity for each trial. (B) Trial-averaged tuning responses do not explain trial-to-trial variability. Tuning average (blue curve in A) is subtracted from mean fluorescence in each trial to compute variance explained across trials. (C) Averaging across time, neurons show highly variable responses during grating and grey trials with a right-skewed distribution due to strong activity in a few trials. Tuned neurons show larger increase in response during gratings than untuned neurons.

To further describe population activity, we computed the time-varying activity during the presentation of a grating and its preceding and following grey presentations. For each trial, we removed neurons that were silent (defined as no fluorescence change 2*S.D. above baseline) and computed the time-varying z-scored fluorescence across neurons (Fig III-3A). Despite the long duration of stimulus presentations, adaptation effects were minimal in these L2/3 neurons, as tuned neurons showed sustained activity throughout the 5 second stimulus presentation. Untuned neurons have weak modulation to the onset and offset of the stimulus and are equally active during the grey period. However, these effects are small in comparison to variability across trials, as indicated by strong overlap between the activity of the two subpopulations. In the awake animal, running speed is known to strongly influence spike rates (Niell and Stryker 2010), and we similarly observe that periods of high population activity are very likely to occur during periods of running (Fig III-3B, anecdotally in III-1B). However, mice did not preferentially run during grating or grey presentations (probability of running 9.7+/-7.7% during gratings; 10.0+/-7.4% during greys; p=0.278 paired t-test). While we used changes in fluorescence for all other analysis, we expanded our comparison of both subpopulations by estimating spike rates using a spike-inference from calcium fluorescence algorithm (Friedrich, Zhou, and Paninski 2017). We found that untuned neurons exhibited an identical firing rate distribution to tuned neurons (Fig III-3C). Therefore, differences between subpopulation dynamics cannot be explained by differences in firing-rates.

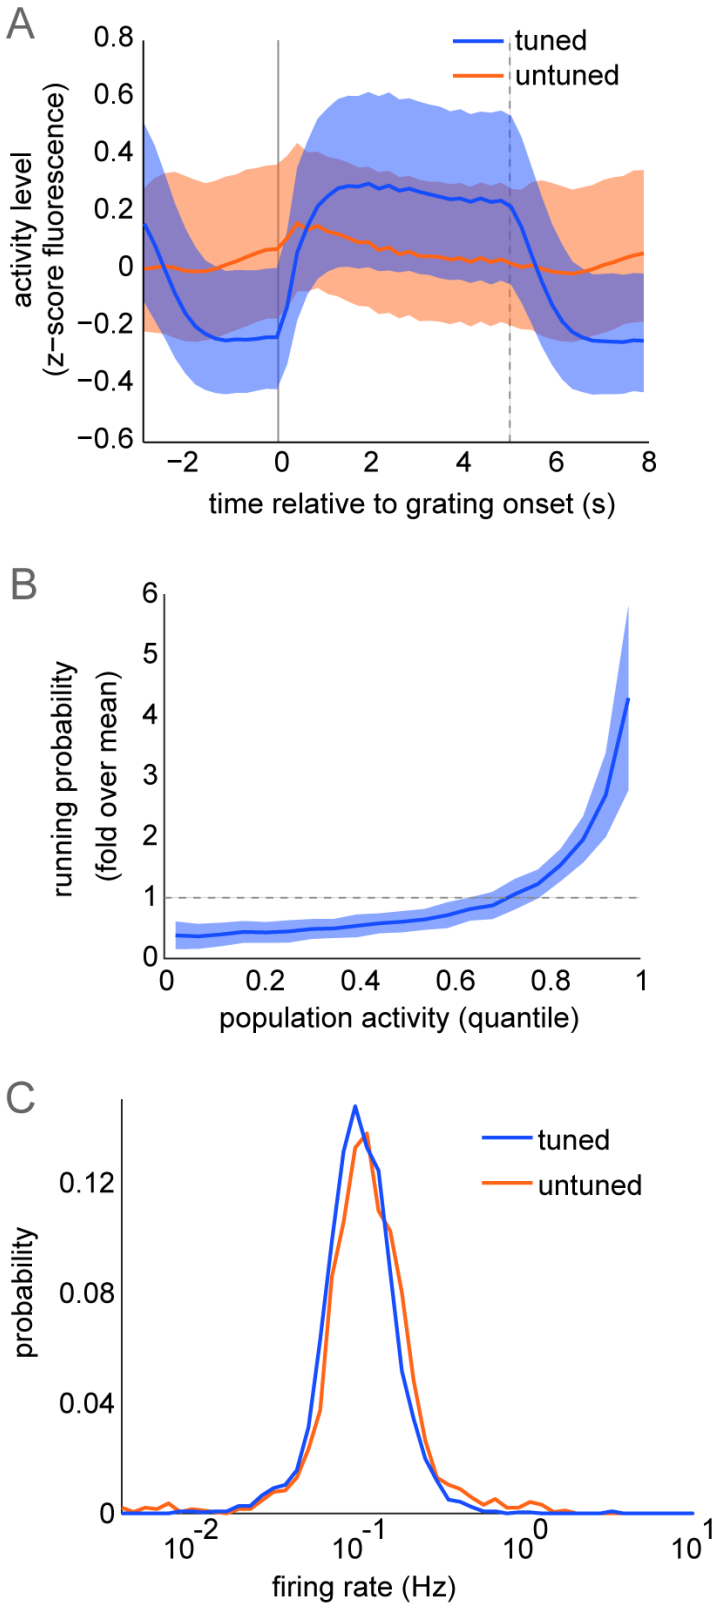

Figure III-3. Population response properties
(A) Mean time-varying fluorescence of active neurons, pooled across datasets, during grating stimulus and preceding, following grey periods in tuned, untuned subpopulations (median+/- quartiles across pooled neurons). The 3-second grey period is shown before and after the grating response. Grating onset, offset shown in solid, dotted lines respectively. (B) Total population fluorescence split into equal quantiles to compute probability of running (speed>0.1cm/s). Plot shows mean+/- std across datasets (n=23). (C) Inferred firing rate distributions computed across all imaging frames from spike trains inferred from fluorescence.

Correlation structure during drifting grating and grey screen trials

Untuned neurons are a large proportion of total neurons, exhibit similar spike rates to tuned neurons, and are likely to contribute to correlation structure in the population. To begin describing how dynamics are affected by stimuli in populations containing tuned and untuned neurons, we first analyzed pairwise correlations during grating and grey presentations. For each pair of neurons, we computed the correlation coefficient between the mean fluorescence (averaged over time) in either grating or grey trials. We did not remove signal-dependent responses, nor did we shuffle responses to eliminate simultaneous cofluctuations, therefore, these correlations are a combination of signal and noise correlations. Overall, within-subpopulation correlations are weak (0.014+/-0.027 tuned; 0.033+/-0.047 untuned), and beween-subpopulation activity is slightly anti-correlated (-0.019+/-0.017). Comparing mean pairwise correlations across all pairs according to their subpopulation, only within-subpopulation correlations are affected by grating stimuli, while between-subpopulation correlations are unchanged between stimulus and grey conditions (Fig III-4A). Tuned neurons show a strong decrease in mean correlations during gratings, as seen in macaques (Smith and Kohn 2008). Conversely, untuned neurons are more strongly correlated during gratings, and yet correlations between tuned and untuned neurons do not change in magnitude between stimulus and grey conditions. Untuned neuron activity is not directly modulated by the stimulus, so changes within this subpopulation most likely reflect changes in activity from the tuned subpopulation propagating through local synaptic connectivity. However, this occurs without a change in mean correlation between tuned and untuned neurons, bringing to question the mechanism involved. The pairwise correlations in tuned neurons during the stimulus are a function of their preferred grating directions, as expected

for signal correlations (Fig III-4B). Similarly tuned neurons show strong correlations, while orthogonally tuned neurons show negative correlations. This structure is not present during activity in grey periods, however. This is surprising, because if local connectivity underlies correlations in the grey condition, one should expect the structure seen during gratings to remain in part because similarly tuned neurons are more likely to be connected (Ko et al. 2011).

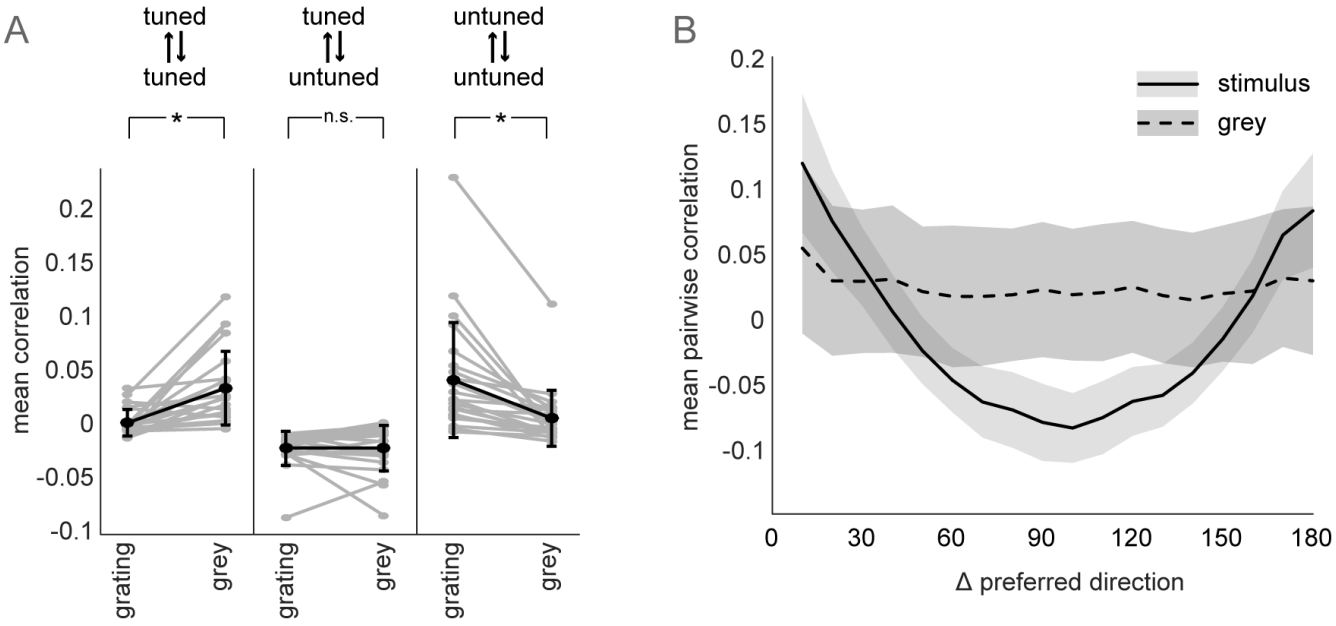

Figure III-4. Correlations during grating and grey stimuli
(A) Mean activity correlation between stimulus conditions, within and between subpopulations (mean+/-std in black; individual datasets in grey). Significance determined with Tukey-Kramer corrected one-way ANOVA at alpha=0.01 (tuned-tuned p=0.004; tuned-untuned p=0.999; untuned-untuned p=0.001) (B) Correlations within grating and grey trials for tuned neurons (mean+/-variance pooling neurons with given difference in tuning over datasets).

Measuring trial-to-trial covariability

Overall correlations in populations including untuned neurons begins to reveal properties of local population activity, but in order to study the sources and structure of trial-to-trial shared variability, researchers attempt to remove the stimulus-dependent portion of responses leaving

only variability, or 'noise' (Cohen and Kohn 2011). Correlated fluctuations between the remaining responses are therefore often called 'noise correlations.' However traditional noise correlation analysis was not appropriate in this case for the following three reasons: 1) the untuned neuronal subpopulation has no stimulus-dependent response; 2) in tuned neurons stimulus-driven response explains only a small portion of overall variability in tuned neurons (Fig III-2B), making traditional noise-correlation analysis unsatisfactory for these neurons as well; 3) finally, V1 populations in mice, cat, and macaque are characterized by global cofluctuations common to every neuron (Lin et al. 2015; Ecker et al. 2014; Schölvinck et al. 2015) such as covariance driven by running (Fig III-1B, III-3B). Therefore, in order to study pairwise noise correlations within and between subpopulations we used partial-correlation analysis that allowed us to account for stimulus-driven responses and population-wide co-fluctuations in both tuned and untuned neurons.

Visual stimuli were presented in 5-minute blocks. Each block of visual stimuli contained three repetitions of each direction in pseudo-random order and corresponding luminance matched grey periods. While the order of grating stimuli were pseudo-random in each block, the order was maintained between blocks. For each pair of neurons, the average activity across all remaining blocks, other than the block considered at that time, represented the two stimulus-dependent responses capturing tuning properties, when present. Additionally, we accounted for the mean within-block population-wide activity of all remaining neurons. Consequently, we were able to compute a partial correlation coefficient in each block between the activity of every pair of neurons, controlling for stimulus responses and population co-activity (Fig III-5A). The mean

partial-correlation across blocks is taken as the final correlation strength and entered as an edge weight into the functional connectivity matrix. We also added directionality to the partial-correlation by examining the mean cross-correlogram across blocks for the neuron pair. If the peak value occurs at lag 0 (i.e. within the same imaging frame), the edge was bidirectional, otherwise the edge was in the direction of positive lag. Lags greater than 500ms were thrown out and correlations set to zero. This resulted in a functional network described by a directed weight matrix. Though we interpret these partial-correlations as equivalent to noise correlations, the correlation matrices are different from traditional noise correlations in two important ways. First, many pairs of neuron correlations are exactly zero (51.7+/-7.2%), and second, non-zero correlations are asymmetric (Fig III-5B). This allows us to analyze these matrices from a graph-theoretic perspective representing the functional partial correlations as a weighted, directed graph. Graph representations of pairwise edges allow us to analyze population-wide statistical features of the correlation structure. Overall, partial-correlation strengths, synonymous with edge weights, were long-tailed, centered slightly above zero (Fig III-5C), similar to noise-correlations observed elsewhere (Montijn, Vinck, and Pennartz 2014). As expected, tuned edge weights were similar to signal correlations (Montijn, Vinck, and Pennartz 2014; Cohen and Newsome 2008), with similarly tuned neurons having larger edge weights on average (Fig III-5D).

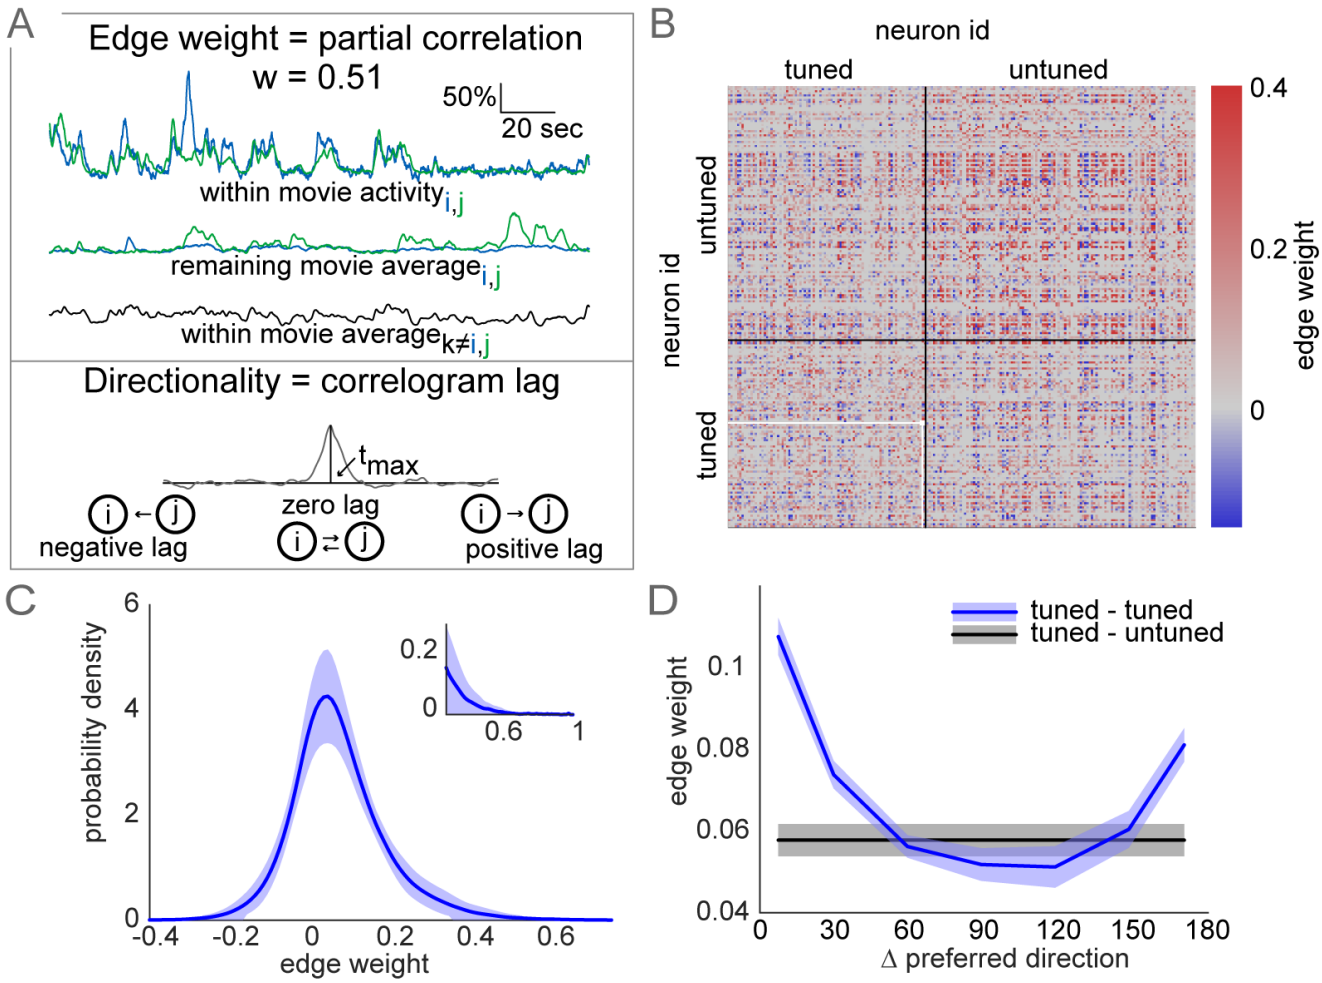

Figure III-5. Partial correlation matrices representing trial-by-trial covariability
(A) Edge weight estimated by mean partial correlation. Fluorescence traces of a neuron pair (i, j) within one movie, and three variables accounted for in partial correlation: remaining-movie average activity of i, j and within-movie average activity of remaining population (top). Edge direction set by offset of peak (t_{max}) in neuron i, j correlogram (mean across movies). (B) Partial-correlation matrix sorted by subpopulation showing dense within- and between-population correlations. Neuron pair entry from (A) in white. (C) Kernel-density estimate distribution of all edge weights (mean+/-std across datasets). Inset shows right-tail of the strongest edges. (D) Tuned edge weights match overall stimulus correlations, despite accounting for stimulus response, and slightly weaker between-subpopulation edge weights (mean+/-sem over datasets).

Graph structure of partial-correlations

We next analyzed the partial-correlations within and between tuned and untuned subpopulations.

The graphs exhibit dense correlations with varying strengths among subpopulations (Fig III-7A).

To analyze biases in edge strengths, we thresholded the matrices at increasing values, setting all edges below each threshold to zero. Among all edges, within-tuned connections are more likely, while within-untuned and between-subpopulations are slightly less common (Fig III-6, Fig III-7B; within-tuned 54.2+/-8.6%, within-untuned 44.1+/-6.0% between 43.9+/-6.2%). Between-subpopulation connections remain the least likely at higher thresholds, but among the strongest edges, within-untuned connections are the most likely. We then recomputed partial-correlation matrices, exclusively using frames during the grey condition or during grating condition to see how correlation strengths were affected, despite controlling for mean stimulus-dependent activity. The two resulting matrices did not have significantly different edge strength (stimulus 0.12+/-.02; grey 0.12+/-.03; p=0.75 one-way ANOVA). However, probability of connection was higher within grey frames (grating 45.3+/-6.5%; grey 58.0+/-8.2%; p=5.8*10^-7 one-way ANOVA), indicating a higher degree of interconnectivity in the population in the absence of stimuli. When analyzing magnitude of edge strength differences between grey and grating matrices, we found that all connection types changed similarly with a mean near zero (Fig III-7C; within-tuned -0.011+/-0.097; within-untuned -0.005+/-0.103; between -0.008+/-0.093).

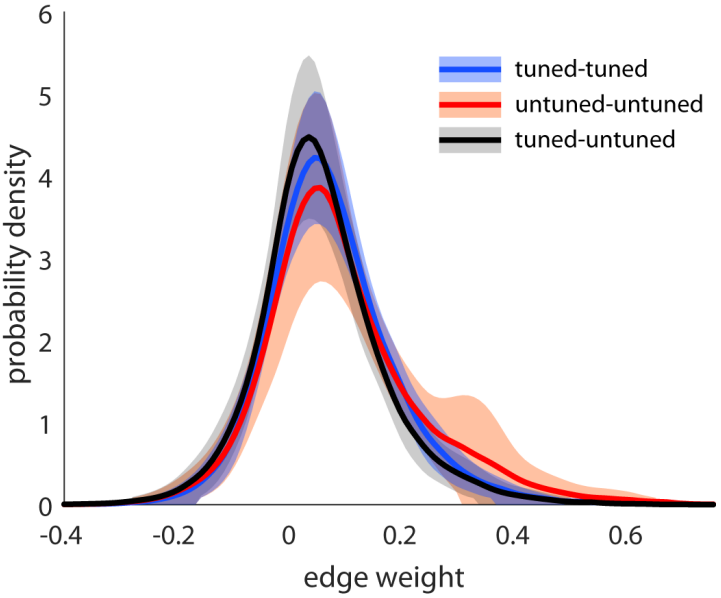

Figure III-6. Partial correlation between tuned and untuned subpopulations
Probability density kernel estimates of edge weights for three categories of connections. Mean and standard deviation shown across all datasets (n=23).

Since population dynamics are partly constrained by synaptic connectivity (Luczak and MacLean 2012), we evaluated whether there was a spatial component to the weight matrices. Edge probability fell monotonically with distance between neurons, similar to traditional noise correlations (Montijn, Vinck, and Pennartz 2014) and synaptic connections (Perin, Berger, and Markram 2011; Oswald et al. 2009). Notably, the decay in connection probability is slower within tuned neurons compared to within untuned neurons (Fig III-7D). Between-population decay lies in the middle. If the spatial structure of untuned correlations is exclusively driven by local connectivity, then bottom-up sensory drive is the most likely source for the longer range of functional correlations among tuned neurons. Furthermore, the mean lag (delay of the cross-correlogram peak used to determine edge directionality) is greater over longer distances and accumulated evenly across subpopulation connection type (Fig III-7E). Assuming a linear change in lag over space, these data suggest the speed of functional correlations in this

preparation is roughly 25 mm/s. Despite allowing for directional edges, nearly half of all edges were bidirectional (42.5+/-16.8%). As a function of edge strength, bidirectional edges are more prominent within-subpopulations, and are strongly biased toward the strongest weights (Fig III-7F). Thresholding at increasing edge strengths sparsifies the matrices, so we normalized the bidirectional edge counts by the probability of bidirectional edges assuming connections are placed randomly. Among all edges, bidirectional edges occur less often than random. The strongest edges, however, are roughly 5 times more likely to be bidirectional than random. Overall, zero-lag connections are less frequent between tuned and untuned neurons, suggesting a transmission or propagation of information, rather than simultaneous representation of the same information between subpopulations.

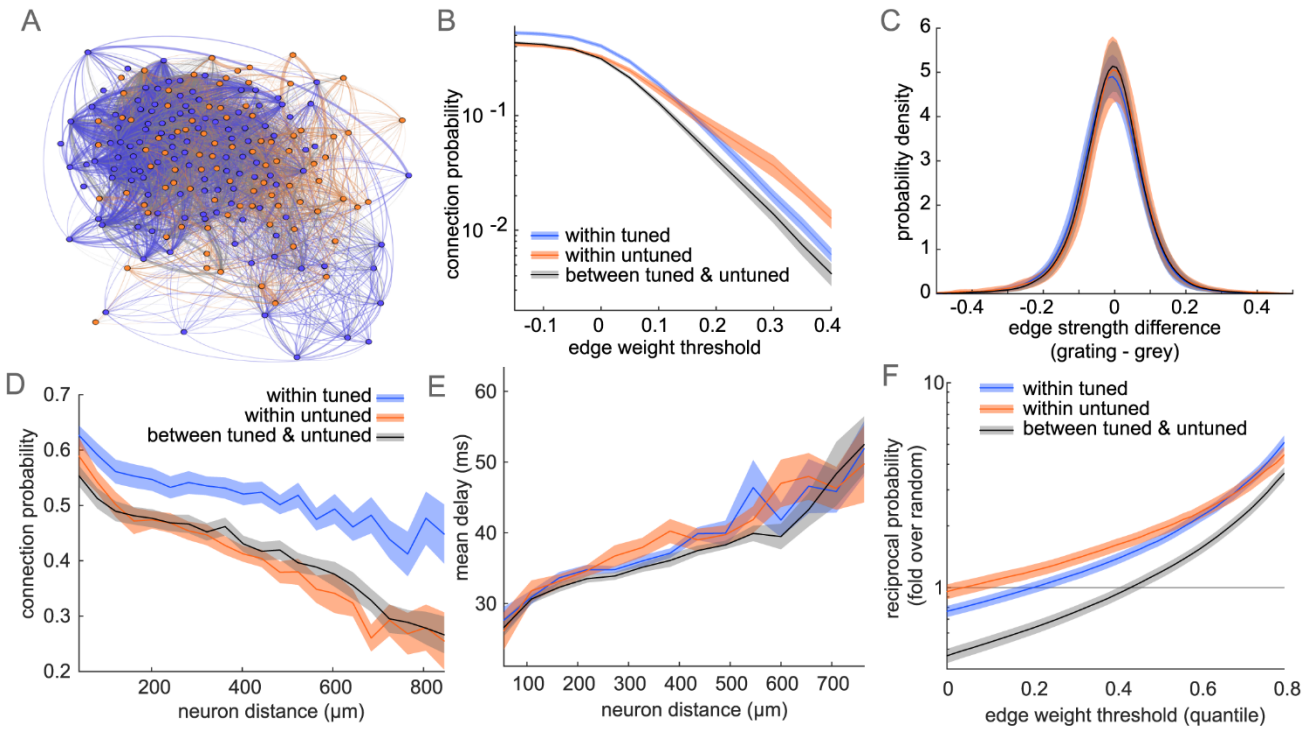

Figure III-7. Directed graph structure of partial-correlation matrices
(A) Example graph representation of matrix in (5B), nodes colored by subpopulation, edges colored by connection type as in (B-F), edge thickness represents absolute-value correlation strength. (B) Probability of connection by connection type for matrices with edges weaker than an increasing threshold set to zero. (C) Difference in edge strength between matrices using exclusively grating or grey frames. (D) Connection probability as a function of distance between neurons. (E) Mean lag of edges over neuron distance (F) Probability of a bidirectional edge at increasing thresholds, normalized by bidirectional probability assuming independence of edges. Plots in (B-F) show mean+/-sem across datasets.

Modeling trial-to-trial variability from local population activity

Because trial-averaged tuning poorly captures trial-to-trial responses we asked whether information in the local population could better explain V1 trial-to-trial response variability. Pairwise correlations have been shown to capture a significant portion of the complexity in population activity (Schneidman et al. 2006). We tested whether the activity of a neuron could be modeled from the activity of its neighbors that had a non-zero correlation. Since correlation coefficients capture the linear relationship between neuron coactivity, we used a simple linear

combination of the input neuron activity with partial-correlation coefficients (edge strengths) as weights. This model gave a time-varying prediction of the fluorescence of a given neuron, which was then rescaled by an offset and a gain to account for different numbers of input neurons (Fig III-8A). In many cases, this model resulted in highly accurate predictions of activity. Mean squared error of the reconstruction was small, and often near optimal compared to weights estimated by regression (Fig III-8B). Tuned neurons were slightly better-modeled on average than untuned neurons (tuned MSE 0.014 median, 0.037 inter-quartile range; untuned MSE 0.017 median, 0.046 inter-quartile range), possibly because of the additional stimulus-dependent information captured by other tuned inputs. The ratio of MSE to mean-squared fluorescence from the true trace, subtracted from one, gives the percent variance explained of the model for each neuron. The optimal reconstructions from regression explained 77.2+/-15.2% of the variance in activity across all neurons (n=4531). Using partial-correlation coefficients as weights performed near this upper-bound, at 65.8+/-17.1%. We tested the robustness of the partial-correlation based predictions modeling 5 minutes of fluorescence change, corresponding to one block, using a functional network built using 45 min, corresponding to the nine other blocks, of non-overlapping recordings to recomputing the weights. We note that one block, i.e. five minutes of stimulus epochs, corresponded to three stimuli at each of the 12 directions and interleaved grey periods. We then tested our predictions on the left-out 5 minute dataset (repeated to leave out each dataset once). On this cross-validation procedure, the average performance on the left-out epoch reached 55.4+/-18.9% variance explained.

We selectively removed either tuned or untuned input neurons to evaluate the relative contributions to model predictions. The accuracy of prediction decreased more when removing within-subpopulation inputs as compared to removing between-subpopulation inputs, but the decrease is small and distributions across neurons strongly overlap (Fig III-8C). Within-subpopulation inputs are more frequent, however, so removing a larger fraction of inputs should be expected to have a larger impact on our ability to model neuronal fluorescence. On a single neuron basis, correlations between tuned and untuned neurons both contribute to predicting time-varying activity. We next asked whether our model based on local population activity could also predict trial-averaged tuning properties. For tuned neurons, we used the modeled fluorescence traces to recompute the mean fluorescence in each grating direction. The average responses in direction space were added together to obtain a mean tuning vector. The modeled fluorescence had very similar tuning properties to the data, as measured by the cosine similarity between model and data tuning vectors (Fig III-8D). The modeled activity was constructed from a neuron's input edges; we also asked if a neuron's outgoing edges were related to tuning properties, which may indicate neurons that are good at decoding the stimulus (strong tuning) were also good at decoding activity in its neighboring neurons (strong outgoing edges). However, there was nearly zero correlation between the tuning strength of tuned neurons and its mean outgoing connection strength ($r=0.04$, $p=0.02$). Local population activity therefore contains information to capture trial-to-trial variability as well as trial-average stimulus response properties, but a neuron's dependence on the stimulus does not affect its local population correlations.

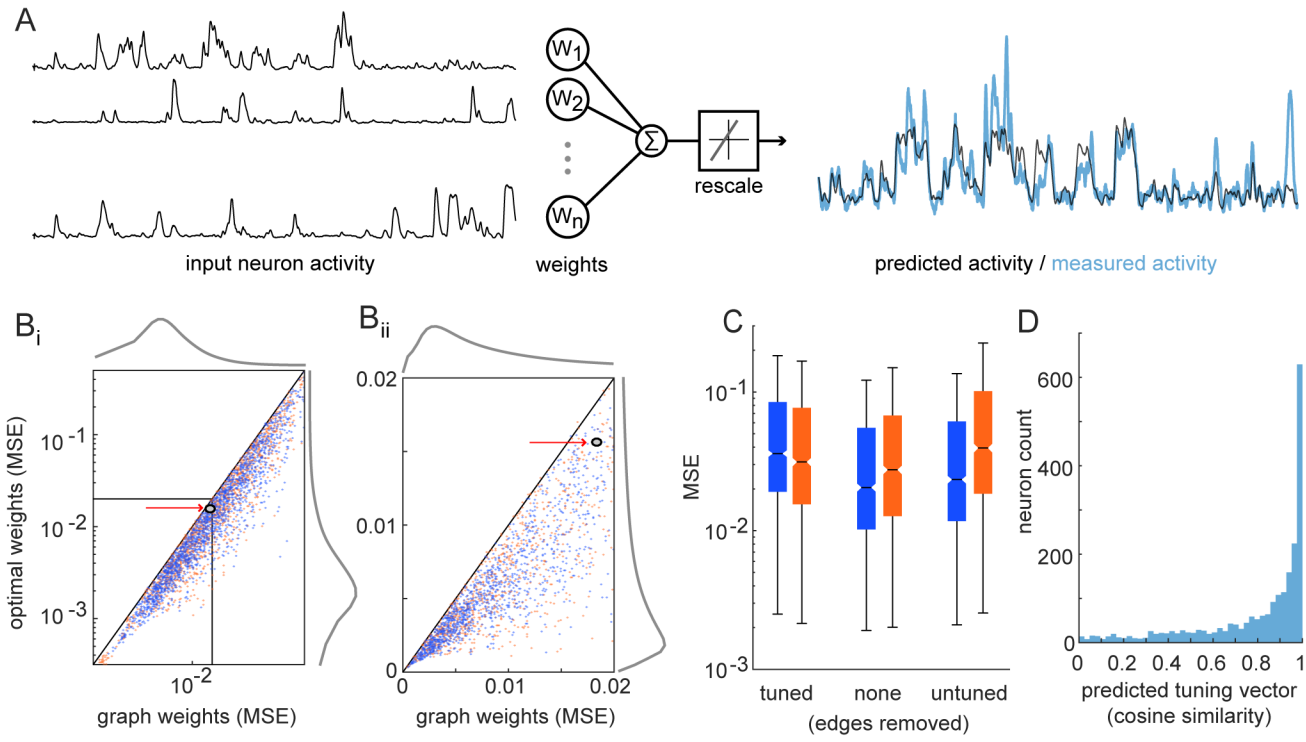

Figure III-8. Modeling time-varying fluorescence from partial-correlation graphs
(A) Illustration showing linear combination of incoming-edge neighbors fluorescence traces and rescaling to predict a given neuron's activity. (Bi) Reconstruction quality across tuned and untuned neurons, quantified by mean-squared error (MSE) of reconstruction, compared to MSE using optimal weights (marginals in grey; arrows show neuron from (A) (Bii) zoom in on neurons with best reconstruction. (C) Increase in MSE when selectively removing within- or between-subpopulation connections (boxplots show median, quartiles, with whiskers extending to most extreme value within 1.5*inter-quartile range outside quartiles). (D) Similarity of mean tuning vector between tuned neuron and reconstructed average activity.

Pairwise and motif contributions to model performance

To investigate how individual neurons' inputs contribute to reconstruction of activity, we selectively removed input neurons based on their edge strength and measured the increase in reconstruction error (MSE), normalized by total mean squared fluorescence in the neuron. As expected, the strongest edges contribute the most to activity reconstruction with the strongest 25% of weights containing over half of the reconstruction capability, whereas removing half of

the weakest edges had no discernible effect (Fig III-9A). Interestingly, randomly removing edges does not linearly reduce prediction performance. This suggests a level of redundancy in the predictive information of input neuron activity. Accounting for the cumulative weight removed, we still found worse performance when removing the strongest edges first, and removing half of the total weight using only the weakest edges has minimal effect on reconstruction error (Fig III-9B). Normalizing for the total weight removed reveals a nearly-linear increase in reconstruction error when removing the strongest weights, suggesting that these neurons may hold independent information from the remaining input pool.

To address the possibility of synergistic or redundant information between input neurons, we analyzed connections between triplets of neurons termed 'motifs' in graph theory literature. Triplet connection motifs are built up from pairwise connectivity and collectively represent higher-order connectivity patterns that cannot be captured by individual edges, and can have strong implications for computation and information propagation within graphs (Timme et al. 2014). We looked at the clustering of triplet motifs for each type of triangle that can be formed with directed edges (Fig III-9C) (Fagiolo 2007). Clustering is a measure of how many motifs are present among a neuron's neighbors given its input and output connections. In comparison to Erdos-Reyni graphs, which have uniform, low levels of clustering across motifs, cycles of edges in the data are less clustered on average, with all other motifs showing elevated clustering. The middle-man motif shows the strongest clustering. Similar results have been found in activity generated by simulated and ex vivo neural networks, although fan-in clustering was higher than middle-man (Chambers and MacLean 2016). To address the possibility that triplet motifs are

responsible for explaining more of the neuronal response than pairwise edges alone, we analyzed the relationship between motif clustering in single neurons and the variance explained by the predicted activity. Because the magnitude of clustering is different across motifs, we first z-scored clustering coefficients, then computed the mean across neurons with different levels of variance explained. Neurons with the best reconstruction showed higher clustering of middle-man motifs and lower cycle clustering, relative to other neurons in the population (Fig III-9D). Total clustering, as well as fan-in and fan-out (not shown for clarity), had weak, positive correlation with variance explained. Together, these relationships map directly onto the overall prevalence of the graph motifs, suggesting that the graph structure has an important function in representing population information.

Because the motif with the highest mean clustering was also most indicative of model performance, we hypothesized that the partial-correlation structure might underlie our ability to predict neuron activity from its local population. Interestingly, across fields of view, the total variance explained in the population increased with number of neurons imaged (Fig III-9E; r=0.58). To compute population variance explained, we took the sum of all squared prediction errors (residuals) across neurons in a field of view, divided by the sum of the squared population fluorescence activity, and subtracted from one. This improvement of variance explained across the population with more recorded neurons suggests that, in addition to motifs, total neurons sampled in the population determines our ability to measure a neurons' single trial dependence on its local population. Moreover, the linear trend had no discernible plateau, so representing the local population may require recording from more than 300 neurons simultaneously.

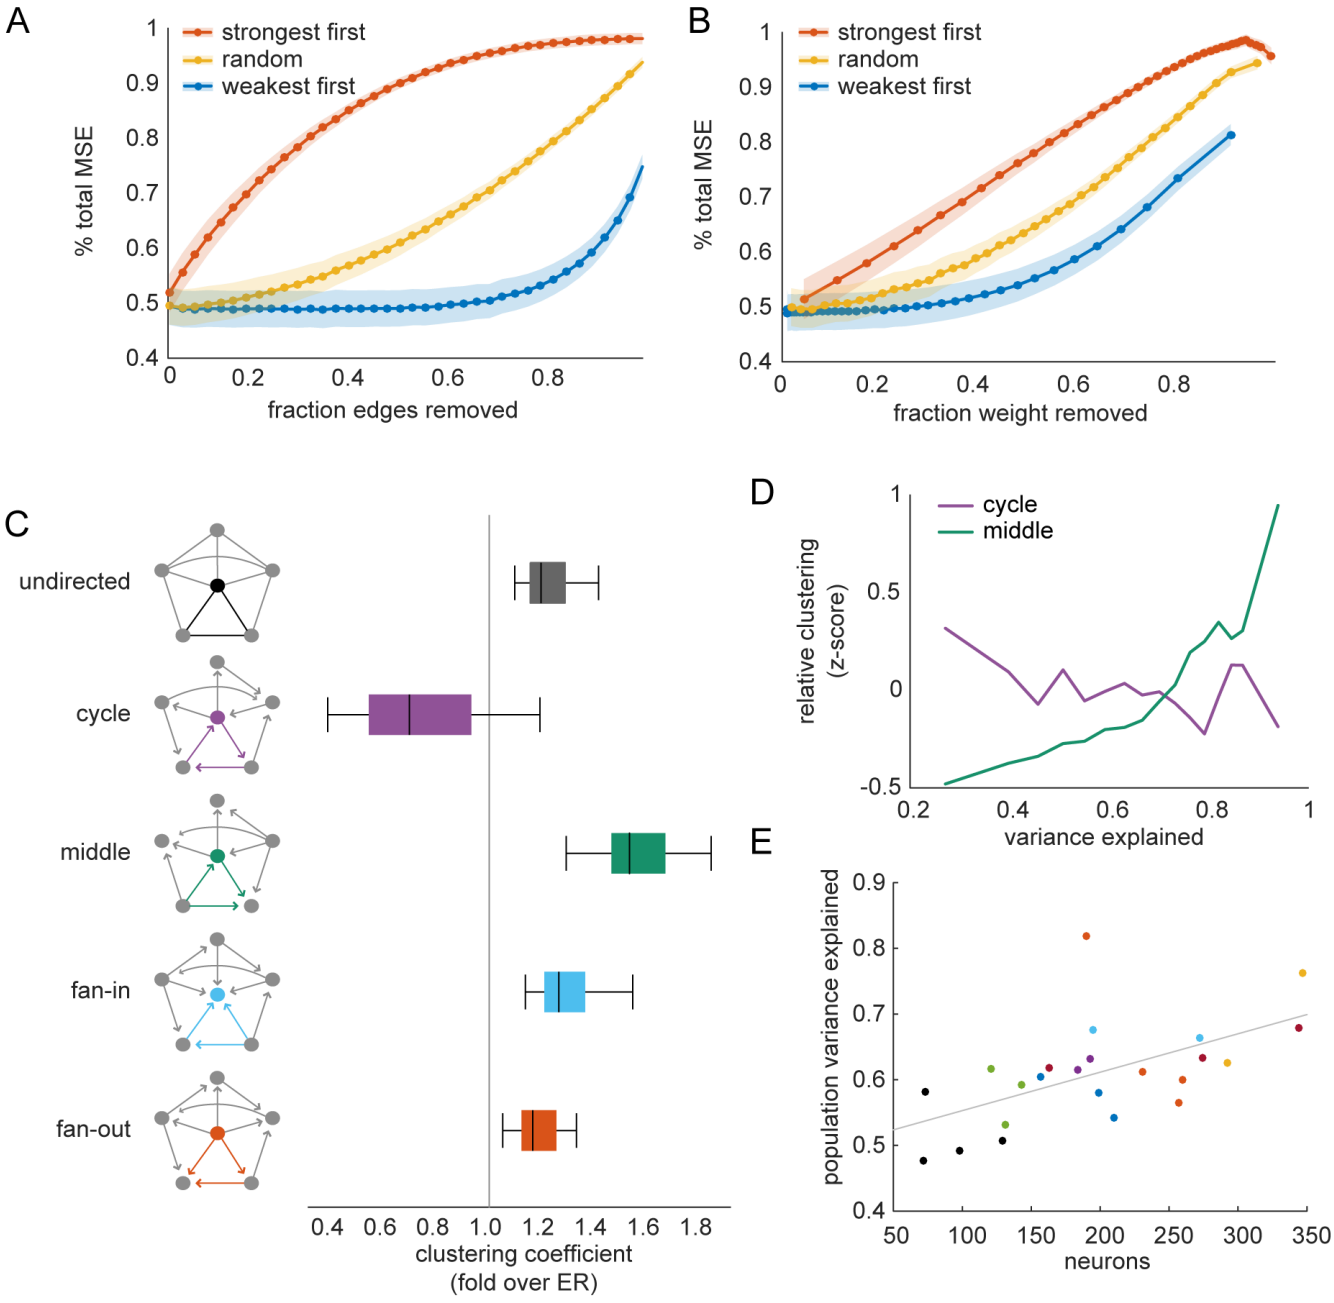

Figure III-9. Pairwise and triplet motif contribution to model performance
(A) Increase in mean-squared error (MSE) by number of edges removed, normalized by neuron total MSE (mean+/-variance across all neurons). (B) Increase in MSE by cumulative weight removed (mean+/-variance). (C) Clustering of triplet motifs in graphs (illustrated in left column) for each type of directed triplet motif. Box-plots as in Figure III-8. Mean clustering was normalized by sparsity-matched Erdos-Reyni graphs. (D) Mean clustering coefficient, z-scored by mean/std for each motif type in (C), at different levels of model prediction showing neurons with best reconstruction had highest middle-man clustering. (E) Total population variance

explained for each dataset (colored by mouse identity) shows better performance with more neurons imaged (linear trend in grey).

Discussion

Summary and correspondence with previous findings

We sought to describe the interrelationships within local populations of V1 neurons, including tuned and untuned neurons, as they relate to single-trial responses to grating stimuli. We used two-photon imaging to record from L2/3 excitatory populations constitutively expressing calcium indicator GCaMP6s. Neurons with similar response properties showed stronger co-variability on average, but across the entire population there was a broad distribution of correlations driven, presumably, by a confluence of sensory drive and activity in the local population. The functional correlations in the recorded populations were sufficient to predict activity in individual neurons, far surpassing predictions from tuning characteristics alone. The dependence on local population activity reinforces theories of layer 2/3 acting under strong modulation with sparse activity and weaker dependence on sensory drive than layer 4 (Petersen and Crochet 2013). We summarized the structure of correlations as directed, sparse matrices in order to analyze population dynamics from a graph theoretic perspective. We demonstrate that a simple population model capable of predicting single-trial neural responses is also able to accurately predict trial-averaged tuning responses, a key feature of V1 function. We found that the prevalence a specific triplet connectivity motif, built up from pairwise correlations, corresponded with our ability to predict activity on single-trials. This result could not have been

observed from only studying pairwise correlations and motivates the continued use of graph theory to study neural population dynamics.

The single-neuron response properties in our data replicate imaging and electrophysiological results in awake recordings of V1 including proportion of tuned neurons (Durand et al. 2016; Reimer et al. 2014; Montijn, Vinck, and Pennartz 2014), and variance explained by the mean tuning curve (Reimer et al. 2014). Both tuned and untuned neurons have low firing rates on average, though estimation of firing rates from calcium imaging is not always straightforward. We note that single neuron properties, including firing rate and trial-to-trial variance, can change substantially between animal models and in different states of anesthesia (Ecker et al. 2014; Schölvinck et al. 2015).

Population structure of pairwise correlations

We found that the magnitude of signal correlations between tuned and untuned subpopulations do not change between grey and grating stimulus conditions, while within subpopulation correlations do change. Perhaps ubiquitous changes would be more expected, or exclusive changes among neurons modulated by the stimulus. Within subpopulation correlations change in opposite directions, however, and could serve as a mechanism to balance changes between the subpopulations.

We found that the spatial organization of the network is a strong determinant of correlation structure with correlations decaying over distance, consistent with paired patch clamp recordings

(Perin, Berger, and Markram 2011) and the correlation structure of activity in isolated preparations (Sadovsky and MacLean 2013). Correlation matrices were computed as an asymmetric partial correlation coefficient, accounting for stimulus and population effects, to allow the incorporation of untuned neurons into traditional noise correlation analyses. While this approach differs from standard noise correlation estimates, the magnitude of correlations and dependence on tuning similarity replicate previous results measuring noise correlations (Montijn, Vinck, and Pennartz 2014). For this reason, we interpret the partial correlations as measuring trial-to-trial covariability, and as being broadly equivalent to noise correlations. Tuned neurons, which combine feedforward sensory inputs with recurrent inputs, show a slower spatial decay of correlational values than untuned neurons, the latter of which presumably are driven more exclusively by recurrent, local inputs. These data show than subdividing the population by their response to grating stimuli can differentiate rates of spatial decay within the network.

Noise correlations are hypothesized to be driven in part by shared synaptic input (Doiron et al. 2016). Spatial dependencies of feedforward and recurrent connectivity can qualitatively change the spiking activity in network models (Rosenbaum et al. 2017). The relatively small diameter of the fields of view imaged here (<1mm vs 10mm) does not allow us to differentiate between the two modes of activity predicted by these models. These data suggest, however, that functional recurrent correlations have less spatial extent than feed-forward functional correlations. To our knowledge, analogous synaptic connectivity estimates do not exist for mouse V1.

In addition to spatial structure of correlations, we found an increase in mean temporal delay of correlations over distance. This delay spans roughly 20-50 ms in our field of view and is therefore likely to reflect timescales of functional correlations rather than monosynaptic transmission delays. We have previously found that functional correlations are indicative of synaptic connections, in some cases, when considering synaptic integration rather than synaptic delays (Chambers and MacLean 2015). The implied speed of this delay accumulation is much faster than propagating LFP waves observed in macaque M1 (Rubino, Robbins, and Hatsopoulos 2006) after normalizing for total cortical surface area (Rakic 2009). The propagation of beta-oscillations and temporal delays of functional correlations likely have different underlying mechanisms, which may explain the different speeds.

Prediction of single-trial responses from pairwise correlations

We were able to extract a substantial amount of predictive information from the local population using a simple, linear model. Other approaches have successfully predicted single-trial responses from ongoing population activity at a larger spatial scale, averaging over many neurons in a population (Arieli et al. 1996). In contrast, we use a large, unbiased sampling of the local population with single cell resolution to predict variability on single trials. Alternative models incorporating known neuronal nonlinearities (Pillow et al. 2008), or more sophisticated or biologically-relevant predictive models (Truccolo, Hochberg, and Donoghue 2010), may have improved performance. Consequently, the increase in total variance explained with population size may show different trends with alternative models. Nevertheless, we chose this modeling approach to maintain ease of interpretation and utilize the linear correlation coefficients

estimated from the data in a straightforward manner. The linear models had good performance, and allowed us to remove input neurons without laboriously re-training the models. This straightforward modeling approach may not capture all of the information present in the local population, but its performance sets a lower bound on predictive information in the population. A small subset of neurons has high mean-squared error of predictions, and we speculate that we have not sampled enough of the population to predict these neurons given that total variance explained scales with population size. Neuron activity was only predicted from the in-degree neurons, rather than the entire population, reducing the number of parameters by 52+/-22% across neurons, though in either case, parameters scale by order N. These results substantiate previous models showing that the collection of pairwise relationships in large populations can explain complex activity patterns (Schneidman et al. 2006). The predictions of single-trial activity that are obtained from local activity are, on one hand, a description of inter-neuronal dependencies in recurrent networks, and on the other hand, are capable of recapitulating trial-averaged tuning properties, extending these dependencies to stimulus encoding. We further explored this idea by identifying a specific motif of pairwise correlations underlying accurate predictions of neuron activity on single-trials.

Functional consequences of triplet motifs

The middle-man motif underlying the most accurate predictions is a specific pattern of pairwise correlations and represents a higher order feature of covariability. Local populations have been shown to contain such high-order correlations, but were not seen at larger spatial scales (Ohiorhenuan et al. 2010; Ohiorhenuan and Victor 2011); further experimentation is necessary to

test whether functional triplet motifs occur across thousands of microns. The finding that middle-man motifs underlie the best predictions of activity may be initially surprising from the connection pattern of the motif. Compared to the fan-in motif, which has two input connections, the middle-man has one input and one output connection, and acts to route connections from its input to its target. However, the motifs were quantified using the clustering coefficient, which normalizes motif count by the total number of possible motifs (i.e. high fan-in clustering doesn't correspond to high in-degree). The functional significance of any given triangle motif clustering is unknown and is likely dependent on the underlying system represented by the graph. In a neural network, middle-man clustering may indicate a hub-like property common in the brain (van den Heuvel and Sporns 2013), having a combination of convergent and divergent correlations. The cycle motif similarly has a combination of convergence and divergence, yet its clustering has a negative effect on prediction accuracy. The difference in the two motifs is the direction of connections between neighbors. In these networks, it may be that cycle clustering reflects recurrent, redundant correlations reducing our ability to predict activity from the population. Conversely, the middle-man motif is isometric to fan-in and fan-out motifs, and could allow for transfer of information between motifs, and in turn increasing predictive power. Finally, this result demonstrates that in addition to providing insights into synaptic mechanisms underlying dynamics (Chambers and MacLean 2016), network science can also provide insights into predictions of single trial neural responses as we have demonstrated here.

V1 is the first stage in which visual information is encoded in densely recurrent cortical networks. Thus, in order to understand activity patterns in V1, one must take into account visual

drive as well as local network activity. We have provided a quantitative comparison of the relative influences of these two factors in awake, ambulating mice. Local network effects dominate on single trials, highlighting the importance of investigating cortical computation from a population perspective in order to understand how information is encoded in single trials. Populations of neurons exhibit emergent properties beyond the sum of their individual neurons and connections, and we use the analytic framework of graph theory to begin unraveling this emergent structure.

Materials and Methods

Animals and surgery

All procedures were performed in accordance and approved by the Institutional Animal Care and Use Committee at the University of Chicago. Data was collected from C57BL/6J mice of either sex (n=4 female; 4 male) expressing transgene Tg(Thy1-GCaMP6s)GP4.12Dkim (Jackson Laboratory) between ages P84 – P191. After induction of anesthesia with isoflurane (induction at 4%, maintenance at 1-1.5%), a 3mm diameter cranial window was implanted above V1 by stereotaxic coordinates and cemented in place alongside a custom titanium headbar. Mice recovered for at least 8 days before intrinsic signal imaging to identify V1 followed by two-photon data collection.

Intrinsic Signal V1 identification

Boundaries of V1 were identified by intrinsic signal imaging post-surgery (Kalatsky and Stryker 2003) (Fig III-1A, left). Mice were anesthetized with isoflurane and head-clamped under a CCD

camera (Qimaging Retiga-SRV). A vertical white-bar stimulus (100% contrast, 0.125Hz) was repeatedly presented on an LED monitor (AOC G2460) approximately 20cm from the contralateral eye while capturing cortical reflectance under 625nm illumination. The retinotopic mapping of V1 was then estimated at each pixel from the phase of peak reflectance driven by increases in activity-dependent blood flow.

Data collection

Two-photon imaging was collected from awake, head-fixed mice on a linear treadmill. Running speed was measured with a rotary encoder attached to the treadmill axle. A L2/3 field of view (roughly 800μm diameter) in V1 was identified with galvanometer-mirror raster scanning (Cambridge Technologies; 6215H). Once a suitable field of view was found, raster scans (1Hz) were continuously acquired for roughly 10 minutes alongside visual stimulation. Neurons (n=72–347 per field of view) were then automatically identified using custom image processing software for imaging during visual stimulation using Heuristically Optimal Path Scanning (Sadovsky et al. 2011) at 25–33 Hz. All imaging was performed at 910nm (Coherent; Chameleon Ultra) with a 20X 1.1NA Olympus objective and GaAsP PMT (Hamamatsu; H10770A-40). Field of view size was estimated by fitting circles to a single raster scan of 15μm fluorescent microbeads and used for each dataset, though true field of view size may vary up to 8% across datasets from realignment of laser beam path.

Visual presentation

Drifting grating stimuli were presented on an ASUS VG248QE, 20cm from the contralateral eye at 60Hz; 60cd/m2. The mean luminance was measured and gamma correction was performed and confirmed using a luminance meter. Square-wave gratings were shown at 80% contrast, 2Hz, 0.04 cyc/deg at 12 evenly spaced directions. Gratings were presented for 5s, interleaved with 3s mean-luminance grey screen. Three repetitions of each orientation were presented in a pseudo-random order, resulting in a roughly 5min stimulus movie. The grating order was preserved between movie presentations, and mice were shown 8-11 repetitions of the movie (24-33 repetitions of each direction).

Data acquisition and pre-processing

Stimulus presentation was monitored with a photodiode (Thorlabs) and synchronized with running speed and imaging frames at 2kHz. For each neuron, baseline fluorescence was estimated from raw fluorescence by thresholding to eliminate spike-induced fluorescence transients and smoothed with a 4th-order, 81-point Savitzky-Golay filter. Fluorescence time-series were then normalized to percent change from baseline (dF/F_0) using this time-varying baseline. Spike inference from fluorescence traces was performed using the OASIS algorithm (Friedrich, Zhou, and Paninski 2017) implemented using software made freely available (github.com/j-friedrich/OASIS). Inference outputs probability of spiking at each time point. As commonly done (Vogelstein et al. 2010), probabilities were thresholded to obtain a binary spike train.

Response properties and tuning classification

Neurons were classified as visually responsive if the mean response to any grating was significantly greater than the mean response across all grey periods by Dunnett-corrected one-way ANOVA (alpha = 0.01). In these analyses, each trial response is the mean fluorescence across the entire grating presentation (5000ms), or the last half of the grey presentation (1500ms) to allow for fluorescence from grating responses to decay. Responsive neurons were then tested for statistically significant orientation- or direction-tuned responses according to the trial vectors in orientation or direction spaces (see Mazurek, Kager, and Van Hooser 2014 for more detailed methods). For significantly tuned responses, tuning curves were then fit with an asymmetric-circular Gaussian to significantly tuned neurons. Tuning curve parameters (baseline, tuning width, peak amplitudes, and preferred direction) were fit repeatedly using randomized initial conditions. The parameter set that minimized mean-squared error was maintained.

Signal correlations

For each neuron, the mean responses in each trial (using the same time windows as for tuning classification) for either stimulus or grey trials were used as response vectors. Typically, these response vectors had 360 elements (12 directions, 30 trials each). The correlation coefficient between each pair of these vectors was used to compute a pairwise correlation matrix for the grating and grey conditions. We did not shuffle responses, therefore these values measure the combination of signal and noise correlations.

Partial correlation matrices

For each pair of neurons, pairwise-correlation was computed as the mean partial correlation between their fluorescence across movies while accounting for three variables. This was

computed with a built-in MATLAB function (partialcorr.m). We computed these correlations on time-varying traces of activity, rather than time-averaged trial activity as done in previous correlation analysis. Controlling for a single variable can be computed with the following equation, where $r_{x,y}$ is the correlation coefficient between time-varying fluorescence traces *x(t)*, *y(t), and* $r_{x,y|z}$ denotes the partial correlation between *x(t)* and *y(t)*, controlling for *z(t)*:

$$r_{x,y|z} = \frac{r_{x,y} - r_{x,z}r_{y,z}}{\sqrt{1-r_{x,z}^2}\sqrt{1-r_{y,z}^2}}$$

Here, *x(t)* is the fluorescence trace of the 'input' neuron, *y(t)* is the fluorescence trace of the 'output' neuron, and *z(t)* represents the three control fluorescence traces. Controlling for more than one variable can be achieved by successive iterations of this procedure. This was computed for each movie, or 5-minute block of grating presentations that was repeated in each experiment. The first two control variables are the mean response of the two neurons in all other movies, accounting for the stimulus-driven responses. This is similar to normalizing by the mean response as in traditional noise correlation estimates. The third control variable is the within-movie mean fluorescence of all other neurons and was included to control for population-wide covariability, for example running speed effects. Furthermore, the cross-correlogram between the two neurons' fluorescence traces, averaged across movies, was used to compute directionality of the correlation. The time-lag of the cross-correlogram global maximum determined the direction and lag of the edge. If the lag was zero, the correlation was bidirectional. If the lag was greater than 500ms (roughly 14 imaging frames), no edge was included.

Graph analysis

The partial-correlation matrix could equivalently be analyzed as a directed, weighted graph. Open source software (Gephi) was used for visualization, with node layout determined by the Yifan-Hu algorithm and tuned by hand. Edge weights less than 0.05 were set to zero for visualization clarity. Erdos-Reyni (ER) null graphs were generated for each graph to match the mean connection probability. The mean directed clustering coefficient across nodes was calculated across 50 ER graphs and averaged for comparison with data. Clustering coefficients were computed with binary matrices (nonzero weights set to one).

Modeling neuron fluorescence

To model neuron responses, we used a linear weighting of the fluorescence of every in-degree using the weights in the partial-correlation matrix. At each time point, a weighted sum was calculated, resulting in a time-varying predicted fluorescence trace. Because different numbers of input neurons and varying calcium transient amplitudes, the modeled trace was then fit with a linear offset and a gain to minimize mean-squared error with the true fluorescence trace. These two parameters were not changed when input neurons were removed (7C, 8A, 8B). For tuned neurons, we also recomputed the trial-averaged tuning response of modeled activity. The fluorescence over a grating presentation was averaged, then trials were averaged over directions to obtain a mean direction-response. The sum of these vectors in direction space gave the model-estimated tuning vector, and the cosine similarity with the data-derived tuning vector was used to quantify the reconstruction of the tuning properties. Cosine similarity was computed in direction or orientation space according to each neuron's tuning properties. To compute total population

variance explained, modeled traces were subtracted from the data traces to obtain residuals, and the ratio of total sum of squares across neurons were subtracted from one as

$$1-\frac{\sum_i \sum_t (X_i(t)-\tilde{X}_i(t))^2}{\sum_i \sum_t (X_i(t))^2}$$

where $X_i(t)$ is the time-varying fluorescence trace of neuron *i,* and $X_i(t)$ tilde is the predicted trace.

Optimal weight estimation

Optimal weights for all incoming edges were computed for each neuron by LASSO regression as

$$\min_{\beta,\beta_0}\left(\frac{1}{2N}\sum_{i=1}^{N}\left(y_i-\beta_0-x_i^T\beta\right)^2+\lambda\sum_{j=1}^{p}\left|\beta_j\right|\right)$$

For weights β, and offset β_0, with modeled neuron fluorescence as y and input neuron activity as x. Weight estimation and 5-fold cross validation to estimate MSE standard error was performed with MATLAB R2016a implementation. The maximum regularization parameter (λ) whose mean-squared error did not exceed the standard error of the minimum MSE was used to find the set of optimal weights.

References

Alon, Uri. 2007. “Network Motifs: Theory and Experimental Approaches.” *Nature Reviews Genetics* 8 (6): 450.

Arieli, Amos, Alexander Sterkin, Amiram Grinvald, and A. D. Aertsen. 1996. “Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses.” *Science* 273 (5283): 1868–1871.

Averbeck, Bruno B., Peter E. Latham, and Alexandre Pouget. 2006. “Neural Correlations, Population Coding and Computation.” *Nature Reviews Neuroscience* 7 (5): 358.

Buonomano, Dean V., and Wolfgang Maass. 2009. “State-Dependent Computations: Spatiotemporal Processing in Cortical Networks.” *Nature Reviews Neuroscience* 10 (2): 113.

Chambers, Brendan, and Jason N. MacLean. 2016. “Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks.” *PLoS Computational Biology* 12 (8): e1005078.

Cohen, Marlene R., and Adam Kohn. 2011. “Measuring and Interpreting Neuronal Correlations.” *Nature Neuroscience* 14 (7): 811.

Cohen, Marlene R., and John HR Maunsell. 2009. “Attention Improves Performance Primarily by Reducing Interneuronal Correlations.” *Nature Neuroscience* 12 (12): 1594.

Cohen, Marlene R., and William T. Newsome. 2008. “Context-Dependent Changes in Functional Circuitry in Visual Area MT.” *Neuron* 60 (1): 162–173.

Doiron, Brent, Ashok Litwin-Kumar, Robert Rosenbaum, Gabriel K. Ocker, and Krešimir Josić. 2016. “The Mechanics of State-Dependent Neural Correlations.” *Nature Neuroscience* 19 (3): 383.

Durand, Séverine, Ramakrishnan Iyer, Kenji Mizuseki, Saskia de Vries, Stefan Mihalas, and R. Clay Reid. 2016. “A Comparison of Visual Response Properties in the Lateral Geniculate Nucleus and Primary Visual Cortex of Awake and Anesthetized Mice.” *Journal of Neuroscience* 36 (48): 12144–12156.

Ecker, Alexander S., Philipp Berens, R. James Cotton, Manivannan Subramaniyan, George H. Denfield, Cathryn R. Cadwell, Stelios M. Smirnakis, Matthias Bethge, and Andreas S. Tolias. 2014. “State Dependence of Noise Correlations in Macaque Primary Visual Cortex.” *Neuron* 82 (1): 235–248.

Ecker, Alexander S., Philipp Berens, Andreas S. Tolias, and Matthias Bethge. 2011. “The Effect of Noise Correlations in Populations of Diversely Tuned Neurons.” *Journal of Neuroscience* 31 (40): 14272–14283.

Fagiolo, Giorgio. 2007. “Clustering in Complex Directed Networks.” *Physical Review E* 76 (2): 026107.

Friedrich, Johannes, Pengcheng Zhou, and Liam Paninski. 2017. “Fast Online Deconvolution of Calcium Imaging Data.” *PLoS Computational Biology* 13 (3): e1005423.

Heeger, David J. 1992. “Normalization of Cell Responses in Cat Striate Cortex.” *Visual Neuroscience* 9 (2): 181–197.

Heuvel, Martijn P. van den, and Olaf Sporns. 2013. “Network Hubs in the Human Brain.” *Trends in Cognitive Sciences* 17 (12): 683–696.

Hu, Yu, Joel Zylberberg, and Eric Shea-Brown. 2014. “The Sign Rule and beyond: Boundary Effects, Flexibility, and Noise Correlations in Neural Population Codes.” *PLoS Computational Biology* 10 (2): e1003469.

Hubel, David H., and Torsten N. Wiesel. 1962. “Receptive Fields, Binocular Interaction and Functional Architecture in the Cat’s Visual Cortex.” *The Journal of Physiology* 160 (1): 106–154.

Kalatsky, Valery A., and Michael P. Stryker. 2003. “New Paradigm for Optical Imaging: Temporally Encoded Maps of Intrinsic Signal.” *Neuron* 38 (4): 529–545.

Ko, Ho, Sonja B. Hofer, Bruno Pichler, Katherine A. Buchanan, P. Jesper Sjöström, and Thomas D. Mrsic-Flogel. 2011. “Functional Specificity of Local Synaptic Connections in Neocortical Networks.” *Nature* 473 (7345): 87.

Lin, I.-Chun, Michael Okun, Matteo Carandini, and Kenneth D. Harris. 2015. “The Nature of Shared Cortical Variability.” *Neuron* 87 (3): 644–656.

Luczak, Artur, and Jason N. MacLean. 2012. “Default Activity Patterns at the Neocortical Microcircuit Level.” *Frontiers in Integrative Neuroscience* 6: 30.

Lund, Jennifer S., Alessandra Angelucci, and Paul C. Bressloff. 2003. “Anatomical Substrates for Functional Columns in Macaque Monkey Primary Visual Cortex.” *Cerebral Cortex* 13 (1): 15–24.

Mazurek, Mark, Marisa Kager, and Stephen D. Van Hooser. 2014. “Robust Quantification of Orientation Selectivity and Direction Selectivity.” *Frontiers in Neural Circuits* 8: 92.

McLelland, Douglas, Pamela M. Baker, Bashir Ahmed, and Wyeth Bair. 2010. “Neuronal Responses during and after the Presentation of Static Visual Stimuli in Macaque Primary Visual Cortex.” *Journal of Neuroscience* 30 (38): 12619–12631.

Meshulam, Leenoy, Jeffrey L. Gauthier, Carlos D. Brody, David W. Tank, and William Bialek. 2017. “Collective Behavior of Place and Non-Place Neurons in the Hippocampal Network.” *Neuron* 96 (5): 1178–1191.

Montijn, Jorrit S., Guido T. Meijer, Carien S. Lansink, and Cyriel MA Pennartz. 2016. “Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective.” *Cell Reports* 16 (9): 2486–2498.

Montijn, Jorrit S., Martin Vinck, and Cyriel Pennartz. 2014. “Population Coding in Mouse Visual Cortex: Response Reliability and Dissociability of Stimulus Tuning and Noise Correlation.” *Frontiers in Computational Neuroscience* 8: 58.

Moreno-Bote, Rubén, Jeffrey Beck, Ingmar Kanitscheider, Xaq Pitkow, Peter Latham, and Alexandre Pouget. 2014. “Information-Limiting Correlations.” *Nature Neuroscience* 17 (10): 1410.

Niell, Cristopher M., and Michael P. Stryker. 2010. “Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex.” *Neuron* 65 (4): 472–479.

Nigam, Sunny, Masanori Shimono, Shinya Ito, Fang-Chin Yeh, Nicholas Timme, Maxym Myroshnychenko, Christopher C. Lapish, Zachary Tosi, Pawel Hottowy, and Wesley C. Smith. 2016. “Rich-Club Organization in Effective Connectivity among Cortical Neurons.” *Journal of Neuroscience* 36 (3): 670–684.

Ohiorhenuan, Ifije E., Ferenc Mechler, Keith P. Purpura, Anita M. Schmid, Qin Hu, and Jonathan D. Victor. 2010. “Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks.” *Nature* 466 (7306): 617.

Ohiorhenuan, Ifije E., and Jonathan D. Victor. 2011. “Information-Geometric Measure of 3-Neuron Firing Patterns Characterizes Scale-Dependence in Cortical Networks.” *Journal of Computational Neuroscience* 30 (1): 125–141.

Okun, Michael, Nicholas A. Steinmetz, Lee Cossell, M. Florencia Iacaruso, Ho Ko, Péter Barthó, Tirin Moore, Sonja B. Hofer, Thomas D. Mrsic-Flogel, and Matteo Carandini. 2015. “Diverse Coupling of Neurons to Populations in Sensory Cortex.” *Nature* 521 (7553): 511.

Olshausen, Bruno A., and David J. Field. 2006. “What Is the Other 85 Percent of V1 Doing.” *23 Problems in Systems Neuroscience*, 182–221.

Oswald, Anne-Marie M., Brent Doiron, John Rinzel, and Alex D. Reyes. 2009. “Spatial Profile and Differential Recruitment of GABAB Modulate Oscillatory Activity in Auditory Cortex.” *Journal of Neuroscience* 29 (33): 10321–10334.

Perin, Rodrigo, Thomas K. Berger, and Henry Markram. 2011. “A Synaptic Organizing Principle for Cortical Neuronal Groups.” *Proceedings of the National Academy of Sciences* 108 (13): 5419–5424.

Petersen, Carl CH, and Sylvain Crochet. 2013. “Synaptic Computation and Sensory Processing in Neocortical Layer 2/3.” *Neuron* 78 (1): 28–48.

Pillow, Jonathan W., Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M. Litke, E. J. Chichilnisky, and Eero P. Simoncelli. 2008. “Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population.” *Nature* 454 (7207): 995.

Rakic, Pasko. 2009. “Evolution of the Neocortex: A Perspective from Developmental Biology.” *Nature Reviews Neuroscience* 10 (10): 724.

Reimer, Jacob, Emmanouil Froudarakis, Cathryn R. Cadwell, Dimitri Yatsenko, George H. Denfield, and Andreas S. Tolias. 2014. “Pupil Fluctuations Track Fast Switching of Cortical States during Quiet Wakefulness.” *Neuron* 84 (2): 355–362.

Rosenbaum, Robert, Matthew A. Smith, Adam Kohn, Jonathan E. Rubin, and Brent Doiron. 2017. “The Spatial Structure of Correlated Neuronal Variability.” *Nature Neuroscience* 20 (1): 107.

Rubino, Doug, Kay A. Robbins, and Nicholas G. Hatsopoulos. 2006. “Propagating Waves Mediate Information Transfer in the Motor Cortex.” *Nature Neuroscience* 9 (12): 1549.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. 2011. “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.” *Journal of Neurophysiology* 106 (3): 1591–1598.

Sadovsky, Alexander J., and Jason N. MacLean. 2013. “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry.” *Journal of Neuroscience* 33 (35): 14048–14060.

Schneidman, Elad, Michael J. Berry II, Ronen Segev, and William Bialek. 2006. “Weak Pairwise Correlations Imply Strongly Correlated Network States in a Neural Population.” *Nature* 440 (7087): 1007.

Schölvinck, Marieke L., Aman B. Saleem, Andrea Benucci, Kenneth D. Harris, and Matteo Carandini. 2015. “Cortical State Determines Global Variability and Correlations in Visual Cortex.” *Journal of Neuroscience* 35 (1): 170–178.

Seriès, Peggy, Peter E. Latham, and Alexandre Pouget. 2004. “Tuning Curve Sharpening for Orientation Selectivity: Coding Efficiency and the Impact of Correlations.” *Nature Neuroscience* 7 (10): 1129.

Shadlen, Michael N., and William T. Newsome. 1998. “The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding.” *Journal of Neuroscience* 18 (10): 3870–3896.

Smith, Matthew A., and Adam Kohn. 2008. “Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex.” *Journal of Neuroscience* 28 (48): 12591–12603.

Timme, Nicholas, Wesley Alford, Benjamin Flecker, and John M. Beggs. 2014. “Synergy, Redundancy, and Multivariate Information Measures: An Experimentalist’s Perspective.” *Journal of Computational Neuroscience* 36 (2): 119–140.

Tolhurst, David J., J. Anthony Movshon, and Andrew F. Dean. 1983. “The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex.” *Vision Research* 23 (8): 775–785.

Truccolo, Wilson, Leigh R. Hochberg, and John P. Donoghue. 2010. “Collective Dynamics in Human and Monkey Sensorimotor Cortex: Predicting Single Neuron Spikes.” *Nature Neuroscience* 13 (1): 105.

Vogelstein, Joshua T., Adam M. Packer, Timothy A. Machado, Tanya Sippy, Baktash Babadi, Rafael Yuste, and Liam Paninski. 2010. “Fast Nonnegative Deconvolution for Spike Train Inference from Population Calcium Imaging.” *Journal of Neurophysiology* 104 (6): 3691–3704.

Watts, Duncan J., and Steven H. Strogatz. 1998. “Collective Dynamics of ‘Small-World’ Networks.” *Nature* 393 (6684): 440.

Yu, Shan, Hongdian Yang, Hiroyuki Nakahara, Gustavo S. Santos, Danko Nikolić, and Dietmar Plenz. 2011. “Higher-Order Interactions Characterized in Cortical Activity.” *Journal of Neuroscience* 31 (48): 17514–17526.

Zohary, Ehud, Michael N. Shadlen, and William T. Newsome. 1994. “Correlated Neuronal Discharge Rate and Its Implications for Psychophysical Performance.” *Nature* 370 (6485): 140.

Zylberberg, Joel. 2017. “Untuned but Not Irrelevant: A Role for Untuned Neurons in Sensory Information Coding.” *BioRxiv*, 134379.

CHAPTER IV: GENERAL DISCUSSION

Summary of Findings

I implemented state of the art experimental techniques to simultaneously record from hundreds of neurons in order to understand how populations of interconnected neurons cooperatively generate patterns of activity. I combined well-established procedures for implantation of cranial windows in adult mice, which provided optical access to cerebral cortex *in vivo*, with an innovating scanning technique developed in my laboratory (Sadovsky et al. 2011) to achieve high scan rates. Through these windows, I recorded roughly 200 neurons at 30Hz on average—a combination of quantity and speed rarely rivaled with two photon microscopy. I studied patterns of activity observed spontaneously and during visual stimulation alongside previously acquired *in vitro* data of spontaneous activity in larger populations with slightly longer framerates (Sadovsky and MacLean 2013). In both cases, activity in the local population was sufficient to explain a large extent of trial-to-trial variability. Due to differences in the experimental and analysis methodology between Chapters II and III, the interpretations diverge slightly.

In a slice preparation of primary auditory cortex, emergent dynamics occur spontaneously. The local network under these conditions show highly variable, but far from completely random, patterns of activity. Across the entire network, I showed that the progression of mean spike times across neurons (mean global progression, or MGP) could be reproduced by Poisson neurons with spike rates matching the population spike histogram. Sequences of spikes were identified from pairs of strongly connected neurons in the functional network. These sequences occurred more reliably beyond what could be accounted for by the population spike histogram. Neurons

participated in many sequences, and the neurons with the highest membership occurred at the end of sequence, demonstrating an overall convergence. On single trials, sequences probabilistically unfolded throughout population activity, highlighting a key feature of Hebbian assembly-like dynamics in neocortex. Despite the overall convergence, the subsets of sequences active on single trials displayed a balance of convergence and divergence which I speculate could play a role in activity propagation.

In Chapter III, I collected data from awake, ambulating mice viewing drifting grating stimuli. An understudied aspect of primary visual cortex is the contributions and impact of neurons that do not modulate their activity, or are untuned, to a given stimulus parameter such as grating direction. The untuned neurons constitute roughly half of the local population in mouse V1 and show a similar distribution of firing rate as tuned neurons. The subpopulations of tuned and untuned neurons shows distinct profiles of correlations within and between one another. Interestingly, the correlations changed within the subpopulations comparing activity between visual stimulation and a mean-luminance grey screen, but the correlations between subpopulations were unchanged. By studying the structure of functional networks from the tuned and untuned local populations, I found that a specific motif (the middle man triangle) of functional connectivity was correlated with how well I could predict a neuron's activity on single trials. Clustering of a different motif has been shown to support synaptic integration leading to post-synaptic spike generation and spike propagation (Chambers and MacLean 2016). Here I show that motif clustering leads to improved predictive information capable of recapitulating single-trial activity as well as tuning properties in tuned neurons.

The functional network, built from the set of pairwise interactions in the network, helped to reveal the structure of sequence spike patterns. This pattern is built from pairwise relationships but has more demanding statistical requirements, representing a form of higher order interactions. The pairwise interactions themselves can organize into higher-order patterns, or motifs, which was explored in the context of neural representations of visual stimuli *in vivo*. Interactions within a population of neurons can be complex, but in its simplest form, the interactions can be represented by the set of pairwise interactions. Here, I have taken two distinct approaches to extend these pairwise interactions into more complex relationships comprising larger groups of neurons. These approaches both show that coordination between groups of neurons in the local population are useful in describing trial-to-trial variability.

Describing Variability in Microcircuit Dynamics

Experimentalists have a restricted view of the brain, typically limited to one brain area, recording only a subset of neurons at a time in the region. Repeated presentations of the same stimulus evoke variable responses, and ongoing brain activity preceding and during the stimulus is helpful in explaining this variability. To ultimately understand how perception and behavior takes place in real time, models of neural encoding must explain the sources of this variability and how it impacts perception and behavior. One way to accomplish this is to extend representations of neural activity beyond average responses to describe dynamics on single trials. While the brain has dense connectivity across and within brain areas, the local population in a few hundred micrometers of a given neuron is an important factor for capturing variability in population

dynamics. Though many other factors, such as brain state (Kisley and Gerstein 1999), attention (Cohen and Maunsell 2009), or possibly motor efference copies, are known to affect responses, relative contributions of these factors are not yet known. In this thesis, I took two complementary approaches to describe local population activity and capture variability across trials. These descriptions help to quantify how much variability is left over after accounting for local population dependencies.

In Chapter II, I studied the capacity and propensity for the local population to generate reliable sequences of spikes. While restricting all possible spike patterns to only sequences of spikes decreases the space of potential patterns, there are still millions of possible sequences of length three from a population of several hundred, and this scales combinatorially with sequence length. I leveraged the information in the functional network to identify candidate sequences to test for statistical significance. From the set of sequences I identified, only a subset were present on single trials. This allowed the variability across trials to be resolved as specific subsets of sequences. Outside the scope of this work, though an interesting future direction, is the possibility of predicting how sequences will diverge and converge given the spiking history in the population. With respect to the mean global progression, successful prediction of spike times from prior sequences would substantiate the sequence spike pattern as a way to capture population variability. Besides stimulus preferences, the factors determining when and how often will a neuron spike on a given trial are not fully understood. One significant factor is the ongoing activity in the local population. Previous data supports this claim (Arieli et al. 1995; Arieli et al.

1996), though the authors were analyzing voltage sensitive dyes that could not resolve individual neurons.

Chapter III dissects the patterns of covariability that facilitate predictions of single-trial activity from the local population with single cell resolution. The pairwise relationships studied were similar to noise correlations—shared fluctuations across trials that are not stimulus dependent—though they were computed differently to accommodate untuned neurons that inherently have no stimulus dependence. Weighting local neuron activity by the strength of these correlations could predict time-varying activity. A question emerging from this result is whether local activity is causing spikes in the given neuron through recurrent connections or whether they are receiving common inputs that drive correlated activity. Inputs are often lagged by several time frames, according to the method for computing the correlations, but many neighboring neurons have zero lag. Inputs with zero lag are unlikely to be causal, though it is not impossible since the temporal resolution data is above 30ms—well within the time frame for synaptic transmission. The analysis is fundamentally correlational, however, and it is very unlikely that these predictions are predominantly causal. A key result from this chapter is that accurate predictions rely on clustering of a specific triplet motif. The mechanism underlying how this motif beneficially impacts predictive power will need to be addressed in future work. While graph theoretical approaches are generalizable, interpretations and mechanisms require domain specificity. The mechanisms underlying specific motifs depend on what the nodes and edges in the graph represent, and in the case of neural networks, this remains a new area of research. Nonetheless,

structured interrelationships in the local population, whether through the lens of spike patterns or motifs of covariability, are necessary to capture trial-to-trial variability.

Multineuronal Spike Patterns

Spatiotemporal dynamics in neural populations are not completely random, and one framework for describing the structure in the dynamics is with spike patterns. The specific forms of these patterns, which must repeat in some fashion, are difficult to identify, as they can take place with repeated features over time, across neurons, or across trials. A pattern must be repeated to truly be a pattern, but population activity is noisy, so identifying a repetition depends strongly on assumptions made by the researcher. Similar to rate coding versus spike-time coding depending on the time scale in question, a pattern can repeat with millisecond precision or it can repeat only approximately depending on a chosen time scale. All this assumes a specific pattern in mind as well, which is not necessarily a safe assumption. Perhaps only two specific spike patterns are biologically motivated and been extensively studied—spike sequences and spike synchrony. However, neural networks can learn to detect arbitrary spike patterns (Kappel et al. 2015), and some researchers have taken an assumption-free approach, attempting to identify arbitrary spike patterns in the data (Gansel and Singer 2012; Harrison, Amarasingham, and Truccolo 2015; Torre et al. 2016).

Without making strong assumptions about the neural code, identification of repeated spike patterns could help reveal principles of neural computation by simply observing what is present in the data. However, recording spike trains in parallel leads to a combinatorial explosion in the

number of possible patterns, such that testing each one becomes infeasible. Even with the computational power to search for repeats of every possible pattern, statistical tests to validate whether the patterns could have arisen by chance are difficult to design and can be prone to errors (Mokeichev et al. 2007; Roxin, Hakim, and Brunel 2008). Despite these difficulties, progress continues to be made developing tools to identify spike patterns in population recordings (Gansel and Singer 2012; Harrison, Amarasingham, and Truccolo 2015; Torre et al. 2016). The functional implications of identified patterns, other than synchrony and sequences, are currently understudied, however.

Spike sequences represent a fundamental property of neural activity—propagation of spikes. In mammals, the physiological mechanism underlying a stimulus followed by a behavioral response is carried out by a cascade of neuronal spikes propagating from sensory to motor systems. This spike propagation is facilitated by spike-time dependent plasticity, which reinforces the likelihood of spike propagation between two neurons that spike in sequence. At a larger scale, spike-time dependent plasticity serves as a synaptic mechanism for an influential hypothesis in neuroscience, the Hebbian cell assembly and the assembly phase sequence (Hebb 1961). Spike synchrony is necessary for spike propagation under most conditions, since synapses are weak relative to the depolarization necessary for action potential generation (Diesmann, Gewaltig, and Aertsen 1999; Kumar, Rotter, and Aertsen 2010). For this reason, spike synchrony has been interpreted as the instantiation of a single cell assembly, representing an ensemble of functionally related neurons that spike together. Synchronous spikes as a pattern can be studied at two levels of detail. Simply summing spikes in the time bin will give an aggregate measure of spike

synchrony without regard for which neurons are active (Ostojic, Brunel, and Hakim 2009; Louis, Borgelt, and Grün 2010). However, neurons are diverse and can be individually impactful, so maintaining neuron identity can be advantageous. Synchronous spike patterns that maintain neuron identity are often referred to as binary words (Osborne et al. 2008; Palmer et al. 2015), and can be a powerful representation for investigating neural coding.

At the scale of the microcircuit, spike propagation is viewed as a progression of activity across the population (Pastalkova et al. 2008; Harvey, Coen, and Tank 2012; Luczak, McNaughton, and Harris 2015). Chapter II shows how spike sequences in small ensembles of neurons can resolve variability in this population progression. Structured spatiotemporal patterns can alternatively be represented as more tractable, lower-dimensional manifolds (Gallego et al. 2017; Elsayed and Cunningham 2017), alleviating the difficulty in pattern identification. However, a downstream brain area is only capable of acting on the spikes and cannot access the manifold representations. Therefore, the implementation of neural computation must act on the spikes themselves. In general, the forms and diversity of spike patterns are not known, but synchronous spikes and sequences of spikes represent the two fundamental patterns from which more complex patterns are built.

Functional Networks

The synaptic connections that form microcircuits occur between pairs of neurons. Analogous to the physical, synaptic connections that constitute the structural network, functional networks are defined by the pairwise relationships between neural activity patterns. These relationships can be

take many different forms with an increasing diversity of metrics used on two sets of spike trains (Chambers et al. 2017). Given a metric, the resulting measure of spike train relationships between every pair of neurons forms a square matrix, often referred to as a 'weight matrix.' Functional networks serve as a convenient and powerful tool to study relationships in population dynamics. Chapters II and III take advantage of functional networks to different ends. Chapter II uses a time-lagged normalized count weight matrix to extend pairwise relationships to search for larger ensembles of neurons. Chapter III uses a partial correlation weight matrix to measure pairwise activity relationships and directly harnesses the weights and patterns of functional connections to predict a neurons' activity on single-trials.

The most common metric to compare pairwise activity is the Pearson correlation coefficient, though there are many different ways to process the data prior to computing the correlation, with strong differences in results (Cohen and Kohn 2011). These correlations are symmetric resulting in a symmetric weight matrix. Other metrics, such as transfer entropy (Timme et al. 2014), result in an asymmetric weight matrix. From correlational and information-theory based metrics like mutual information or transfer entropy, researchers can also apply normalizations to these raw weight matrices. For example, the weight between two neurons can be normalized by the root mean squared correlation of the two neurons with all other neurons in the population (Chambers et al. 2017). This will down-weight correlations from neurons which seem to be broadly correlated with the population, emphasizing unique pairwise relationships. While the underlying synaptic network plays a role in shaping population activity patterns, how the structural network influences functional relationships is still understudied (Chambers et al. 2017). Weight matrices

are also synonymous with adjacency matrices, defining the edges in a graph. Using a graph theoretic framework, functional networks can be studied from a variety of perspectives (Sadovsky and MacLean 2013; van den Heuvel and Sporns 2013; Nigam et al. 2016). However, the sets of possible analyses on graphs strongly depends on the properties of the edges. Weighted edges, as are present in many functional networks, cannot be analyzed in the same way as unweighted edges. Historically, mathematicians have focused on the development of analytic tools in unweighted networks, due to their simplicity, though efforts are moving toward generalizing these tools (Fagiolo 2007; Abdallah 2011). Directed and undirected edges also play a large role in the analysis, with directed-edge graphs, or asymmetric weight matrices, having larger relevance to functional networks. The functional networks used in this thesis are both weighted, directed matrices.

The choice of a metric to construct a functional network should be informed by multiple factors, such as the data acquisition method and the overarching scientific goals. In chapter II, I constructed functional networks out of normalized, frame-lagged spikes between two neurons. This pairwise spiking relationship represents the building block of longer chains of lagged spikes, or spike sequences. The weight matrix in this case was used in generating candidate ensembles of neurons that were more likely to spike in sequence than a random ordering of neurons. In chapter III, a functional network was created using a partial correlation based metric. Rather than using the functional networks to make an informed search through potential sequences, the partial correlations, controlling for signal-dependent activity and population-wide cofluctuations, were used to directly predict single-trial activity. The weights in the functional

graph yielded predictions that were close to optimal weights, and interestingly, the variance in prediction quality was correlated with a motif of functional connections. Despite the differences in experimental preparation, measurement of the weight matrices, and purpose of the functional network, the complete set of pairwise relationships were crucial in describing variability of activity in the population across trials.

Future directions

Chapters II and III both resolve the variability in population dynamics, but they do so using different analytic toolsets and under different experimental conditions. This research is highly interdisciplinary, allowing for further research along many directions. A straightforward extension of these approaches would be to apply one study's analysis to the other's experimental protocol. Graph theory has already been applied to the functional networks in Chapter II (Sadovsky and MacLean 2013), but for different purposes. The method to create edges between neurons also differed and was similar to the normalized count method (Pajevic and Plenz 2009). Different calculations of functional connections reveal distinct properties of the dynamics (Chambers et al. 2017), so the partial-correlation based method used in Chapter III may lead to novel results about how the *in vitro* network produces patterned activity. The *in vitro* preparation can be used to study stimulus-evoked properties of neurons through electrode stimulation of intact thalamocortical projections (Runfeldt, Sadovsky, and MacLean 2014; Sederberg, Palmer, and MacLean 2015), but the response properties are less well defined than orientation tuning in response to drifting gratings. Nevertheless, comparison of spontaneous and stimulus-evoked responses has shown evidence that they exist within the same subspace of activity (Luczak,

Barthó, and Harris 2009). The underlying strategy of computing covariability while controlling for stimulus and population-wide effects could be applied to the *in vitro* microcircuit. Alternatively, *in vivo* activity in V1 or other cortical areas could be distilled into its ensembles of sequences. This approach inherently captures time-varying activity, allowing for the study of sequence composition during time-varying drifting gratings. This approach could help explain how the dynamics of populations represent simple visual stimuli by comparing sequences at the onset of a grating to those immediately prior to grating offset, as well as those during the transitions between grey screens and gratings. The interplay between tuned and untuned subpopulations deserves further exploration, as well. One possible way to extend the approach of Chapter III is with different visual stimuli. The transition between different stimuli, and its effect on dynamics can be investigated with the framework of spike sequences. For example, grating stimuli that abruptly change their direction may have diverse effects on dynamics, depending on the orientation. An acute direction change is likely to activate a similar subpopulation of tuned neurons, while an obtuse direction change will activate a distinct subpopulation. How the same microcircuit behaves during these transitions may be helpful in unravelling how more complex, dynamic stimuli are represented. It is likely that the ongoing activity in the untuned subpopulation at the point of transition plays an important role in shaping how the activity patterns are routed to different tuned subpopulations. Other stimulus parameters beyond grating direction modulate visual cortical neurons. For example, the nature and influence of the untuned neurons in the context of spatial frequency and temporal frequency tuning is an interesting future direction.

Animals interact with a dynamically changing world and the visual system is widely believed to be adapted to represent ethologically relevant visual stimuli (Marr 1970). Another future direction of this work will be to record population activity in response to natural stimuli. Populations have different spiking statistics in response to these stimuli (Felsen et al. 2005), and characterizing the trial-to-trial variability in response to these stimuli will have ethological relevance. It is possible that deep understanding of microcircuit representations of complex, dynamic stimuli requires more sophisticated analytic tools. An advantage of the functional graph representation of dynamics is its relative simplicity. These graphs summarize time-varying dynamics as a static weight matrix. Temporal graphs, on the other hand, have edges that change over time (Michail 2015), naturally extending functional graphs to a dynamic representation, but their flexibility comes at the cost of complexity.

Current technology in genetic engineering and optical imaging opens avenues for new experiments to extend the results in this thesis. Specific cell classes and subtypes can be specifically targeted for imaging or manipulation. Possible targets include all inhibitory neurons or a specific subclass, such as parvalbumin-expressing interneurons or somatostatin neurons. Among excitatory neurons, subclasses can be defined by their projection sites, such as cortico-cortical and cortical-thalamic projecting neurons, or pyramidal neurons in a specific cortical layer. In M1, cortico-spinal neurons have direct impact on voluntary motor control; understanding how the local microcircuitry impacts their activity is necessary to define the specific functions of cortical motor areas. Pharmacological and optical techniques exist for cell-specific manipulation (Rogan and Roth 2011). Typically, genetically-targeted neurons are

silenced, though they can also be activated. Finally, improved optical techniques can provide a more complete picture of the local microcircuit. Spatial light modulators and acousto-optic deflectors for two-photon microscopy give a three-dimensional view into the microcircuit with high temporal resolution (Nikolenko et al. 2008; Grewe et al. 2010). Data collected with this setup will avoid a limitation of Chapter III, which is restricted to layers 2/3.

Closing

Neocortical activity patterns take shape within populations of recurrently connected neurons. Particularly at the scale of several hundred microns, the local population determines and is dependent on the activity of its constituent neurons. These local populations comprise microcircuits that are hypothesized to be generalizable functional component of cerebral cortex. Rapidly developing technologies are providing unprecedented views into living and working brains, however, there is no current consensus for how to describe the observed activity in large populations. Descriptions beyond single cells and pairs of cells, encompassing ensembles of neurons are likely necessary to capture the complexity of population spike patterns. The difficulty in finding a fundamental neural code stems in part because there appears to be multiple valid encoding schemes, ranging from pure firing rate codes to spike-timing codes. An information-representation scheme may be specific to some cortical areas. Alternatively, firing rate and time-based schema transition smoothly across different timescales, and different timescales can simultaneously represent different features of sensory information, known as a multiplexed code (Panzeri et al. 2010).

The potential for a multiplexed code motivates a deeper exploration of new descriptions and representations of multineuronal spike patterns. The current understanding of neural encoding is not complete or sufficient, and new encoding schemes are not mutually exclusive in a multiplexed code. The interrelationships between neurons can reveal the boundaries of these encoding schemes, delineating the patterns a population does and does not produce. By extending pairwise relationships to spike sequences and interaction motifs among larger ensembles of neurons, I was able to capture trial-to-trial variability of population activity. These are promising approaches for describing neural representations of stimuli and motor commands on single trials. Through multidisciplinary efforts, we are nearing the cusp of bridging our understanding of single-neuron physiology to large-scale, biological substrates of behavior and psychology. The neocortical microcircuit is the setting for this revolutionary union of neuroscientific disciplines.

References

Abdallah, Sherief. 2011. "Generalizing Unweighted Network Measures to Capture the Focus in Interactions." *Social Network Analysis and Mining* 1 (4): 255–69.

Arieli, A., D. Shoham, R. Hildesheim, and A. Grinvald. 1995. "Coherent Spatiotemporal Patterns of Ongoing Activity Revealed by Real-Time Optical Imaging Coupled with Single-Unit Recording in the Cat Visual Cortex." *Journal of Neurophysiology* 73 (5): 2072–93.

Arieli, Amos, Alexander Sterkin, Amiram Grinvald, and A. D. Aertsen. 1996. "Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses." *Science* 273 (5283): 1868–1871.

Chambers, Brendan, Maayan Levy, Joseph B. Dechery, and Jason N. MacLean. 2017. "Ensemble Stacking Mitigates Biases in Inference of Synaptic Connectivity." *Network Neuroscience* 2 (1): 60–85.

Chambers, Brendan, and Jason N. MacLean. 2016. “Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks.” *PLoS Computational Biology* 12 (8): e1005078.

Cohen, Marlene R., and Adam Kohn. 2011. “Measuring and Interpreting Neuronal Correlations.” *Nature Neuroscience* 14 (7): 811.

Cohen, Marlene R., and John HR Maunsell. 2009. “Attention Improves Performance Primarily by Reducing Interneuronal Correlations.” *Nature Neuroscience* 12 (12): 1594.

Diesmann, Markus, Marc-Oliver Gewaltig, and Ad Aertsen. 1999. “Stable Propagation of Synchronous Spiking in Cortical Neural Networks.” *Nature* 402 (6761): 529.

Elsayed, Gamaleldin F., and John P. Cunningham. 2017. “Structure in Neural Population Recordings: An Expected Byproduct of Simpler Phenomena?” *Nature Neuroscience* 20 (9): 1310–18.

Fagiolo, Giorgio. 2007. “Clustering in Complex Directed Networks.” *Physical Review E* 76 (2): 026107.

Felsen, Gidon, Jon Touryan, Feng Han, and Yang Dan. 2005. “Cortical Sensitivity to Visual Features in Natural Scenes.” *PLOS Biology* 3 (10): e342.

Gallego, Juan A., Matthew G. Perich, Lee E. Miller, and Sara A. Solla. 2017. “Neural Manifolds for the Control of Movement.” *Neuron* 94 (5): 978–984.

Gansel, Kai S., and Wolf Singer. 2012. “Detecting Multineuronal Temporal Patterns in Parallel Spike Trains.” *Frontiers in Neuroinformatics* 6: 18.

Grewe, Benjamin F., Dominik Langer, Hansjörg Kasper, Björn M. Kampa, and Fritjof Helmchen. 2010. “High-Speed *in Vivo* Calcium Imaging Reveals Neuronal Network Activity with near-Millisecond Precision.” *Nature Methods* 7 (5): 399–405.

Harrison, Matthew T., Asohan Amarasingham, and Wilson Truccolo. 2015. “Spatiotemporal Conditional Inference and Hypothesis Tests for Neural Ensemble Spiking Precision.” *Neural Computation* 27 (1): 104–150.

Harvey, Christopher D., Philip Coen, and David W. Tank. 2012. “Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task.” *Nature* 484 (7392): 62.

Hebb, D. O. 1961. *Organization of Behavior. New York: Science Editions*. Inc.

Heuvel, Martijn P. van den, and Olaf Sporns. 2013. “Network Hubs in the Human Brain.” *Trends in Cognitive Sciences* 17 (12): 683–696.

Kappel, David, Stefan Habenschuss, Robert Legenstein, and Wolfgang Maass. 2015. “Network Plasticity as Bayesian Inference.” *PLOS Computational Biology* 11 (11): e1004485.

Kisley, Michael A., and George L. Gerstein. 1999. “Trial-to-Trial Variability and State-Dependent Modulation of Auditory-Evoked Responses in Cortex.” *Journal of Neuroscience* 19 (23): 10451–10460.

Kumar, Arvind, Stefan Rotter, and Ad Aertsen. 2010. “Spiking Activity Propagation in Neuronal Networks: Reconciling Different Perspectives on Neural Coding.” *Nature Reviews Neuroscience* 11 (9): 615.

Louis, Sebastien, Christian Borgelt, and Sonja Grün. 2010. “Complexity Distribution as a Measure for Assembly Size and Temporal Precision.” *Neural Networks*, Analysis and Modeling of Massively Parallel Neural Signals, 23 (6): 705–12.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. 2009. “Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations.” *Neuron* 62 (3): 413–425.

Luczak, Artur, Bruce L. McNaughton, and Kenneth D. Harris. 2015. “Packet-Based Communication in the Cortex.” *Nature Reviews Neuroscience* 16 (12): 745.

Marr, D. 1970. “A Theory of Neocortex.” *Proc Roy Soc Lond B* 176 (1043): 161–234.
Michail, Othon. 2015. “An Introduction to Temporal Graphs: An Algorithmic Perspective.” *ArXiv:1503.00278 [Cs]*, March.

Mokeichev, Alik, Michael Okun, Omri Barak, Yonatan Katz, Ohad Ben-Shahar, and Ilan Lampl. 2007. “Stochastic Emergence of Repeating Cortical Motifs in Spontaneous Membrane Potential Fluctuations in Vivo.” *Neuron* 53 (3): 413–425.

Nigam, Sunny, Masanori Shimono, Shinya Ito, Fang-Chin Yeh, Nicholas Timme, Maxym Myroshnychenko, Christopher C. Lapish, et al. 2016. “Rich-Club Organization in Effective Connectivity among Cortical Neurons.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 36 (3): 670–84.

Nikolenko, Volodymyr, Brendon O. Watson, Roberto Araya, Alan Woodruff, Darcy S. Peterka, and Rafael Yuste. 2008. “SLM Microscopy: Scanless Two-Photon Imaging and Photostimulation Using Spatial Light Modulators.” *Frontiers in Neural Circuits* 2.

Osborne, Leslie C., Stephanie E. Palmer, Stephen G. Lisberger, and William Bialek. 2008. “The Neural Basis for Combinatorial Coding in a Cortical Population Response.” *Journal of Neuroscience* 28 (50): 13522–31.

Ostojic, Srdjan, Nicolas Brunel, and Vincent Hakim. 2009. “Synchronization Properties of Networks of Electrically Coupled Neurons in the Presence of Noise and Heterogeneities.” *Journal of Computational Neuroscience* 26 (3): 369.

Pajevic, Sinisa, and Dietmar Plenz. 2009. “Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches.” *PLOS Computational Biology* 5 (1): e1000271.

Palmer, Stephanie E., Olivier Marre, Michael J. Berry, and William Bialek. 2015. “Predictive Information in a Sensory Population.” *Proceedings of the National Academy of Sciences* 112 (22): 6908–13.

Panzeri, Stefano, Nicolas Brunel, Nikos K. Logothetis, and Christoph Kayser. 2010. “Sensory Neural Codes Using Multiplexed Temporal Scales.” *Trends in Neurosciences* 33 (3): 111–20.

Pastalkova, Eva, Vladimir Itskov, Asohan Amarasingham, and György Buzsáki. 2008. “Internally Generated Cell Assembly Sequences in the Rat Hippocampus.” *Science* 321 (5894): 1322–1327.

Rogan, Sarah C., and Bryan L. Roth. 2011. “Remote Control of Neuronal Signaling.” *Pharmacological Reviews* 63 (2): 291–315.

Roxin, Alex, Vincent Hakim, and Nicolas Brunel. 2008. “The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons.” *Journal of Neuroscience* 28 (42): 10734–10745.

Runfeldt, Melissa J., Alexander J. Sadovsky, and Jason N. MacLean. 2014. “Acetylcholine Functionally Reorganizes Neocortical Microcircuits.” *Journal of Neurophysiology* 112 (5): 1205–1216.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. 2011. “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.” *Journal of Neurophysiology* 106 (3): 1591–1598.

Sadovsky, Alexander J., and Jason N. MacLean. 2013. “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry.” *Journal of Neuroscience* 33 (35): 14048–14060.

Sederberg, Audrey J., Stephanie E. Palmer, and Jason N. MacLean. 2015. “Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity.” *Journal of Neurophysiology* 113 (7): 2921–2933.

Timme, Nicholas, Wesley Alford, Benjamin Flecker, and John M. Beggs. 2014. “Synergy, Redundancy, and Multivariate Information Measures: An Experimentalist’s Perspective.” *Journal of Computational Neuroscience* 36 (2): 119–140.

Torre, Emiliano, Carlos Canova, Michael Denker, George Gerstein, Moritz Helias, and Sonja Grün. 2016. “ASSET: Analysis of Sequences of Synchronous Events in Massively Parallel Spike Trains.” *PLoS Computational Biology* 12 (7): e1004939.