THE UNIVERSITY OF CHICAGO

EMERGENT COOPERATION SUPPORTS SYNAPTIC INTEGRATION, SELECTING ROUTES FOR ACTIVITY IN CORTICAL 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
BRENDAN CHAMBERS

CHICAGO, ILLINOIS
AUGUST 2016

COPYRIGHT

Chapter III is under peer review.

TABLE OF CONTENTS

Chapter 1. On the collective synaptic dynamics of neocortical populations

Chapter 2. Multineuronal activity patterns identify selective synaptic connections under realistic experimental constraints

Chapter 3. Higher-order synaptic interactions coordinate dynamics in recurrent networks

LIST OF FIGURES

ABSTRACT

The nervous system underlies sensation and behavioral response, linking perception to action across diverse animal systems. Within a given species, a multitude of nested recurrent loops transform incoming sensory input into ecologically appropriate motor patterns. A key structure of interest is the mammalian neocortex, because neocortex is intimately associated with sophisticated, flexible response selection. Mammalian neocortex is characterized by stereotyped anatomical organization, suggesting the possibility of generic algorithms for processing information across sensory modalities and species. Ultimately, understanding information processing in the neocortex is paramount for progress in human technology, medicine, and philosophy.

Connectivity is the substrate for spiking activity in neocortex. The relationship between connectivity and emergent activity is sometimes assumed to be straightforward, but in fact, given the state of current knowledge, knowing connectivity is not sufficient for predicting coordinated population responses. The complexity and diversity of synaptic mechanisms defy simple interpretations.

Therefore, it is crucial to map population activity itself, and begin to delineate the reliable organizing principles that characterize the flow of activity through neocortical populations. Activity mapping is the key to eventual progress in understanding computation itself. In this thesis, I investigate whether activity mapping is practical under typical experimental constraints, and explore what activity mapping can tell us about the organization of neuronal firing. Building on those results, I then discover a surprising but intuitive relationship between a distributed connectivity scheme and higher-order correlations in population dynamics.

CHAPTER 1

ON THE COLLECTIVE SYNAPTIC DYNAMICS OF NEOCORTICAL POPULATIONS

INTRODUCTION

In its entirety, the nervous system needs to generate bodily responses and movement on the basis of sensory inputs and internal state. The brain is an interconnected system characterized by complexity at multiple scales. Its circuitry processes information via nested, recurrent networks, with a given subnetwork containing a multiplicity of *input* and *output neurons* routing ongoing activity. Computation is presumed to be implemented through relationships between input and output activity. Therefore, understanding how activity propagates through local network is a crucial prerequisite for progress in understanding neocortical information processing.

Neurons do not transmit activity faithfully like telephone relays. Individual synaptic connections are weak, and collective responses are difficult to predict even when connection patterns are known *a priori*. But populations of neurons do respond with surprising regularities. In fact, responses in local neocortex are selective and organized even in the absence of input activity. Activity in isolated tissue is surprisingly similar to evoked activity in vitro[1] and in vivo[2]. These population firing patterns are a crucial area of study, because they are the substrate through which inputs shape output firing. Local neocortical populations are an important bridge between cellular- and systems-level function.

---

[1] MacLean et al., “Internal Dynamics Determine the Cortical Response to Thalamic Stimulation.”

[2] Luczak, Barthó, and Harris, “Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations.”

FOUNDATIONAL INVESTIGATIONS OF POPULATION RESPONSES

In 1972, Phillip W. Anderson laid the foundation for complex-systems science when he described *emergent properties* for the first time in a publication titled *More is Different*. At the time, Anderson was employed at Bell Laboratories and Cambridge University. He had already revolutionized condensed matter physics in a handful of ways, but he had yet to be awarded his Nobel prize. He argued that large systems of interacting elements can respond very differently as a collective than would the constituent elements in isolation. Out of spontaneous interactions, through broken symmetry, quantitative differences can turn into qualitative differences. The discovery of emergent phenomena needed "no new knowledge of fundamental laws". In fact, they "would have been extremely difficult to derive synthetically from those laws" governing lower-level components[3].

Collections of neurons can interact in ways that are difficult to predict from observations of single cells or connections. For sixty-five years, since Donald Hebb published his theory of cell assemblies[4], evidence has accumulated as to whether and how selective groups of neurons cooperate to process information. Some of the most important work in this area to date was undertaken by a small number of under-appreciated investigators, beginning their careers at the time Anderson published *More is Different*.

In the late 1960s, George Gerstein and Donald Perkel had been investigating pairwise timing regularities in multi-unit recordings[5]. In mapping pairwise cross-correlations over many stimulus presentations, Gerstein and Perkel noticed the presence of reliable lagged timing relationships at multiple time-scales. They were attempting to link those timing patterns to

---

[3] Anderson, "More Is Different."

[4] Hebb, *The Organization of Behavior*.

[5] Gerstein and Perkel, "Simultaneously Recorded Trains of Action Potentials."

underlying connection patterns[6]. Using simulations and intuition, they argued that some of these timing relationships were informative about underlying connectivity, whereas others reflected statistical features of the stimulus itself. Of the connectivity-driven correlation structure, Gerstein and Perkel argued that some features were likely to reflect monosynaptic excitatory connections, while others were more likely to arise from shared common input and polysynaptic relationships. Contemporarily, Bruce Knight worked to interpret single-neuron recordings in the context of encoding within larger populations[7], working to extend rate-code frameworks in light of new understanding about population responses. Ultimately, these early efforts to map network responses were limited by the recording technology available, which could not provide high-density sampling of interconnected populations. Too prescient for popular success at the time, their work deserves wider recognition in the neuroscience community.

These pioneers of population neuroscience laid the foundation for our contemporary understanding, including this thesis work. Technological progress is opening new avenues to expand on the research programs laid out decades ago, which sound utterly modern today:

> The basic problems requiring investigation are those of identifying…functional groups [of neurons] and of characterizing their activity in terms of information processing or production of behavior…Because functional grouping implies some degree of temporal organization in the activities of the constituent neurons, it follows that experimental evidence of the existence of such groups may be furnished by quantitative characterizations of the mutual temporal relationships in the electrical activity of the involved neurons. The further understanding of the operating principles of neuronal groups, once identified, is crucially dependent on detailed knowledge of the underlying circuitry.

---

[6] Perkel, Gerstein, and Moore, “Neuronal Spike Trains and Stochastic Point Processes”; Gerstein and Perkel, “Mutual Temporal Relationships among Neuronal Spike Trains”; Perkel et al., “Nerve-Impulse Patterns.”

[7] Knight, “Dynamics of Encoding in a Population of Neurons”; Knight, “The Relationship between the Firing Rate of a Single Neuron and the Level of Activity in a Population of Neurons Experimental Evidence for Resonant Enhancement in the Population Response.”

> Detailed description of functional connectivity is only partially furnished by classical neuroanatomical techniques, or by a combination of electrophysiological signal tracing and such newer methods as fluorescent dye injection and scanning electron microscopy. The approach described in this paper is intended as an adjunct to these more familiar methods; it can yield functional "circuit diagrams" of the connections among the observed neorons…These techniques, statistical in nature, must be imbedded in an appropriate experimental design in order to allow inferences of biological significance to be drawn[8].

To complement the data that was available, these investigations combined network simulations with experimental data, a practice that bridged traditional disciplinary frameworks. Gunter Palm theorized about Hebbian assemblies[9], yet also supervised an oft-cited experiment to estimate the number of neurons and connections in the human brain[10]. While designing simulations, the Palm laboratory had realized that elementary data on connections per cubic millimeter were not available—in this way simulations helped interpret experiments, and experiments improved the realism of simulations, in an iterative process of mutual benefit that continues to the present day.

Throughout the 1980s and 90s, these investigators continued to make progress in understanding the structure of population responses. Simulations continued to play a key role in interpreting their data[11]. With Ad Aertson, Gerstein worked to improve methods for inferring underlying connectivity from simultaneous recordings[12]. In addition to pairwise timing relationships, they discovered the presence of higher-order response patterns in their data[13],

---

[8] Gerstein and Perkel, "Mutual Temporal Relationships among Neuronal Spike Trains."

[9] Palm, "On the Storage Capacity of an Associative Memory with Randomly Distributed Storage Elements"; Palm, "From Neural Dynamics to Cell Assemblies."

[10] Schüz and Palm, "Density of Neurons and Synapses in the Cerebral Cortex of the Mouse."

[11] Aertsen, Erb, and Palm, "Dynamics of Functional Coupling in the Cerebral Cortex."

[12] Aertsen and Gerstein, "Evaluation of Neuronal Connectivity."

[13] Gerstein and Aertsen, "Representation of Cooperative Firing Activity among Simultaneously Recorded Neurons."

meaning that neuronal interactions could not be fully understood at the pairwise levels, but also entailed coordinated timing across distributed synaptic connections. Gerstein, Aertson and Palm noticed that patterns of *effective connectivity* were dynamic, progressively changing after stimulus onset[14], and coordinated multi-neuron responses were found to contain patterns of unexpected synchrony[15].

SPIKING DYNAMICS IN ACTIVE NETWORKS

Networks of neocortical neurons exhibit irregular dynamics[16], with high variability in inter-spike intervals compared to their mean[17]. Although it can be useful to understand the consequences of random input timing[18], the irregular firing observed in vivo is difficult to achieve with random inputs; rather, irregular firing is consistent with temporally coordinated presynaptic inputs[19]. Intracellular membrane recordings in vivo reveal non-stationary frequency content at timescales relevant for synaptic integration[20]. Because timing impacts postsynaptic integration (e.g. determining whether output neurons efficaciously recruit their targets), it is important to

---

[14] Aertsen et al., "Dynamics of Neuronal Firing Correlation"; Aertsen and Gerstein, "Dynamic Aspects of Neuronal Cooperativity."

[15] Lindsey et al., "Repeated Patterns of Distributed Synchrony in Neuronal Assemblies."

[16] Brunel, "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons"; Destexhe, Rudolph, and Paré, "The High-Conductance State of Neocortical Neurons in Vivo."

[17] Calvin and Stevens, "Synaptic Noise and Other Sources of Randomness in Motoneuron Interspike Intervals"; Mainen and Sejnowski, "Reliability of Spike Timing in Neocortical Neurons"; Reich et al., "Response Variability and Timing Precision of Neuronal Spike Trains in Vivo"; Stein, Gossen, and Jones, "Neuronal Variability."

[18] Fourcaud-Trocmé et al., "How Spike Generation Mechanisms Determine the Neuronal Response to Fluctuating Inputs"; Burkitt, "A Review of the Integrate-and-Fire Neuron Model"; Vilela and Lindner, "Comparative Study of Different Integrate-and-Fire Neurons."

[19] Softky and Koch, "The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs"; Stevens and Zador, "Input Synchrony and the Irregular Firing of Cortical Neurons."

[20] Azouz and Gray, "Cellular Mechanisms Contributing to Response Variability of Cortical Neurons in Vivo."

characterize these non-random dynamical features on a single-trial basis and link them to their underlying synaptic mechanisms.

During ongoing activity, neurons are known to operate in a state of elevated conductance, with a depolarized average membrane potential[21]. Input currents from inhibitory and excitatory neurons are related by a stable ratio in their magnitudes[22]. Approximate balance in input contributions from excitatory and inhibitory neurons supports criticality[23], linked to responsiveness and dynamic range[24]. Balanced inputs with non-stationary magnitudes can support diverse correlation timescales[25], consistent with correlation dynamics observed in vivo. Fluctuations caused by synaptic input occur at multiple timescales, with synchronous excitatory bombardment especially efficacious generating an action potential[26]. Synchronous excitation is particularly effective when it occurs in a brief window preceding onset of inhibition[27].

Neurons do not fire independently; they influence one another causally and exhibit correlated timing. Of course, high average correlations have the potential to limit computation by reducing the entropy of cortical dynamics[28]. Pathological levels of correlation manifest as epilepsy and debilitating conditions of hyper-excitability. At the other extreme, completely uncorrelated

---

[21] Destexhe et al., "Fluctuating Synaptic Conductances Recreate in Vivo-like Activity in Neocortical Neurons."

[22] Shu, Hasenstaub, and McCormick, "Turning on and off Recurrent Balanced Cortical Activity"; Haider et al., "Neocortical Network Activity in Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition"; Denève and Machens, "Efficient Codes and Balanced Networks."

[23] Poil et al., "Critical-State Dynamics of Avalanches and Oscillations Jointly Emerge from Balanced Excitation/Inhibition in Neuronal Networks."

[24] Shew et al., "Neuronal Avalanches Imply Maximum Dynamic Range in Cortical Networks at Criticality."

[25] Litwin-Kumar et al., "Balanced Synaptic Input Shapes the Correlation between Neural Spike Trains."

[26] Azouz and Gray, "Adaptive Coincidence Detection and Dynamic Gain Control in Visual Cortical Neurons In Vivo."

[27] Wehr and Zador, "Balanced Inhibition Underlies Tuning and Sharpens Spike Timing in Auditory Cortex."

[28] Montani et al., "The Role of Correlations in Direction and Contrast Coding in the Primary Visual Cortex."

population responses (for all timescales of measurement) would imply a total absence of causal synaptic interactions. Information processing as a cortical system would be impossible in a population of completely independent neurons. Though sparse, pairwise and higher-order correlations are necessary hallmarks of sparse, causal synaptic interactions—and thereby, information processing.

The neuroscience community is seeing resurging interest in multi-neuron correlations[29]. Efforts have been made to distinguish correlations induced by stimulus-features from correlations induced by internal variability[30], following a theme laid out by Gerstein and Perkel[31]. The issue of relative orientation between these two sources of correlation is central to interpreting single neuron firing rates, with respect to their impact on an ideal decoder[32].

Among possible approaches to quantifying multi-neuron timing relationships, simultaneous spike-count correlations have received the most attention (e.g. [33]). In virtuoso work that combined *in vivo* imaging with *in vitro* whole-cell recordings, simultaneous correlations were verified to be revealing of likely underlying connectivity[34]. This was the central argument made by George Gerstein and Donald Perkel in 1972.

---

[29] Josić et al., “Stimulus-Dependent Correlations and Population Codes”; Renart et al., “The Asynchronous State in Cortical Circuits”; Ecker et al., “Decorrelated Neuronal Firing in Cortical Microcircuits”; Middleton et al., “Neural Correlation Is Stimulus Modulated by Feedforward Inhibitory Circuitry”; Cohen and Kohn, “Measuring and Interpreting Neuronal Correlations.”

[30] Averbeck, Latham, and Pouget, “Neural Correlations, Population Coding and Computation”; Rothschild, Nelken, and Mizrahi, “Functional Organization and Population Dynamics in the Mouse Primary Auditory Cortex.”

[31] Gerstein and Perkel, “Simultaneously Recorded Trains of Action Potentials.”

[32] Sederberg, Palmer, and MacLean, “Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity”; Luczak, Barthó, and Harris, “Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations.”

[33] de la Rocha et al., “Correlation between Neural Spike Trains Increases with Firing Rate.”

[34] Ko et al., “Functional Specificity of Local Synaptic Connections in Neocortical Networks.”

Even with high-quality empirical evidence, the modern neuroscience community remains cautious of arguments that patterns of synaptic connectivity impact population dynamics. Caution is appropriate, because there are many details left to understand about relationships between connectivity and dynamics. But the basic premise that dynamics arise from the propagation of activity through the synaptic network should not be seen as speculative[35]. This is one reason Gerstein et al. deserve to be more widely read.

CONNECTION RULES IN LOCAL SYNAPTIC NETWORKS

In contrast to our lack of progress in understanding typical patterns of activity, tremendous progress has been made understanding connection rules at the scale of area-specific cortical maps[36]. Here, as an example of progress, I discuss some recent consensus findings based on investigation of primary motor cortex (M1), which has become an important model system for characterizing cortical connectivity[37].

Based largely on work from the laboratory of Gordon Shepherd Jr, a classification scheme linking cell-type identity with projection targets is emerging. These findings have emphasized three major excitatory subclasses: Intertelencephalic neurons (IT) project to ipsi- and contralateral cortical targets. Pyramidal tract neurons (PT) send projections to subcortical targets and feature branching axons to thalamic nuclei, possibly as efference copies[38]. These PT neurons participate in higher-order thalamic loops and are characterized by Class 1 glutamatergic connections and

---

[35] Salinas and Sejnowski, “Correlated Neuronal Activity and the Flow of Neural Information”; Ohiorhenuan et al., “Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks.”
[36] Lefort et al., “The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex.”
[37] Harris and Shepherd, “The Neocortical Circuit.”
[38] Guillery and Sherman, “Branched Thalamic Afferents.”

tufted dendrites. Corticothalamic neurons (CT) reside in layer 6 and are associated with gain control and gating. CT neurons are perhaps clumsily named, since they are merely a subclass of all thalamus-projecting cells: the CT population does not include those L5B neurons with branching axons. In motor cortex, CT neurons receive a major projection from frontal areas, which tempts speculation about volitional modulation of gain.

M1 receives thalamic input from posterior thalamus nuclei (PO) including POm and VA/VL, sources of predominantly-sensory versus predominantly-motor related input activity. Axons from PO predominantly target L2/3 and L5a, while axons from VA/VL additionally target L5B and L6[39]. Interestingly, corticothalamic neurons in L6 project to VL, but they do not seem to target the VL cells with direct projections to M1. Similarly, those VL projections do not seem to directly target corticothalamic neurons in L6[40]. Branching axons from M1 pyramidal tract outputs target PO.

Important progress has also been made in understanding connectivity at the microcircuit level. The self-similar laminar organization of neocortex suggests the presence of repeated modular information processing circuits[41]. Indeed, regularities of connection direction have been found across layers, with L4 -> L2/3 -> L5 connections suggesting orderly flows of activity. Anatomical columns have been proposed as the long-sought generic cortical microcircuit[42], but it is still unclear whether these proposals imply that columnar connectivity is the appropriate fundamental-unit for understanding cortical activity. For example, activity propagates without

---

[39] Hooks et al., "Organization of Cortical and Thalamic Input to Pyramidal Neurons in Mouse Motor Cortex."

[40] Yamawaki and Shepherd, "Synaptic Circuit Organization of Motor Corticothalamic Neurons."

[41] Mountcastle, "Modality and Topographic Properties of Single Neurons of Cat's Somatic Sensory Cortex."

[42] Douglas and Martin, "Neuronal Circuits of the Neocortex"; Bastos et al., "Canonical Microcircuits for Predictive Coding."

regard for columnar orientation in dissected cortical tissue[43], single-whisker stimulation evokes distributed activity across multiple differently tuned barrel columns[44], and tuning properties can be mixed between near neighbors[45]. Columnar architecture may reflect developmental constraints and cell lineage rules for connectivity[46] without necessarily being a fundamental constituent of cortical activity[47].

A major breakthrough in understanding synaptic microcircuits resulted from studies performing multiple whole-cell recordings to test the presence and strength of synaptic connections between random pairs[48]. For L5 pyramidal neurons in juvenile rats, connection probabilities were as high as 0.2 for close neighbors, falling to 0.1 at radii of 0.15 mm, and 0.05 at radii of 0.3 mm[49]. Bidirectional connections were more abundant than expected in random networks, as were clustered neighborhoods[50]. However, the consequences of these connection patterns for active networks are not well understood.

Even though synaptic connections are relatively dense in local populations, population responses are sparse and selective[51]. It is difficult to predict emergent activity patterns even when synaptic connectivity is known. Synaptic connection patterns without the context of recent activity can elicit misleading predictions about the flow of activity, for example by leading to false

---

[43] Sadovsky and MacLean, “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry.”

[44] Ferezou, Bolea, and Petersen, “Visualizing the Cortical Representation of Whisker Touch.”

[45] Kaschube, “Neural Maps versus Salt-and-Pepper Organization in Visual Cortex”; Bonin et al., “Local Diversity and Fine-Scale Organization of Receptive Fields in Mouse Visual Cortex.”

[46] Li et al., “Clonally Related Visual Cortical Neurons Show Similar Stimulus Feature Selectivity.”

[47] Horton and Adams, “The Cortical Column.”

[48] Song et al., “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits”; Perin, Berger, and Markram, “A Synaptic Organizing Principle for Cortical Neuronal Groups.”

[49] Perin, Berger, and Markram, “A Synaptic Organizing Principle for Cortical Neuronal Groups.”

[50] Song et al., “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits.”

[51] Barth and Poulet, “Experimental Evidence for Sparse Firing in the Neocortex.”

expectations of disinhibition[52]. Ongoing activity plays a crucial role in shaping firing in the synaptic network.

In Chapter 3, I describe a clustering motif underlying reliably coordinated recruitment. Activity preferentially occurs at synaptic locations characterized by this feature in random models, and recruitment maps based on experimental recordings activity are enriched in the same directed clustering motif. These results succeed in mapping a synaptic connection motif onto its likely consequences for propagating activity. This connection motif tends to synchronize its inputs, so that their coordinated effect preferentially elicits postsynaptic action potentials. These results represent progress in predicting propagating activity on the basis of underlying connection patterns.

STUDYING CONNECTIVITY AND DYNAMICS TOGETHER

Neocortical activity is patterned by local connectivity

Connection probabilities are highest among near neighbors in the neocortex. It is likely that the numerous local connections around a given neuron do more than merely relay afferent activity to output neurons. Dissociated model systems have proven to be a fruitful method for studying local connections, encompassing neighbor interactions while facilitating the kind of experimental access that is impossible or much more difficult to achieve in vivo. These preparations divorce local connectivity from non-local influences including feedback modulation from distant cortical areas, and they enable complete experimental control over the neuromodulatory milieu. The

---

[52] Fisher et al., "A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit."

dynamics arising in these simplified systems are informative of in vivo function, because they illuminate structured responses arising from local synaptic interactions.

Progress in understanding local network dynamics has been driven by the advent of large dense population recordings, particularly two-photon imaging in acute tissue and dense multi-electrode recordings of cultured populations. These experimental preparations have enabled the large sample sizes and high signal-to-noise-ratios necessary for monitoring population responses. Even more critically, both preparations isolated neocortical networks from ongoing inputs and neuromodulation. Acute slices of sensory cortex become spontaneously active in the presence of oxygen and artificial cerebrospinal fluid[53]. Like evoked responses, spontaneous events are selective in their activation of the local population with event-specific tuning in a subset of the population. Non-event-turned neurons accounted for the majority of active cells, which were scattered heterogeneously among inactive cells[54]. These dynamics disappeared when synaptic transmission is blocked pharmacologically[55].

In cultured populations as well as acute slices, passive recordings reveal a clear transition from occasional isolated spikes to self-sustaining, patterned activity[56]. These ignition events arise from a small number of sites, thought to reflect features of their connection to the rest of the neuronal network[57]. Ignition dynamics give rise to spontaneous activity—activity intrinsic to neocortex, patterned by local connectivity.

In healthy slices 450 um in width, thalamically evoked activity encompasses a highly

---

[53] Ivanov and Zilberter, "Critical State of Energy Metabolism in Brain Slices."
[54] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry."
[55] Cossart, Aronov, and Yuste, "Attractor Dynamics of Network UP States in the Neocortex."
[56] Beggs and Plenz, "Neuronal Avalanches in Neocortical Circuits."
[57] Orlandi et al., "Noise Focusing and the Emergence of Coherent Activity in Neuronal Cultures."

similar subset of the population compared to spontaneous activity. Even in the absence of external input, the local cortical network imposes structure on responses in space and time. In fact, when thalamic connections are entirely cut, the active population of cortical neurons still recapitulates both intact spontaneous activity and thalamically evoked activity[58]. Thus, patterned firing observed in the ex vivo system is known to reflect propagation of activity through local neocortical connectivity. In vivo, auditory input was observed to elicit firing vectors that occupied a subspace of firing vectors explored during spontaneous activity[59]. These sources of evidence suggest that the role of inputs is to bias local response patterns, influencing propagating activity and impacting the timing and rate of output neuron activity. The organization of spontaneous activity reveals building blocks for network activity in vivo.

Mapping multi-neuron patterns

To someday understand how inputs control local network activity to shape output firing, it is crucial to study the organization of propagating activity. Activity is known to contain sequential structure at multiple timescales[60], and sequence identity seems to causally impact behavioral choices[61]. There is ongoing debate about the relationship between different timescales, but it is

---

[58] MacLean et al., “Internal Dynamics Determine the Cortical Response to Thalamic Stimulation.”

[59] Luczak, Barthó, and Harris, “Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations.”

[60] Luczak et al., “Sequential Structure of Neocortical Spontaneous Activity in Vivo”; Ito et al., “Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model.”

[61] Harvey, Coen, and Tank, “Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task.”

likely that both non-linear dendritic processing[62] and propagating activity[63] can map sequential firing patterns onto output firing (or silence).

Another approach which is gaining traction is to map pairwise timing relationships, to generate *functional circuit diagrams*[64]. As activity-mapping becomes more popular, specialized inference approaches are being newly developed and refined[65]. These functional maps are a natural abstraction for studying propagating activity, but ground truth underlying connectivity is usually not available experimentally. In Chapter 2, I investigate the hypothesis that functional relationships can reveal sites of likely synaptic connectivity, with special attention to feasible experimental conditions. Functional connectivity has also been shown to reorganize during neuromodulation[66], opening new avenues to understanding modulatory function. This promises to be a growth area for the neuroscience community.

Correlations are known to be non-stationary in time and span multiple timescales. Thus, timescales inherent to computing correlation levels (width of the Gaussian convolution, binning of data) influence the interpretability of measurements. In Chapter 2, I argue that reliable timing

---

[62] Hawkins and Ahmad, "Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex."
[63] Laje and Buonomano, "Robust Timing and Motor Patterns by Taming Chaos in Recurrent Neural Networks."
[64] Ito et al., "Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model"; Stetter et al., "Model-Free Reconstruction of Excitatory Neuronal Connectivity from Calcium Imaging Signals"; Gururangan, Sadovsky, and MacLean, "Analysis of Graph Invariants in Functional Neocortical Circuitry Reveals Generalized Features Common to Three Areas of Sensory Cortex."
[65] Mishchencko, Vogelstein, and Paninski, "A Bayesian Approach for Inferring Neuronal Connectivity from Calcium Fluorescent Imaging Data"; Gerhard et al., "Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone"; Orlandi et al., "First Connectomics Challenge."
[66] Runfeldt, Sadovsky, and MacLean, "Acetylcholine Functionally Reorganizes Neocortical Microcircuits."

relationships are most informative of underlying connectivity when they are measured at timescales similar to those governing synaptic integration.

The importance of simulation

The key to progress in understanding multi-neuron patterning is close interaction between experiment and simulation. Experimental methods are becoming ever more incisive, more powerful to support or invalidate hypotheses about function. Computational models, because they are constructed based on controlled design choices, identify gaps in our understanding and point the way for new experiments. Simulations of cortical networks offer omniscient knowledge of a schematized system, providing access to the numerous distributed synaptic bombardments making up simulated activity. In Chapter 3, this access is used to map the flow of spikes through simulated networks. These results make a testable prediction about patterns of propagating activity in vivo.

Network simulations are continually evolving, in tension between simplicity and realism. Importantly, real neurons are sensitive to coordinated presynaptic inputs[67]. Non-linear integrative features of neurons are important for producing realistic dynamics. Synchrony plays an important role in shaping population responses by impacting postsynaptic integration. In this work, model neurons are leaky-integrators with the addition of conductance-based connections. For a postsynaptic neuron, input currents scale their effect on postsynaptic voltage based on their reversal potentials. As a result, all else being equal, synchronous inputs sum more effectively than temporally separated inputs. The leaky aspect of membrane potential rules also contributes

---

[67] Azouz and Gray, "Dynamic Spike Threshold Reveals a Mechanism for Synaptic Coincidence Detection in Cortical Neurons in Vivo."

sensitivity to input synchrony. Sensitivity to coincident inputs are an important feature of model design underlying the findings in the third section.

PARALLEL PERSPECTIVES IN MACHINE LEARNING

Theories of computation by neurons have always been hard to test in practice. Machine learning, a sister discipline to neuroscience, operates in a complementary paradigm by emphasizing practical analysis of data. The roots of machine learning drew inspiration from neuroscience, and the two fields have been fruitfully cross-pollinating ever since. Modern artificial neural networks are revolutionizing the human world. These models are distinct from simulations of spiking neurons, and are certain to be less powerful computationally. But they are also tractable, and their record-breaking performances speak for themselves. Just as investigators of mice should also learn from cat studies, and tree shrews, and ravens, there are valuable comparative insights to be gained from progress in artificial neural networks.

The brain still wildly outstrips any man-made intelligence, and machine learning has much still to learn from neuroscience. Early artificial neurons reflected efforts to understand stimulus integration in neocortex. Those first binary neuron models summated a scaled vector and compared it to a threshold for output activity, with inhibitory input operating as a veto[68]. These artificial neurons could be assembled into collections that operated as logic gates. By corollary, these units could in principle implement general computers. The Perceptron model introduced independently weighted inputs and graded inhibition[69]. Networks of perceptrons were more flexible than McCulloch-Pitt neurons, free to tune their function by changing individual weights. Connection

---

[68] McCulloch and Pitts, “A Logical Calculus of the Ideas Immanent in Nervous Activity.”

[69] Rosenblatt, *The Perceptron, a Perceiving and Recognizing Automaton Project Para*; Rosenblatt, “The Perceptron.”

topologies defining causal interactions among perceptrons were simple, by comparison to the arcane topologies and large numbers of units needed for McCulloch-Pitt networks. Perceptrons were arranged into three layers with feedforward connection between adjacent layers. These networks were trained with backpropagation[70]. Three-layer artificial networks can perform arbitrarily flexible functions, at least if they are allowed an unconstrained number of nodes per layer[71]. The difference between uniform weighting versus heterogeneous weighting had major consequences for the design of network topology.

Artificial neural networks with more than three layers were speculated to potentially outperform their three-layer cousins, given equal training data. However, backpropagation was not feasible for training deep networks, with each layer adding additional nonlinearities to the error signal. After decades of declining interest in artificial networks, increases in computational resources suddenly yielded explosive progress in *deep learning*[72]. It was even shown that a Hebbian procedure could be used to pre-train these networks[73]. Standardized connection topologies were introduced, further increasing performance and efficiency of implementation[74]. Deep feedforward networks are now regularly breaking records in computer vision.

Still, mammalian nervous systems routinely accomplish feats that elude human designed algorithms. Artificial neurons typically do not integrate inputs over multiple time-steps, and lack recurrent connections. As a result, these systems may not be as appropriate for processing dynamic

---

[70] Williams and Hinton, "Learning Representations by Back-Propagating Errors."

[71] Hornik, Stinchcombe, and White, "Multilayer Feedforward Networks Are Universal Approximators."

[72] LeCun, Bengio, and Hinton, "Deep Learning."

[73] Hinton, "Learning Multiple Layers of Representation."

[74] Krizhevsky, Sutskever, and Hinton, "ImageNet Classification with Deep Convolutional Neural Networks"; Simonyan and Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition."

stimuli. Yet natural vision, control of movement, auditory scene analysis—typical ethological problems need to be sensitive to input structure at multiple temporal scales.

Recurrent networks are potentially a more powerful network design for information processing. Intriguingly, recurrent networks can be trained with Hebbian rules to implement auto-encoding[75]. Hopfield networks have become a computational model for associative memory capacity[76], since they can converge to stable attractor states and also cycle through dynamic attractor patterns.

Newer recurrent network models bear less resemblance to populations of neurons, with bypasses and signal switching gates[77]. These cutting edge designs are accomplishing amazing performance, changing the state-of-the-art in natural language processing [78] including machine translation[79], captioning images[80], and even deploying sequential foveal attention to process images[81].

Artificial recurrent networks raise exciting opportunities for interplay between machine learning and computational neuroscience[82]. Exchanges between machine learning and computational neuroscience are beneficial for both fields. Hallmarks of neocortical activity like dynamic correlations and higher-order patterning are a potential source of new progress in artificial recurrent networks. How do cortical networks balance information about the past with information

---

[75] Hopfield, "Neural Networks and Physical Systems with Emergent Collective Computational Abilities."
[76] Palm and Sommer, "Information Capacity in Recurrent McCulloch–Pitts Networks with Sparsely Coded Memory States."
[77] Graves, "Generating Sequences With Recurrent Neural Networks."
[78] Graves and Jaitly, "Towards End-To-End Speech Recognition with Recurrent Neural Networks."
[79] Sutskever, Vinyals, and Le, "Sequence to Sequence Learning with Neural Networks."
[80] Mao et al., "Explain Images with Multimodal Recurrent Neural Networks."
[81] Ba, Mnih, and Kavukcuoglu, "Multiple Object Recognition with Visual Attention."
[82] Graves, "Adaptive Computation Time for Recurrent Neural Networks."

about the present with predictions about the future[83]? How is activity in recurrent networks decoded downstream[84]? What design features orchestrate routing through synaptic networks? Neuroscience can make a very practical impact on human technology if we can understand how the neural code is manifest through synaptic interactions.

SUMMARY

The brain is a complex system with a multitude of components. In this way, it is a fundamentally different scientific frontier than other historical research areas. Understanding the nervous system requires diverse interdisciplinary perspectives and levels of abstraction. One important scale for future research is to understand patterned firing in local populations. Distributed interactions for information processing have been best studied in machine learning, but artificial neural units are radically simple compared to the computational power and response complexity of a single neuron. Neighboring neurons coordinate their firing in complex beyond-pairwise fashion, the study of which requires dual attention to cellular-level electrophysiology and circuit-level functional networks.

Active synaptic networks are characterized by patterned activity, the emergent result of myriad complex interactions. These interactions are weak and selective, and they are coordinated in ways that are not well understand. The intricate flow of activity through synaptic networks results in information processing, and ultimately gives rise to outward behavior. In this work I investigate topological features of activity patterns using top-down mapping approaches and

---

[83] Palmer et al., “Predictive Information in a Sensory Population.”
[84] Buzsáki, “Neural Syntax.”

simulations of naturalistic networks. Understanding the synaptic dynamics of active circuits is essential for progress in neuroscience.

REFERENCES

Aertsen, A. M., G. L. Gerstein, M. K. Habib, and G. Palm. "Dynamics of Neuronal Firing Correlation: Modulation Of 'effective Connectivity.'" *Journal of Neurophysiology* 61, no. 5 (May 1989): 900–917.

Aertsen, Ad, Michael Erb, and Günther Palm. "Dynamics of Functional Coupling in the Cerebral Cortex: An Attempt at a Model-Based Interpretation." *Physica D: Nonlinear Phenomena* 75, no. 1 (August 1, 1994): 103–28. doi:10.1016/0167-2789(94)90278-X.

Aertsen, Ad M. H. J., and George L. Gerstein. "Dynamic Aspects of Neuronal Cooperativity: Fast Stimulus-Locked Modulations of Effective Connectivity." In *Neuronal Cooperativity*, edited by Professor Dr Jürgen Krüger, 52–67. Springer Series in Synergetics 49. Springer Berlin Heidelberg, 1991. http://link.springer.com/chapter/10.1007/978-3-642-84301-3_4.

Aertson Ad M. H. J., and George L. Gerstein. "Evaluation of Neuronal Connectivity: Sensitivity of Cross-Correlation." *Brain Research* 340, no. 2 (August 12, 1985): 341–54. doi:10.1016/0006-8993(85)90931-X.

Anderson, P. W. "More Is Different." *Science* 177, no. 4047 (August 4, 1972): 393–96. doi:10.1126/science.177.4047.393.

Averbeck, Bruno B., Peter E. Latham, and Alexandre Pouget. "Neural Correlations, Population Coding and Computation." *Nature Reviews Neuroscience* 7, no. 5 (May 2006): 358–66. doi:10.1038/nrn1888.

Azouz, R., and C. M. Gray. "Cellular Mechanisms Contributing to Response Variability of Cortical Neurons in Vivo." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 19, no. 6 (March 15, 1999): 2209–23.

Azouz, Rony, and Charles M. Gray. "Adaptive Coincidence Detection and Dynamic Gain Control in Visual Cortical Neurons In Vivo." *Neuron* 37, no. 3 (February 6, 2003): 513–23. doi:10.1016/S0896-6273(02)01186-8.

Azouz, Rony, and Charles M. Gray. "Dynamic Spike Threshold Reveals a Mechanism for Synaptic Coincidence Detection in Cortical Neurons in Vivo." *Proceedings of the National Academy of Sciences* 97, no. 14 (July 5, 2000): 8110–15. doi:10.1073/pnas.130200797.

Ba, Jimmy, Volodymyr Mnih, and Koray Kavukcuoglu. “Multiple Object Recognition with Visual Attention.” *arXiv:1412.7755 [Cs]*, December 24, 2014. http://arxiv.org/abs/1412.7755.

Barth, Alison L., and James FA Poulet. “Experimental Evidence for Sparse Firing in the Neocortex.” *Trends in Neurosciences* 35, no. 6 (2012): 345–355.

Bastos, Andre M., W. Martin Usrey, Rick A. Adams, George R. Mangun, Pascal Fries, and Karl J. Friston. “Canonical Microcircuits for Predictive Coding.” *Neuron* 76, no. 4 (November 21, 2012): 695–711. doi:10.1016/j.neuron.2012.10.038.

Beggs, John M., and Dietmar Plenz. “Neuronal Avalanches in Neocortical Circuits.” *The Journal of Neuroscience* 23, no. 35 (2003): 11167–11177.

Bonin, Vincent, Mark H. Histed, Sergey Yurgenson, and R. Clay Reid. “Local Diversity and Fine-Scale Organization of Receptive Fields in Mouse Visual Cortex.” *The Journal of Neuroscience* 31, no. 50 (December 14, 2011): 18506–21. doi:10.1523/JNEUROSCI.2974-11.2011.

Brunel, Nicolas. “Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons.” *Journal of Computational Neuroscience* 8, no. 3 (2000): 183–208.

Burkitt, A. N. “A Review of the Integrate-and-Fire Neuron Model: I. Homogeneous Synaptic Input.” *Biological Cybernetics* 95, no. 1 (April 19, 2006): 1–19. doi:10.1007/s00422-006-0068-6.

Buzsáki, György. “Neural Syntax: Cell Assemblies, Synapsembles, and Readers.” *Neuron* 68, no. 3 (November 4, 2010): 362–85. doi:10.1016/j.neuron.2010.09.023.

Calvin, William H., and CHARLES F. Stevens. “Synaptic Noise and Other Sources of Randomness in Motoneuron Interspike Intervals.” *J Neurophysiol* 31, no. 4 (1968): 574–587.

Cohen, Marlene R., and Adam Kohn. “Measuring and Interpreting Neuronal Correlations.” *Nature Neuroscience* 14, no. 7 (July 2011): 811–19. doi:10.1038/nn.2842.

Cossart, Rosa, Dmitriy Aronov, and Rafael Yuste. “Attractor Dynamics of Network UP States in the Neocortex.” *Nature* 423, no. 6937 (May 15, 2003): 283–88. doi:10.1038/nature01614.

de la Rocha, Jaime, Brent Doiron, Eric Shea-Brown, Krešimir Josić, and Alex Reyes. “Correlation between Neural Spike Trains Increases with Firing Rate.” *Nature* 448, no. 7155 (August 16, 2007): 802–6. doi:10.1038/nature06028.

Denève, Sophie, and Christian K. Machens. “Efficient Codes and Balanced Networks.” *Nature Neuroscience* 19, no. 3 (March 2016): 375–82. doi:10.1038/nn.4243.

Destexhe, A, M Rudolph, J. -M Fellous, and T. J Sejnowski. “Fluctuating Synaptic Conductances Recreate in Vivo-like Activity in Neocortical Neurons.” *Neuroscience* 107, no. 1 (November 9, 2001): 13–24. doi:10.1016/S0306-4522(01)00344-X.

Destexhe, Alain, Michael Rudolph, and Denis Paré. “The High-Conductance State of Neocortical Neurons in Vivo.” *Nature Reviews Neuroscience* 4, no. 9 (2003): 739–751.

Douglas, Rodney J., and Kevan A. C. Martin. “Neuronal Circuits of the Neocortex.” *Annual Review of Neuroscience* 27, no. 1 (2004): 419–51. doi:10.1146/annurev.neuro.27.070203.144152.

Ecker, Alexander S., Philipp Berens, Georgios A. Keliris, Matthias Bethge, Nikos K. Logothetis, and Andreas S. Tolias. “Decorrelated Neuronal Firing in Cortical Microcircuits.” *Science* 327, no. 5965 (January 29, 2010): 584–87. doi:10.1126/science.1179867.

Ferezou, Isabelle, Sonia Bolea, and Carl C. H. Petersen. “Visualizing the Cortical Representation of Whisker Touch: Voltage-Sensitive Dye Imaging in Freely Moving Mice.” *Neuron* 50, no. 4 (May 18, 2006): 617–29. doi:10.1016/j.neuron.2006.03.043.

Fisher, Dimitry, Itsaso Olasagasti, David W. Tank, Emre R. F. Aksay, and Mark S. Goldman. “A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit.” *Neuron* 79, no. 5 (September 4, 2013): 987–1000. doi:10.1016/j.neuron.2013.06.041.

Fourcaud-Trocmé, Nicolas, David Hansel, Carl van Vreeswijk, and Nicolas Brunel. “How Spike Generation Mechanisms Determine the Neuronal Response to Fluctuating Inputs.” *The Journal of Neuroscience* 23, no. 37 (December 17, 2003): 11628–40.

Gerhard, Felipe, Tilman Kispersky, Gabrielle J. Gutierrez, Eve Marder, Mark Kramer, and Uri Eden. “Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone.” *PLOS Comput Biol* 9, no. 7 (July 11, 2013): e1003138. doi:10.1371/journal.pcbi.1003138.

Gerstein, G. L., and A. M. Aertsen. “Representation of Cooperative Firing Activity among Simultaneously Recorded Neurons.” *Journal of Neurophysiology* 54, no. 6 (December 1985): 1513–28.

Gerstein, G. L., and D. H. Perkel. “Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation.” *Science (New York, N.Y.)* 164, no. 3881 (May 16, 1969): 828–30.

Gerstein, George L., and Donald H. Perkel. “Mutual Temporal Relationships among Neuronal Spike Trains.” *Biophysical Journal* 12, no. 5 (May 1972): 453–73.

Graves, Alex. "Adaptive Computation Time for Recurrent Neural Networks." *arXiv:1603.08983 [Cs]*, March 29, 2016. http://arxiv.org/abs/1603.08983.

Graves, Alex. "Generating Sequences With Recurrent Neural Networks." *arXiv:1308.0850 [Cs]*, August 4, 2013. http://arxiv.org/abs/1308.0850.

Graves, Alex, and Navdeep Jaitly. "Towards End-To-End Speech Recognition with Recurrent Neural Networks." In *ICML*, 14:1764–1772, 2014. http://www.jmlr.org/proceedings/papers/v32/graves14.pdf.

Guillery, R. W., and S. Murray Sherman. "Branched Thalamic Afferents: What Are the Messages That They Relay to the Cortex?" *Brain Research Reviews* 66, no. 1 (2011): 205–219.

Gururangan, Suchin S., Alexander J. Sadovsky, and Jason N. MacLean. "Analysis of Graph Invariants in Functional Neocortical Circuitry Reveals Generalized Features Common to Three Areas of Sensory Cortex." *PLoS Comput Biol* 10, no. 7 (July 10, 2014): e1003710. doi:10.1371/journal.pcbi.1003710.

Haider, Bilal, Alvaro Duque, Andrea R. Hasenstaub, and David A. McCormick. "Neocortical Network Activity in Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 26, no. 17 (April 26, 2006): 4535–45. doi:10.1523/JNEUROSCI.5297-05.2006.

Harris, Kenneth D., and Gordon M. G. Shepherd. "The Neocortical Circuit: Themes and Variations." *Nature Neuroscience* 18, no. 2 (February 2015): 170–81. doi:10.1038/nn.3917.

Harvey, Christopher D., Philip Coen, and David W. Tank. "Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task." *Nature* 484, no. 7392 (April 5, 2012): 62–68. doi:10.1038/nature10918.

Hawkins, Jeff, and Subutai Ahmad. "Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex." *Frontiers in Neural Circuits* 10 (March 30, 2016). doi:10.3389/fncir.2016.00023.

Hebb, Donald. *The Organization of Behavior*. Wiley & Sons, 1949.

Hinton, Geoffrey E. "Learning Multiple Layers of Representation." *Trends in Cognitive Sciences* 11, no. 10 (October 2007): 428–34. doi:10.1016/j.tics.2007.09.004.

Hooks, Bryan M., Tianyi Mao, Diego A. Gutnisky, Naoki Yamawaki, Karel Svoboda, and Gordon MG Shepherd. "Organization of Cortical and Thalamic Input to Pyramidal Neurons in Mouse Motor Cortex." *The Journal of Neuroscience* 33, no. 2 (2013): 748–760.

Hopfield, John J. "Neural Networks and Physical Systems with Emergent Collective Computational Abilities." *Proceedings of the National Academy of Sciences* 79, no. 8 (1982): 2554–2558.

Hornik, Kurt, Maxwell Stinchcombe, and Halbert White. "Multilayer Feedforward Networks Are Universal Approximators." *Neural Networks* 2, no. 5 (1989): 359–366.

Horton, Jonathan C., and Daniel L. Adams. "The Cortical Column: A Structure without a Function." *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 360, no. 1456 (2005): 837–862.

Ito, Shinya, Michael E. Hansen, Randy Heiland, Andrew Lumsdaine, Alan M. Litke, and John M. Beggs. "Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model." *PLOS ONE* 6, no. 11 (November 15, 2011): e27431. doi:10.1371/journal.pone.0027431.

Ivanov, Anton, and Yuri Zilberter. "Critical State of Energy Metabolism in Brain Slices: The Principal Role of Oxygen Delivery and Energy Substrates in Shaping Neuronal Activity." *Frontiers in Neuroenergetics* 3 (2011): 9. doi:10.3389/fnene.2011.00009.

Josić, Krešimir, Eric Shea-Brown, Brent Doiron, and Jaime de la Rocha. "Stimulus-Dependent Correlations and Population Codes." *Neural Computation* 21, no. 10 (July 27, 2009): 2774–2804. doi:10.1162/neco.2009.10-08-879.

Kaschube, Matthias. "Neural Maps versus Salt-and-Pepper Organization in Visual Cortex." *Current Opinion in Neurobiology*, Neural maps, 24 (February 2014): 95–102. doi:10.1016/j.conb.2013.08.017.

Knight, Bruce W. "Dynamics of Encoding in a Population of Neurons." *The Journal of General Physiology* 59, no. 6 (June 1, 1972): 734–66. doi:10.1085/jgp.59.6.734.

Knight, Bruce W. "The Relationship between the Firing Rate of a Single Neuron and the Level of Activity in a Population of Neurons Experimental Evidence for Resonant Enhancement in the Population Response." *The Journal of General Physiology* 59, no. 6 (June 1, 1972): 767–78. doi:10.1085/jgp.59.6.767.

Ko, Ho, Sonja B. Hofer, Bruno Pichler, Katherine A. Buchanan, P. Jesper Sjöström, and Thomas D. Mrsic-Flogel. "Functional Specificity of Local Synaptic Connections in Neocortical Networks." *Nature* 473, no. 7345 (May 5, 2011): 87–91. doi:10.1038/nature09880.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks." In *Advances in Neural Information Processing Systems*

25, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, 1097–1105. Curran Associates, Inc., 2012. http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

Laje, Rodrigo, and Dean V. Buonomano. "Robust Timing and Motor Patterns by Taming Chaos in Recurrent Neural Networks." *Nature Neuroscience* 16, no. 7 (July 2013): 925–33. doi:10.1038/nn.3405.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep Learning." *Nature* 521, no. 7553 (May 28, 2015): 436–44. doi:10.1038/nature14539.

Lefort, Sandrine, Christian Tomm, J.-C. Floyd Sarria, and Carl CH Petersen. "The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex." *Neuron* 61, no. 2 (2009): 301–316.

Li, Ye, Hui Lu, Pei-lin Cheng, Shaoyu Ge, Huatai Xu, Song-Hai Shi, and Yang Dan. "Clonally Related Visual Cortical Neurons Show Similar Stimulus Feature Selectivity." *Nature* 486, no. 7401 (2012): 118–121.

Lindsey, B. G., K. F. Morris, R. Shannon, and G. L. Gerstein. "Repeated Patterns of Distributed Synchrony in Neuronal Assemblies." *Journal of Neurophysiology* 78, no. 3 (September 1, 1997): 1714–19.

Litwin-Kumar, Ashok, Anne-Marie M. Oswald, Nathaniel N. Urban, and Brent Doiron. "Balanced Synaptic Input Shapes the Correlation between Neural Spike Trains." *PLOS Comput Biol* 7, no. 12 (December 22, 2011): e1002305. doi:10.1371/journal.pcbi.1002305.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. "Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations." *Neuron* 62, no. 3 (2009): 413–425.

Luczak, Artur, Peter Barthó, Stephan L. Marguet, György Buzsáki, and Kenneth D. Harris. "Sequential Structure of Neocortical Spontaneous Activity in Vivo." *Proceedings of the National Academy of Sciences* 104, no. 1 (January 2, 2007): 347–52. doi:10.1073/pnas.0605643104.

MacLean, Jason N., Brendon O. Watson, Gloster B. Aaron, and Rafael Yuste. "Internal Dynamics Determine the Cortical Response to Thalamic Stimulation." *Neuron* 48, no. 5 (December 8, 2005): 811–23. doi:10.1016/j.neuron.2005.09.035.

Mainen, Zachary F., and Terrence J. Sejnowski. "Reliability of Spike Timing in Neocortical Neurons." *Science* 268, no. 5216 (1995): 1503.

Mao, Junhua, Wei Xu, Yi Yang, Jiang Wang, and Alan L. Yuille. “Explain Images with Multimodal Recurrent Neural Networks.” *arXiv:1410.1090 [Cs]*, October 4, 2014. http://arxiv.org/abs/1410.1090.

McCulloch, Warren S., and Walter Pitts. “A Logical Calculus of the Ideas Immanent in Nervous Activity.” *The Bulletin of Mathematical Biophysics* 5, no. 4 (n.d.): 115–33. doi:10.1007/BF02478259.

Middleton, Jason W., Cyrus Omar, Brent Doiron, and Daniel J. Simons. “Neural Correlation Is Stimulus Modulated by Feedforward Inhibitory Circuitry.” *The Journal of Neuroscience* 32, no. 2 (January 11, 2012): 506–18. doi:10.1523/JNEUROSCI.3474-11.2012.

Mishchencko, Yuriy, Joshua T. Vogelstein, and Liam Paninski. “A Bayesian Approach for Inferring Neuronal Connectivity from Calcium Fluorescent Imaging Data.” *The Annals of Applied Statistics*, 2011, 1229–1261.

Montani, Fernando, Adam Kohn, Matthew A. Smith, and Simon R. Schultz. “The Role of Correlations in Direction and Contrast Coding in the Primary Visual Cortex.” *The Journal of Neuroscience* 27, no. 9 (February 28, 2007): 2338–48. doi:10.1523/JNEUROSCI.3417-06.2007.

Mountcastle, Vernon B. “Modality and Topographic Properties of Single Neurons of Cat’s Somatic Sensory Cortex.” *Journal of Neurophysiology* 20, no. 4 (July 1, 1957): 408–34.

Ohiorhenuan, Ifije E., Ferenc Mechler, Keith P. Purpura, Anita M. Schmid, Qin Hu, and Jonathan D. Victor. “Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks.” *Nature* 466, no. 7306 (July 29, 2010): 617–21. doi:10.1038/nature09178.

Orlandi, Javier G., Bisakha Ray, Demian Battaglia, Isabelle Guyon, Vincent Lemaire, Mehreen Saeed, Jordi Soriano, Alexander Statnikov, and Olav Stetter. “First Connectomics Challenge: From Imaging to Connectivity.” In *Neural Connectomics Workshop at ECML 2014*, 1–22, 2014. http://www.jmlr.org/proceedings/papers/v46/orlandi15.pdf.

Orlandi, Javier G., Jordi Soriano, Enrique Alvarez-Lacalle, Sara Teller, and Jaume Casademunt. “Noise Focusing and the Emergence of Coherent Activity in Neuronal Cultures.” *Nature Physics* 9, no. 9 (2013): 582–590.

Palm, Dr Günther. “From Neural Dynamics to Cell Assemblies.” In *Neural Assemblies*, edited by Dr Günther Palm, 104–16. Studies of Brain Function 7. Springer Berlin Heidelberg, 1982. http://link.springer.com/chapter/10.1007/978-3-642-81792-2_12.

Palm, G. “On the Storage Capacity of an Associative Memory with Randomly Distributed Storage Elements.” *Biological Cybernetics* 39, no. 2 (1981): 125–27. doi:10.1007/BF00336738.

Palm, G., and F. T. Sommer. “Information Capacity in Recurrent McCulloch–Pitts Networks with Sparsely Coded Memory States.” *Network: Computation in Neural Systems* 3, no. 2 (January 1, 1992): 177–86. doi:10.1088/0954-898X_3_2_006.

Palmer, Stephanie E., Olivier Marre, Michael J. Berry, and William Bialek. “Predictive Information in a Sensory Population.” *Proceedings of the National Academy of Sciences* 112, no. 22 (June 2, 2015): 6908–13. doi:10.1073/pnas.1506855112.

Perin, Rodrigo, Thomas K. Berger, and Henry Markram. “A Synaptic Organizing Principle for Cortical Neuronal Groups.” *Proceedings of the National Academy of Sciences* 108, no. 13 (2011): 5419–5424.

Perkel, Donald H., George L. Gerstein, and George P. Moore. “Neuronal Spike Trains and Stochastic Point Processes.” *Biophysical Journal* 7, no. 4 (July 1967): 419–40.

Perkel, Donald H., George L. Gerstein, Mark S. Smith, and William G. Tatton. “Nerve-Impulse Patterns: A Quantitative Display Technique for Three Neurons.” *Brain Research* 100, no. 2 (December 19, 1975): 271–96. doi:10.1016/0006-8993(75)90483-7.

Poil, Simon-Shlomo, Richard Hardstone, Huibert D. Mansvelder, and Klaus Linkenkaer-Hansen. “Critical-State Dynamics of Avalanches and Oscillations Jointly Emerge from Balanced Excitation/Inhibition in Neuronal Networks.” *The Journal of Neuroscience* 32, no. 29 (July 18, 2012): 9817–23. doi:10.1523/JNEUROSCI.5990-11.2012.

Reich, D. S., J. D. Victor, B. W. Knight, T. Ozaki, and E. Kaplan. “Response Variability and Timing Precision of Neuronal Spike Trains in Vivo.” *Journal of Neurophysiology* 77, no. 5 (May 1997): 2836–41.

Renart, Alfonso, Jaime de la Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D. Harris. “The Asynchronous State in Cortical Circuits.” *Science* 327, no. 5965 (January 29, 2010): 587–90. doi:10.1126/science.1179850.

Rosenblatt, Frank. *The Perceptron, a Perceiving and Recognizing Automaton Project Para*. Cornell Aeronautical Laboratory, 1957.

Rosenblatt, Frank. “The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain.” *Psychological Review* 65, no. 6 (1958): 386.

Rothschild, Gideon, Israel Nelken, and Adi Mizrahi. “Functional Organization and Population Dynamics in the Mouse Primary Auditory Cortex.” *Nature Neuroscience* 13, no. 3 (March 2010): 353–60. doi:10.1038/nn.2484.

Runfeldt, Melissa J., Alexander J. Sadovsky, and Jason N. MacLean. “Acetylcholine Functionally Reorganizes Neocortical Microcircuits.” *Journal of Neurophysiology* 112, no. 5 (September 1, 2014): 1205–16. doi:10.1152/jn.00071.2014.

Sadovsky, Alexander J., and Jason N. MacLean. “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry.” *The Journal of Neuroscience* 33, no. 35 (August 28, 2013): 14048–60. doi:10.1523/JNEUROSCI.1977-13.2013.

Salinas, Emilio, and Terrence J. Sejnowski. “Correlated Neuronal Activity and the Flow of Neural Information.” *Nature Reviews Neuroscience* 2, no. 8 (August 2001): 539–50. doi:10.1038/35086012.

Schüz, Almut, and Günther Palm. “Density of Neurons and Synapses in the Cerebral Cortex of the Mouse.” *The Journal of Comparative Neurology* 286, no. 4 (August 22, 1989): 442–55. doi:10.1002/cne.902860404.

Sederberg, Audrey J., Stephanie E. Palmer, and Jason N. MacLean. “Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity.” *Journal of Neurophysiology*, February 18, 2015, jn.00885.2014. doi:10.1152/jn.00885.2014.

Shew, Woodrow L., Hongdian Yang, Thomas Petermann, Rajarshi Roy, and Dietmar Plenz. “Neuronal Avalanches Imply Maximum Dynamic Range in Cortical Networks at Criticality.” *The Journal of Neuroscience* 29, no. 49 (December 9, 2009): 15595–600. doi:10.1523/JNEUROSCI.3864-09.2009.

Shu, Yousheng, Andrea Hasenstaub, and David A. McCormick. “Turning on and off Recurrent Balanced Cortical Activity.” *Nature* 423, no. 6937 (May 15, 2003): 288–93. doi:10.1038/nature01616.

Simonyan, Karen, and Andrew Zisserman. “Very Deep Convolutional Networks for Large-Scale Image Recognition.” *arXiv:1409.1556 [Cs]*, September 4, 2014. http://arxiv.org/abs/1409.1556.

Softky, W. R., and C. Koch. “The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs.” *The Journal of Neuroscience* 13, no. 1 (January 1, 1993): 334–50.

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B. Chklovskii. “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits.” *PLoS Biol* 3, no. 3 (2005): e68.

Stein, Richard B., E. Roderich Gossen, and Kelvin E. Jones. “Neuronal Variability: Noise or Part of the Signal?” *Nature Reviews Neuroscience* 6, no. 5 (2005): 389–397.

Stetter, Olav, Demian Battaglia, Jordi Soriano, and Theo Geisel. “Model-Free Reconstruction of Excitatory Neuronal Connectivity from Calcium Imaging Signals.” *PLOS Comput Biol* 8, no. 8 (August 23, 2012): e1002653. doi:10.1371/journal.pcbi.1002653.

Stevens, Charles F., and Anthony M. Zador. “Input Synchrony and the Irregular Firing of Cortical Neurons.” *Nature Neuroscience* 1, no. 3 (July 1998): 210–17. doi:10.1038/659.

Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le. “Sequence to Sequence Learning with Neural Networks.” *arXiv:1409.3215 [Cs]*, September 10, 2014. http://arxiv.org/abs/1409.3215.

Vilela, Rafael D., and Benjamin Lindner. “Comparative Study of Different Integrate-and-Fire Neurons: Spontaneous Activity, Dynamical Response, and Stimulus-Induced Correlation.” *Physical Review E* 80, no. 3 (September 21, 2009): 31909. doi:10.1103/PhysRevE.80.031909.

Wehr, Michael, and Anthony M. Zador. “Balanced Inhibition Underlies Tuning and Sharpens Spike Timing in Auditory Cortex.” *Nature* 426, no. 6965 (2003): 442–446.

Williams, DRGHR, and G. E. Hinton. “Learning Representations by Back-Propagating Errors.” *Nature* 323 (1986): 533–536.

Yamawaki, Naoki, and Gordon MG Shepherd. “Synaptic Circuit Organization of Motor Corticothalamic Neurons.” *The Journal of Neuroscience* 35, no. 5 (2015): 2293–2307.

CHAPTER 2

Multineuronal activity patterns identify selective synaptic connections under realistic experimental constraints

This work was previously published: Chambers B & MacLean JN (2015). Multineuronal activity patterns identify selective synaptic connections under realistic experimental constraints. *Journal of neurophysiology, 114*(3), 1837-1849.

ABSTRACT

Structured multineuronal activity patterns within local neocortical circuitry are strongly linked to sensory input, motor output, and behavioral choice. These reliable patterns of pairwise lagged firing are the consequence of connectivity since they are not present in rate-matched but unconnected Poisson nulls. It is important to relate multineuronal patterns to their synaptic underpinnings, but it is unclear how effectively statistical dependencies in spiking between neurons identify causal synaptic connections. To assess the feasibility of mapping function onto structure we used a network model that showed a diversity of multineuronal activity patterns and replicated experimental constraints on data acquisition. Using an iterative Bayesian inference algorithm, we detected a select subset of monosynaptic connections substantially more precisely than correlation-based inference, a common alternative approach. We found that precise inference of synaptic connections improved with increasing numbers of diverse multineuronal activity patterns in contrast to increased observations of a single pattern. Surprisingly, neuronal spiking was most effective and precise at revealing causal synaptic connectivity when the lags considered by the iterative Bayesian algorithm encompassed the timescale of synaptic conductance and integration (~10ms), rather than synaptic transmission time (~2ms), highlighting the importance of synaptic integration in driving postsynaptic spiking. Lastly, strong synaptic connections were

detected preferentially, underscoring their special importance in cortical computation. Even after simulating experimental constraints, top down approaches to cortical connectivity—from function to structure—identify synaptic connections underlying multineuronal activity. These select connections are closely tied to cortical processing.

INTRODUCTION

Synaptic connections are fundamental to neocortical computation. They are responsible for instantiating and constraining the multineuronal activity patterns that underlie sensation and behavior[1]. If we are to understand information processing, it is crucial to map cortical activity at the level of neurons and their synaptic relationships. For example, paired patch clamp recordings during quiescence have revealed dense connectivity in local excitatory networks[2]. Yet circuit spiking activity is sparse and diverse, in ways that are not easily predicted from connection patterns alone[3]. Functional connectivity maps are an important bridge between static connectivity and dynamic information processing.

It is challenging to link activity to underlying connectivity. A large number of presynaptic inputs contribute to the high conductance, depolarized state of postsynaptic neurons during circuit

---

[1] Harvey, Coen, and Tank, "Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task"; O'Connor et al., "Neural Coding during Active Somatosensation Revealed Using Illusory Touch."

[2] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits"; Neske, Patrick, and Connors, "Contributions of Diverse Excitatory and Inhibitory Neurons to Recurrent Network Activity in Cerebral Cortex."

[3] Barth and Poulet, "Experimental Evidence for Sparse Firing in the Neocortex."

activity[4], limiting the influence of any one synaptic connection [5]. The large number of spikes which comprise a multineuronal activity pattern and the impossibility of monitoring all neurons in neocortex mean that correlations in spiking between neurons are not necessarily indicative of a monosynaptic connection[6]. Nevertheless, correlation has been shown to indicate an enhanced likelihood that neurons are synaptically connected[7]. A select fraction of pairwise spike timing relationships are highly reliable[8] and do not arise by chance. These strong statistical dependencies may be particularly informative of the underlying synaptic connections that produce them.

To establish whether multineuronal activity patterns can be used in a practical way to identify causal synaptic connections, we applied a modified iterative Bayesian algorithm for inferring synaptic connectivity from population activity[9]. To evaluate its performance and inform experimental design, we built a neuronal network model comprised of leaky integrate-and-fire units with conductance-based synapses and heterogeneous weights. Model activity was poised near criticality and spiking was sparse, irregular and asynchronous. Beyond what is expected by chance, we found that model population activity includes reliable lagged timing relationships as a result of its interconnectivity, and we demonstrate that the iterative Bayesian inference method

---

[4] Brunel, "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons"; Destexhe, Rudolph, and Paré, "The High-Conductance State of Neocortical Neurons in Vivo"; MacLean et al., "Internal Dynamics Determine the Cortical Response to Thalamic Stimulation"; Watson, MacLean, and Yuste, "UP States Protect Ongoing Cortical Activity from Thalamic Inputs"; Kumar et al., "The High-Conductance State of Cortical Networks"; Neske, Patrick, and Connors, "Contributions of Diverse Excitatory and Inhibitory Neurons to Recurrent Network Activity in Cerebral Cortex."

[5] Teramae, Tsubo, and Fukai, "Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links"; Chicharro and Panzeri, "Algorithms of Causal Inference for the Analysis of Effective Connectivity among Brain Regions."

[6] Gerstein and Perkel, "Simultaneously Recorded Trains of Action Potentials."

[7] Ko et al., "Functional Specificity of Local Synaptic Connections in Neocortical Networks."

[8] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry."

[9] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."

reveals this structure, relating patterned activity to underlying connectivity. Inference improved in the presence of diverse activity—a hallmark of multineuronal patterns in local cortical networks[10]. The requirement that a neuron generally must integrate more than one input to achieve threshold[11] resulted in inference being most effective when the algorithm considered delays that closely matched the time constant of an excitatory synaptic conductance. Finally, the strongest connections were preferentially detected, consistent with their increased likelihood of driving a postsynaptic action potential.

Despite the fact that there are many impinging synaptic inputs, it is the subset providing drive necessary for the next spike in the pattern that are salient using this method. With temporally proximal pre- and postsynaptic action potentials, these are also the connections that have the capacity to undergo spike-timing dependent plasticity. Because these connections underlie propagating activity, we propose that they are particularly meaningful in the context of their multineuronal patterns, closely tied to cortical processing and learning.

MATERIALS AND METHODS

Simulated cortical networks with random connection topology

Network simulations were implemented in Python using the Brian Brain Simulator[12]. The model was modified from the integrate-and-fire *COBA* model[13]. In conductance-based models,

---

[10] Harvey, Coen, and Tank, "Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task"; Sadovsky and MacLean, "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity."

[11] Magee, "Dendritic Integration of Excitatory Synaptic Input."

[12] Goodman and Brette, "The Brian Simulator."

[13] Vogels and Abbott, "Signal Propagation and Logic Gating in Networks of Integrate-and-Fire Neurons"; Brette et al., "Simulation of Networks of Spiking Neurons."

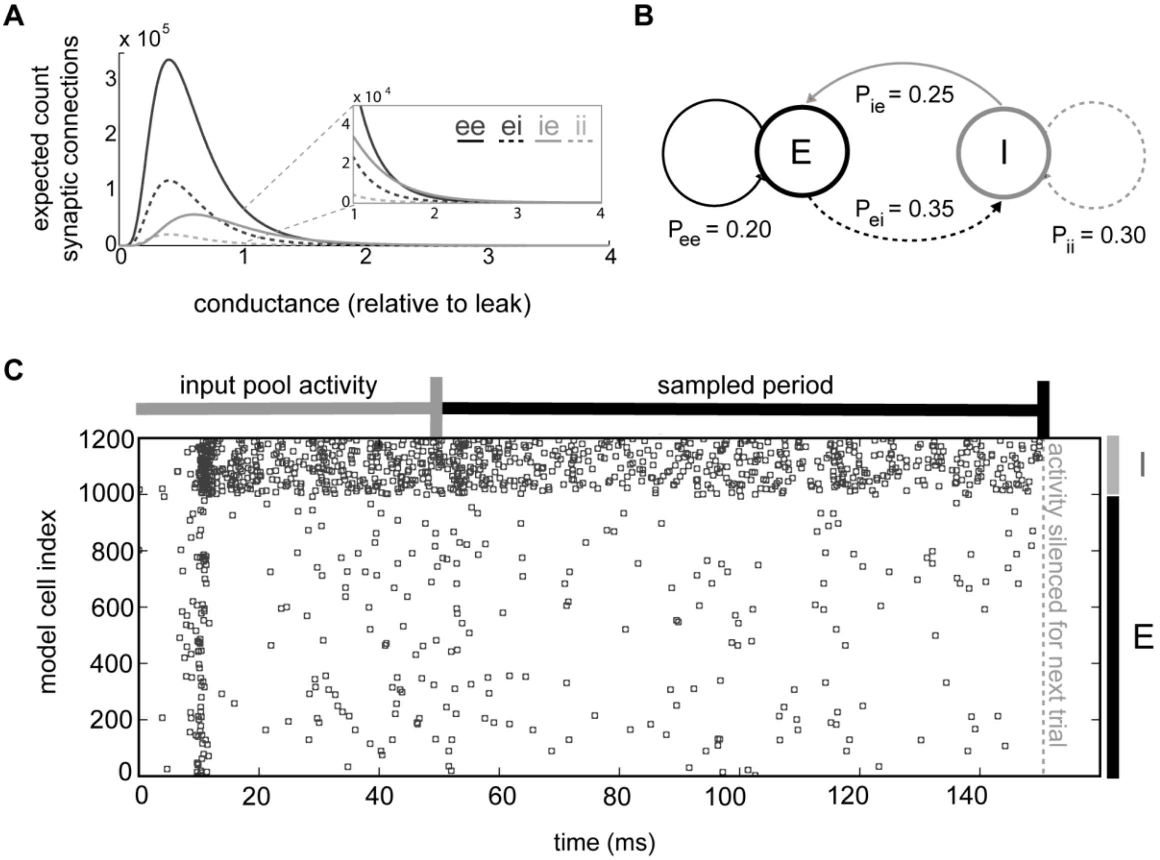

Figure 2-1. Conductance-based integrate and fire network with cell-type dependent connectivity.

(A) Synaptic weights were drawn from a lognormal distribution, with i-e conductances scaled 150% relative to other connections. Excitatory connections to other excitatory cells were most numerous and were isolated as the focus of this study.
(B) Two populations of integrate-and-fire units with conductance-based synapses, excitatory (black) and inhibitory (grey), were connected sparsely and randomly. Connection probabilities were determined by source and target identity (e.g. individual inhibitory projections to excitatory cells existed with probability pie = 0.25).
(C) A single simulation epoch of model network activity showing both excitatory and inhibitory spiking. The network was activated via input units firing Poisson spikes at 15 Hz, connected randomly to the excitatory population. After 50 ms, input activity was silenced and network activity progressed in isolation. During this period, activity was recorded from the excitatory population.

the depolarizing extent of synaptic bombardments changes with respect to inhibitory and excitatory reversal potentials[14]. Units in the model were divided into two groups—excitatory cells (*e*) and inhibitory cells (*i*). Synaptic weights were drawn from a lognormal distribution (Fig. 1A). Inhibitory connections, which constituted twenty percent of the total population, were scaled to be fifty percent stronger than excitatory connections in the mean. A small tonic excitatory drive g_t to all units helped stabilize sparse spiking. Excitatory and inhibitory conductances g_e and g_i were modeled as decaying exponentials triggered by presynaptic spiking:

$$\frac{dg_e}{dt} = -\frac{g_e}{\tau_i}$$

$$\frac{dg_i}{dt} = -\frac{g_i}{\tau_i}$$

Membrane potential dynamics were defined as follows, where v is voltage across the membrane:

$$\frac{dv}{dt} = \frac{g_e(E_e - v) + g_i(E_i - v) + g_t(E_t - v) + g_{leak}(E_{leak} - v)}{\tau_m}$$

Excitatory reversal potential E_e was 0 mV, as was E_t. Inhibitory reversal potential E_i was -90 mV. Reversal potential for leak current E_{leak} was -65 mV. Firing threshold was -48 mV, and post-spike reset was -70 mV. In addition to after spike hyperpolarization induced by the reset

---

[14] Cavallari, Panzeri, and Mazzoni, "Comparison of the Dynamics of Neural Interactions between Current-Based and Conductance-Based Integrate-and-Fire Recurrent Networks."

potential, a 1 ms absolute refractory period was imposed on model neurons. Leak conductance g_{leak} was fixed at 0.20 mS. Tonic depolarizing conductance g_t was equal in magnitude to the leak conductance.

Collective spiking generated spike-driven conductances that dwarfed the leak conductance, in keeping with definitions of high-conductance state[15]. Membrane time constant τ_m was 20 ms; excitatory synaptic time constant τ_e was 10 ms; and inhibitory synaptic time constant τ_i was 5 ms.

Excitatory connectivity was topologically random with $p_{ee} = 0.2$, in correspondence with measurements of dense local excitatory connectivity from L4 in the mouse C2 barrel[16]. Other connection probabilities were chosen to produce overall sparse excitatory spiking and fast promiscuous inhibitory spiking. The probability of an *i* unit impinging on an *e* unit P_{ie} was 0.25; p_{ii} was 0.3; p_{ei} was 0.35 (Fig. 1B).

Simulation protocol

The model network was stimulated using a pool of fifty input units. Each made uniformly weighted connections onto any excitatory cell with probability 1/10 (magnitude 3/5 leak conductance each). To begin a simulation trial, Poisson units spiked over these input connections at 15 Hz (50 ms duration). Network activity was not recorded during the input period. In the recording epoch, input units were silenced, allowing the network to fire according to its internal dynamics (100 ms duration) (Fig. 1C). After 100 ms, activity was silenced and a new trial was initiated.

---

[15] Destexhe, Rudolph, and Paré, "The High-Conductance State of Neocortical Neurons in Vivo."

[16] Lefort et al., "The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex."

The set of connections from input pool to excitatory neurons defined an input topology. After one hundred simulated recordings the input topology was randomly regenerated, stimulating the network using a new random pool of network neurons. After ten input patterns the simulation was terminated, for a total recording time of 100 s.

Finally, replicating experimental constraints, we occluded the majority of the neuronal network and down-sampled the remaining spiking activity. Except where otherwise specified, forty percent of excitatory cells were visible and used for inference.

Cross-validation of activity differences in relation to input context

Activity vectors of length N were constructed from segments of recorded activity such that entry N_i was the total number of spikes by neuron i. Distances between activity vectors were computed using the L_1 norm and normalized to mean total firing between the two vectors, quantifying the extent their activity differed as a fraction of total activity. This procedure was repeated across five 100 s simulation trials, each divided into 10 s epochs with shared input projections. Each epoch was subdivided into non-overlapping 5 s segments. Thus, there were 50 same-input comparisons total. Between-input comparisons were matched in number to same-input comparisons by random selection without replacement.

Reducing the number of possible connection schemes

In a network with N nodes where connections can take on m states, the number of possible connection schemes (*network topologies*) is m^N, a collection too large to consider exhaustively. To simplify the problem, we assume directed pairs fall into two

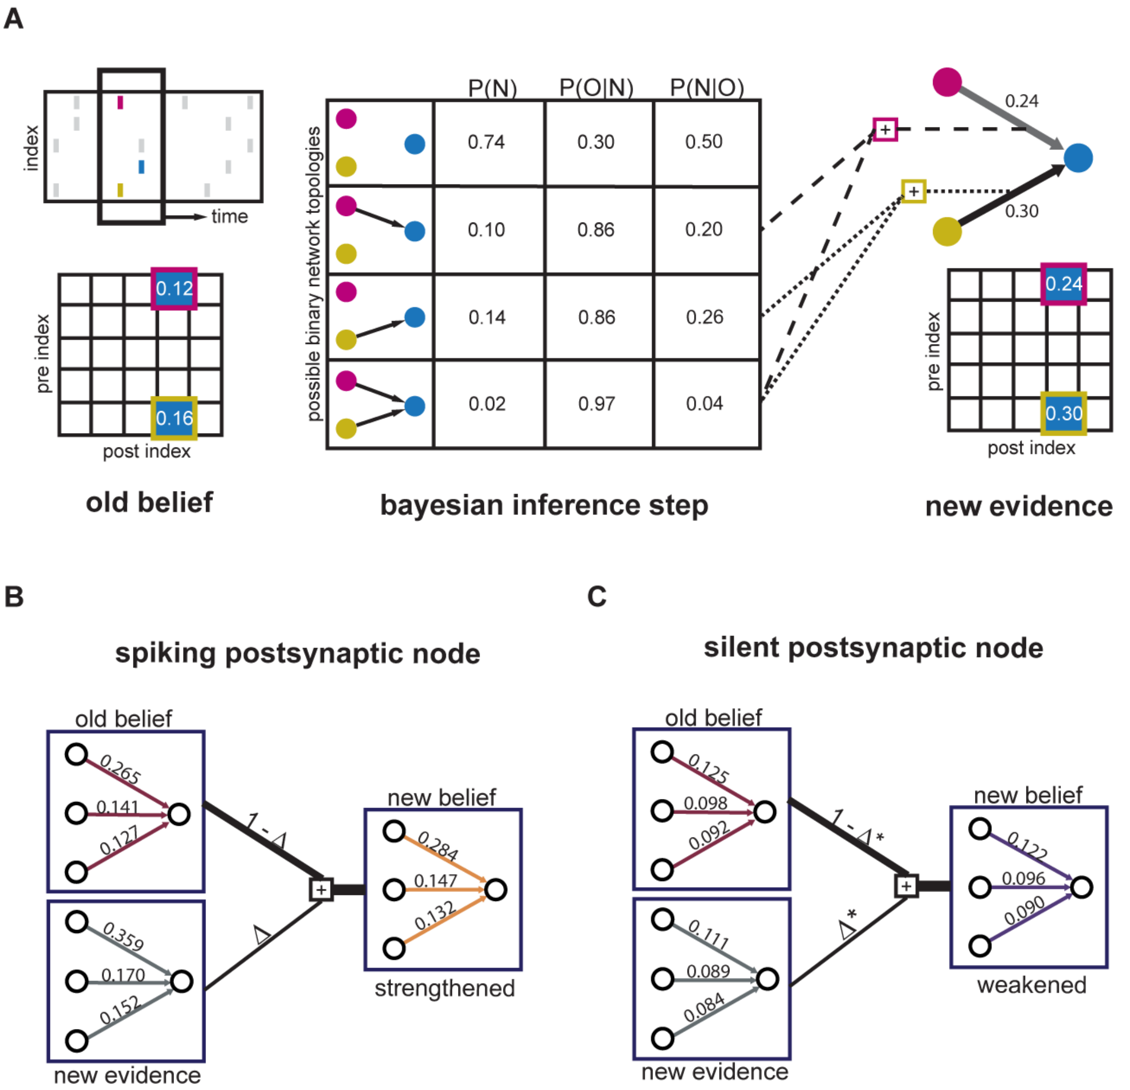

Figure 2-2. Iterative Bayesian inference for heuristically mapping reliable lagged firing.

(A) Inference step for a single observation. Top left. Pick one putative postsynaptic cell and frame (time bin). Consider the active neurons from the previous frame, each of which might have helped recruit the postsynaptic neuron to fire. Bottom left. Look up beliefs about connectivity between putative pre- and postsynaptic cells. Middle. Consider every binary feedforward connection topology on this subset of neurons. Compute their respective posterior probabilities using a likelihood model and Bayes rule. Right. Sum over the posterior distribution to obtain new evidence for the existence of individual connections.
(B) Integrate new evidence with old belief as a weighted sum. In this example, the postsynaptic reference neuron was active, and beliefs about possible presynaptic connectivity increase.
(C) Suppose the postsynaptic reference neuron was quiescent, resulting in decrementing beliefs about connectivity from putative presynaptics. Since quiescence is more common than activity in these recordings, $\Delta^* < \Delta$. An asymmetric learning rate prevents decrements from rapidly outstripping increments.

categories, connected or not connected (m=2). Additionally, during each Bayesian update step (Fig. 2A), only recently active nodes are considered. The number of active nodes in a particular frame (N_{active}) is small in practice because of sparse spiking and fast framerates. Depending on computational resources, a cutoff can be imposed on observations where N_{active} is large; we discarded observations where $N_{\text{active}} > 12$. With these simplifications a tractable number of possibilities are considered at each update step, despite lacking a polynomial-time algorithm for computing the distribution of posterior probabilities over potential topologies[17].

Inference of synaptic connectivity

The matrix W will be a repository for the developing beliefs about synaptic connectivity. At a given iteration, heuristic evidence for the existence of a connection e from potential presynaptic neuron (k) to reference neuron ($post$) can be found at entry $W_e := W_{k,post}$. Before the first observation, entries in W are initialized to a small uniform value. In this work we selected 0.1 for initial meanfield belief.

At observation j, choose a random frame t and a random reference neuron $post$. Let X_{post}^{t} describe whether the postsynaptic reference is active or quiescent in that frame. Call the set of neurons firing during the previous frame Pre, potential presynaptic partners to $post$. Let an observation (O_j) be the conjunction of reference postsynaptic and recently active presynaptic neurons:

---

[17] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."

$$O_j := (X_{post}^{t}, Pre^{t-1})$$

$$X_{post}^{t} := \begin{cases} 1 & if \quad \text{neuron } post \text{ active during frame } t \\ 0 & o.w. \end{cases}$$

$$\mathrm{Pre}^{t-1} := \{k \mid \text{ neuron } k \text{ active during frame } t-1\}$$

Consider a vector of ones and zeros representing the connections from <u>*Pre*</u> to *post* (N_i) describing the *i*th possible binary connection scheme between *Pre* and *post*. In this simplified framework where synapses can only take two states, zeros represent absent connections and ones represent existing connections.

$$N_i \in \{0,1\}^{N_{active}^{t}}$$

Our objective will be obtaining posterior probability $P(N_i|O_j)$ over all possible binary topologies N_i connecting *Pre* to *post*. To that end we define expressions for prior and likelihood using *W*. Assuming independence among connections for tractability, we compute the prior probability of a given possible topology N_i as

$$P(N_i) := \prod_{e \in N_i} W_e^{\,j-1} \prod_{e \notin N_i} (1 - W_e^{\,j-1})$$

For the likelihood expression, we define a simple forward model linking binary topology to postsynaptic activity. We imagine an idealized synaptic connection as transmitting a presynaptic spike to postsynaptic partners with probability α. We have used $\alpha = 0.8$ in this work. In cases of a postsynaptic spike, we model the forward interaction simply as

$$P(O_j^{active} \mid N_i) := 1 - \prod_{e \in N_i} (1 - \alpha)$$

Similarly, in the complementary case of postsynaptic quiescence:

$$P(O_j^{quiescent} \mid N_i) := \prod_{e \in N_i} (1 - \alpha)$$

It is important that α be less than 1, so that presynaptic collaboration is recognized in the likelihood computation. This simple model clearly underestimates the potential for collective and nonlinear presynaptic effects. Its corresponding virtue is that it requires very few parameters be fit to data.

From prior and likelihood, the posterior probability $P(N_i|O_j)$ is obtained using Bayes rule:

$$P(N_i \mid O_j) = \frac{P(O_j \mid N_i)P(N_i)}{P(O_j)} = \frac{P(O_j \mid N_i)P(N_i)}{\sum_l P(O_j \mid N_l)P(N_l)}$$

The marginal $P(O_j)$ is computed as a sum over all binary connection schemes. New evidence w'_e for individual connection e is obtained by summing across the posterior distribution:

$$w'_e = \sum_{i|(e \in N_i)} P(N_i \mid O_j)$$

This evidence is integrated into W as a weighted sum (Fig. 2B). Only the subset of edges considered in the Bayesian step are updated (the edges in N_i). The result of linear combination over the course of multiple observations is exponential smoothing on accumulated evidence.

$$W_e^j = \Delta w'_e + (1-\Delta) W_e^{j-1}$$

Because firing patterns are sparse, to prevent runaway decrementing towards zero in W, the learning rate Δ takes different values for postsynaptic activity versus quiescence (Fig. 2C). In this work we used $\Delta_{\text{active}} = 0.2$ and $\Delta_{\text{quiet}} = 0.05$. Because of the asymmetric learning rate, postsynaptic quiescence is a weak source of evidence compared to postsynaptic activity—in correspondence with biological intuitions. Over the course of many repetitions of the inference step, we drew pairs (*t, post*) without replacement until each visible neuron was sampled for all $t > 1$.

Evaluating performance

Precision was defined as the fraction of true positives to inferred connections (TP / TP + FP). Sensitivity was defined as the fraction of true positives to total connections (TP / TP + FN). The first measure quantifies detection accuracy, while the second quantifies extent of coverage. There is a tradeoff between these two factors, varying with choice of threshold on the belief matrix. High thresholds yield high precision but low sensitivity; low thresholds yield low precision but high sensitivity.

An alternative way to quantify the precision-sensitivity tradeoff is to use receiver operating characteristic (ROC) curves, but ROC curves are not a natural metric in this specific context of synaptic detection[18]. Specifically, since complete coverage of true positives is out of reach given current experimental methodologies, a large portion of the curve tends to fall on the main diagonal. As a result, measures of area under the curve systematically devalue inferred topologies that do not identify synapses over the entire range of functional weights, even if they perform optimally over a restricted range of thresholds. Area under the ROC curve also fails to differentiate between false-positive and false-negative errors, which have very different consequences for interpreting inferred connectivity[19]. Precision and sensitivity lessen this quantification problem and are additionally appealing because they are intuitive measures of detection performance.

We computed chance-level performance by counting the total number of actual synaptic connections in a category of interest (e.g. falling below a cutoff weight) normalized by the number possible (the set of all directed pairs in the sample population).

Preprocessing for correlation-type activity maps

For correlation-based inference, simulated spiking was first convolved with a Gaussian function using standard deviation equal to 10 ms.

Error motifs

Detection errors were analyzed for systematic features. To quantify the prevalence of errors arising from common inputs, we isolated false positives and counted the frequency with which they shared some hidden third input. We compared that error distribution to the frequency with which randomly selected pairs shared hidden inputs.

---

[18] Lobo, Jiménez-Valverde, and Real, "AUC."

[19] Ibid.

Accurate identification of connection direction is important for understanding large-scale features of activity flow (although it is less important for paired patch clamp experiments where pair identity matters more than direction of interaction, which can be recovered from EPSPs instead). We quantified misdirection errors by isolating spuriously inferred pre- and postsynaptic cells and counting the frequency of genuine connection from reference post- to presynaptic. We compared the resulting distribution to randomly chosen directed pairs.

Finally, we hypothesized that slow framerates might lead to spurious connections spanning multiple synapses. We isolated just the strongest ten percent of synaptic connections and computed the frequency that misidentified pre- and postsynaptic cells were connected by a chain (path) of strong connections no more than 4 cells long (3 edges or fewer). These limits served to constrain the otherwise dense connectivity, where otherwise any two neurons would be connected by a short arbitrary path.

RESULTS

Conductance-based network model with heterogeneous synaptic weights

A network model was developed to measure the performance, i.e. benchmark activity maps to excitatory synaptic connections (see Methods). For the model we drew synaptic conductances from a lognormal distribution in order to ensure that connections were weak on average, with a small number of strong connections (Fig. 1A), matching experimental reports[20]. Heterogeneity in synaptic efficacy has important consequences for information transmission[21].

---

[20] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits"; Perin, Berger, and Markram, "A Synaptic Organizing Principle for Cortical Neuronal Groups."

[21] Teramae, Tsubo, and Fukai, "Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links."

Although local neocortical connectivity is highly non-random[22], we used random connectivity in our model (Fig. 1B). A random network design is appropriate for benchmarking because it is more challenging to reconstruct from activity than more realistic clustered topologies[23], avoiding biases that inflate apparent performance.

Network activity was initiated from a small Poisson spiking external input pool. Each input unit projected randomly and independently into the excitatory population. By periodically randomizing the small number of input projections, we sought to achieve a diversity of activity within each network model, similar to spontaneous neocortical data[24] and activity in visual cortex driven by viewing natural scences[25]. Except where noted, model network activity was only recorded for inference after cessation of activity in the input population (Fig. 1C). At the end of the recording period we reset the model and began a new trial.

Simulated spike dynamics were asynchronous and irregular[26]. Asynchrony was measured with spike-rate correlations, by convolving spike times with a Guassian kernel of width $\sigma = 3$ ms. Among excitatory neurons in the recording period, mean correlation coefficient was 0.0019[27]. Irregularity was measured with interspike-intervals, which were observed to have mean squared-coefficient of variation of 0.81, consistent with other reports of irregular activity[28]. Excitatory

---

[22] Perin, Berger, and Markram, "A Synaptic Organizing Principle for Cortical Neuronal Groups"; Watts and Thomson, "Excitatory and Inhibitory Connections Show Selectivity in the Neocortex"; Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."

[23] Kobayashi and Kitano, "Impact of Network Topology on Inference of Synaptic Connectivity from Multi-Neuronal Spike Data Simulated by a Large-Scale Cortical Network Model."

[24] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry"; Sadovsky and MacLean, "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity"; Luczak, Barthó, and Harris, "Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations."

[25] Miller et al., "Visual Stimuli Recruit Intrinsically Generated Cortical Ensembles."

[26] Brunel, "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons."

[27] Kumar et al., "The High-Conductance State of Cortical Networks."

[28] Kumar, Rotter, and Aertsen, "Spiking Activity Propagation in Neuronal Networks."

spiking activity was also characterized by a branching parameter of 0.99 (for 10 ms bins), indicating near-critical dynamics[29]. Firing rates in the excitatory population during the recording period were 1.33 +/- 3.15 Hz (mean +/- std) consistent with findings in awake behaving mice[30]. According to these criteria, simulated spike dynamics were similar to *in vivo* recordings.

Iterative Bayesian inference to map synaptic recruitment

We compared activity in cortical circuits to the structure of their underlying connectivity. Functional weights summarize the frequency of recurring firing patterns, producing a map of circuit activity. In particular, we employed an iterative Bayesian inference method (modified from [31]) which requires few parameters and relatively few epochs of population activity as a result of its simple dynamical model (see Methods). In iterative Bayesian inference, information about firing patterns is accumulated by considering small portions of the entire recording, stitched together over the course of many observations. Bayes Rule is used to parse accumulating beliefs about conditional relationships across multiple iterations (Fig. 2, see methods). A single frame lag-window was used throughout this work when computing functional relationships.

The algorithm is outlined as follows: in a single observation, choose a random frame *t* and choose a (putative) postsynaptic neuron *post*. Consider the set of neurons **Pre** firing during frame *t-1*. These cells are candidates for having driven *post* to fire. Thus, an observation contains evidence about these potential presynaptic partners. Broadly, if *post* is spiking, there is new evidence for connectivity from **Pre** to *post*. Conversely, if *post* is silent, the observation is evidence

---

[29] Beggs and Plenz, “Neuronal Avalanches in Neocortical Circuits”; Haldeman and Beggs, “Critical Branching Captures Activity in Living Neural Networks and Maximizes the Number of Metastable States.”
[30] Crochet et al., “Synaptic Mechanisms Underlying Sparse Coding of Active Touch.”
[31] Pajevic and Plenz, “Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches.”

for a lack of connectivity from **Pre** to *post*. Silence is only weakly informative of lack of connectivity, because synaptic connections are weak in isolation. New evidence is integrated into the prior after each iteration step, gradually updating an overall summary of functional relationships.

Defining recruitment

Far from recruiting neurons nonspecifically, cortical activity flow at the level of the microcircuit is sparse and precise[32]. Isolating activity within a single input context (trials with the same input projection topology), 36±0.98% of model neurons were never active; but with the inclusion of many input contexts over the entire simulation, only 14.5% of model neurons were inactive, recapitulating experimental data[33]. Naturally, unresponsive neurons, which do not participate in propagating circuit activity, do not give rise to functional relationships.

Throughout, *lagged firing* describes any pair of cells active within one frame of each other, regardless of their connectedness. We defined *active synaptic connections* as those connecting pairs where the presynaptic neuron was active at least once. Surprisingly, many active synaptic connections never resulted in action potential generation in the postsynaptic model neuron. As a result, despite the fact that a synaptic connection was active and produced a subthreshold depolarization, such pre post pairs were not a route for propagating activity (which manifests as a multineuronal activity pattern). Synaptic connections which resulted in an action potential in both the pre- and postsynaptic model neurons were defined as *recruiting synaptic connections* (Fig.

---

[32] Kruskal, Li, and MacLean, "Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex"; Dombeck, Graziano, and Tank, "Functional Clustering of Neurons in Motor Cortex Determined by Cellular Resolution Imaging in Awake Behaving Mice."

[33] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry."

3A). Specifically, if a presynaptic partner was active in the lag window preceding postsynaptic firing at least once, the connection was defined as contributing to recruitment. Of course, these functional categorizations depend on the specific multineuronal activity pattern (see below) and choice of the lag window that was considered by the algorithm (see below).

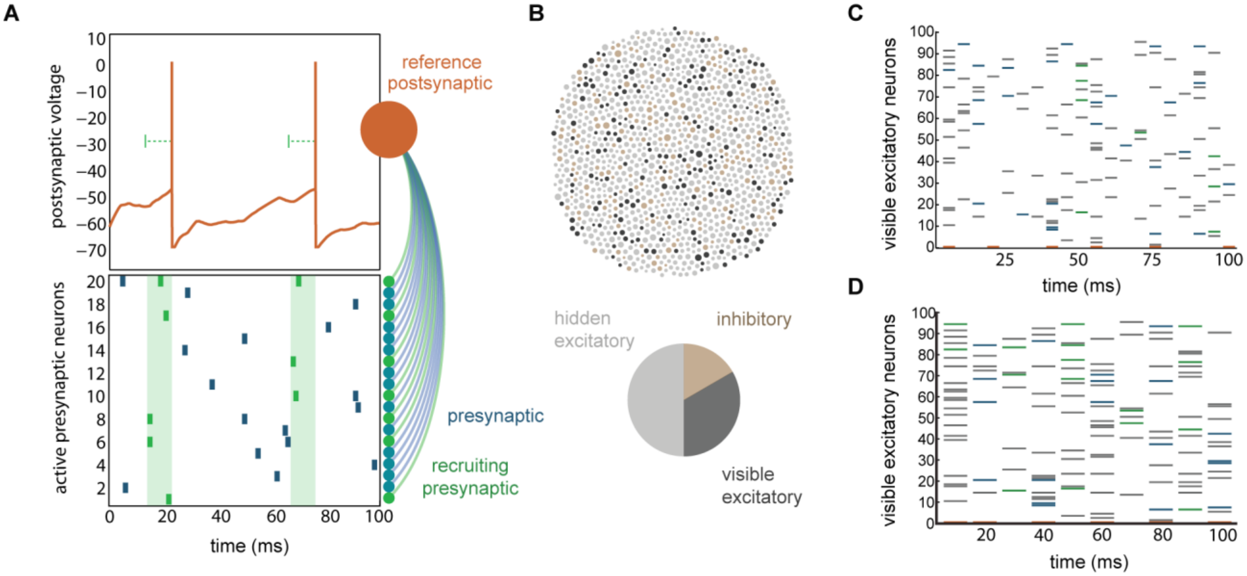

Figure 2-3. Quantifying recruitment and respecting experimental constraints.

(A) Relative to a postsynaptic reference, recruiting presynaptic neurons are those active at least once during an integration window preceding postsynaptic spiking.
(B) Excitatory cells were sampled randomly for recording, with the majority of neurons hidden (60% of excitatory cells hidden), introducing the challenge of hidden statistical structure. Node diameters proportional to their weighted in-degrees (sum absolute value of afferent weights).
(C) Activity subsampled into 5 ms bins and categorized using a single-frame integration window, Recruiting presynaptic cells (green) and other active presynaptic cells (blue) were annotated relative to a postsynaptic reference cell (orange, bottom row). The majority of lagged activity involves cells that are not presynaptic (gray) to the reference.
(D) Activity was also subsampled into 10 ms and 25 ms bins.

Simulating experimental conditions

Even with dense uniform sampling within a large field of view[34], connected neurons reside above, below, and outside the imaged field. Unobserved activity can lead to ambiguity in the reconstruction of synaptic connectivity from spiking because of misleading statistical dependencies[35], particularly in resolving true synaptic connections versus unconnected neurons with shared inputs[36]. To benchmark the iterative Bayesian algorithm while acknowledging experimental constraints, we occluded 60% of excitatory cells within the network model (except where otherwise noted) (Fig. 3B). Further, we down-sampled modeled network activity in time, since even the fastest optics and indicator dyes are slower than an action potential (Fig. 3C). The consequences of estimating spike times from Ca^{2+} signals have been investigated explicitly in a previous study[37].

Poisson null populations verify significance of timing relationships

Whenever many neurons are simultaneously active, spurious apparent structure will exist because of chance coincident spike timing. To confirm the presence of genuine pairwise timing in model activity, we created populations of rate-matched Poisson units. A Poisson population can establish chance frequency of activity patterns because it lacks any internal causal interactions[38]. Each Poisson unit was paired to a corresponding neuron in the model. Poisson rates were individually determined from the number of spikes their model counterparts fired during the

---

[34] Sadovsky et al., “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.”

[35] Chicharro and Panzeri, “Algorithms of Causal Inference for the Analysis of Effective Connectivity among Brain Regions.”

[36] Nykamp, “A Mathematical Framework for Inferring Connectivity in Probabilistic Neuronal Networks”; Nykamp, “Pinpointing Connectivity despite Hidden Nodes within Stimulus-Driven Networks.”

[37] Lütcke et al., “Inference of Neuronal Network Spike Dynamics and Topology from Calcium Imaging Data.”

[38] Roxin, Hakim, and Brunel, “The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons.”

sampling period, on a trial-by-trial basis. Thus, the null population recapitulated firing rate structure and chance interactions without mimicking precise interrelationships in timing.

We applied Iterative Bayesian inference to the connected network activity and the independent Poisson activity sampled over 100 s with 5 ms lag windows. Inferred connectivity matrices from matched Poisson activity were impoverished in total weight (model: 0.17±0.10; Poisson: 0.058±0.074; mean ± standard deviation of individual edge weights). These differences were significant ($p = 1.1 \times 10^{-7}$, random subset of 100 non-zero edges, *Wilcoxon rank-sum*) (Fig. 4A). Strong weights reflect reliable spiking relationships between neurons, and they were present in the randomly connected model but not in the Poisson matched null population. These results support the hypothesis that causal connectivity gives rise to reliable functional relationships beyond what is expected by chance. Functional relationships computed with iterative Bayesian inference for an *ex vivo* dataset collected from mouse somatosensory cortex echoed the simulated-network weight distribution, with a small number of particularly reliable relationships forming an extensive tail (Fig. 4B).

Using timing relationships to predict synaptic connectivity

Applying a threshold to the matrix of functional relationships involved a tradeoff between false positives (Type 1 errors) and false negatives (Type II errors) (FPs and FNs respectively). To dissociate these error types, we assessed performance using two metrics: precision (TP / TP + FP) and sensitivity (TP / TP + FN). First, we considered hypothetical conditions where the excitatory population was fully sampled at 100 Hz. As threshold increased, predictions of connectivity became more precise but less sensitive, capturing a smaller fraction of total connections. This tradeoff became particularly visible after isolating synaptic pairs that fired in

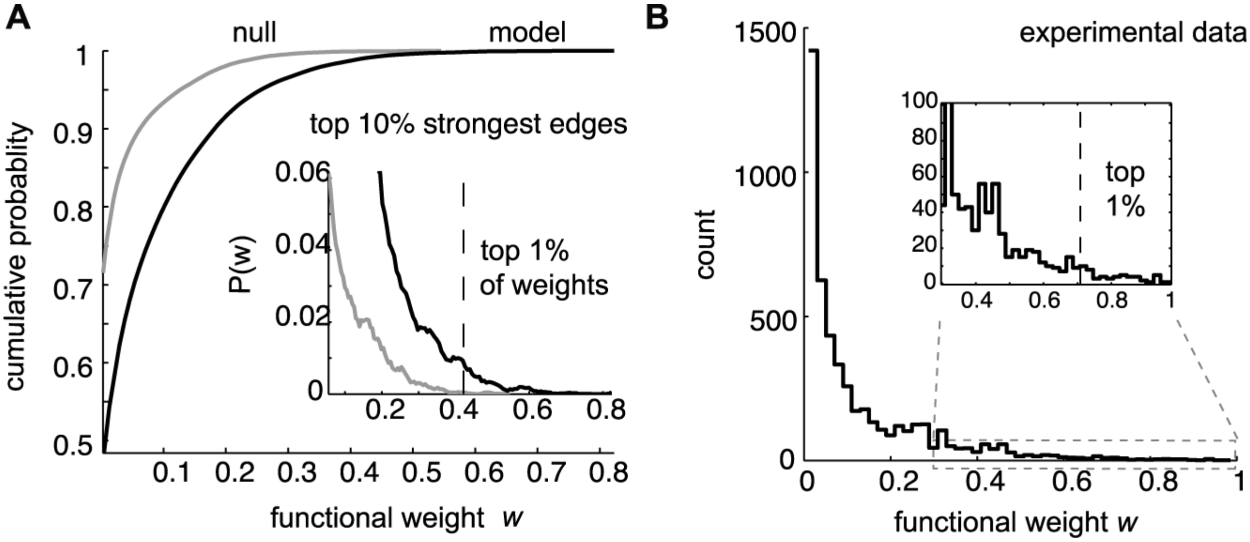

Figure 2-4. Cortical activity includes reliable pairwise timing relationships.

(A) Cumulative distribution of inferred connection strengths for connected model (black) versus Poisson rate matched (grey) populations. Firing rates for each Poisson units were matched on a trial-by-trial basis to each model neuron. Inset. The corresponding probability distribution function, in zoom. Top percentile of non-zero functional weights (dashed).
(B) Comparison to imaging data. Inferred functional weights for activity recorded in mouse somatosensory neocortex.

succession: where presynaptic recruited postsynaptic partner to threshold at least once (*recruiting connections*). Applying a top 1%-level threshold informed by Poisson null comparisons, the bulk of chance-level timing relationships were excluded, prioritizing true positives and avoiding false positives.

Correlation is a simple alternative for computing functional relationships, and it has been shown to be informative about synaptic connectivity[39]. In fact, when the goal is to maximize the number of detected connections in a low-precision regime, correlation outperforms iterative Bayesian inference. For intermediate scenarios, correlation and iterative Bayesian inference perform comparably. However, when the goal is to identify true positives with minimal ambiguity,

---

[39] Cossell et al., "Functional Organization of Excitatory Synaptic Strength in Primary Visual Cortex."

iterative Bayesian inference substantially outperforms correlation (in a high-precision regime) (Figure 5A). Because precision comes at the cost of limited sensitivity, these approaches cannot replace anatomical methods for revealing entire connectomes. Instead, they have the potential to find small active subnetworks of synaptic connections. Inferred iterative Bayesian weights are highly informative of connection probability (Fig. 5B). We find a linear relationship between threshold and precision up to a saturation at perfect performance ($r^2 = 0.98$). This relationship between functional weights and the likelihood of a synaptic connection held true regardless of inference method (correlation alone, $r^2 = 0.89$).

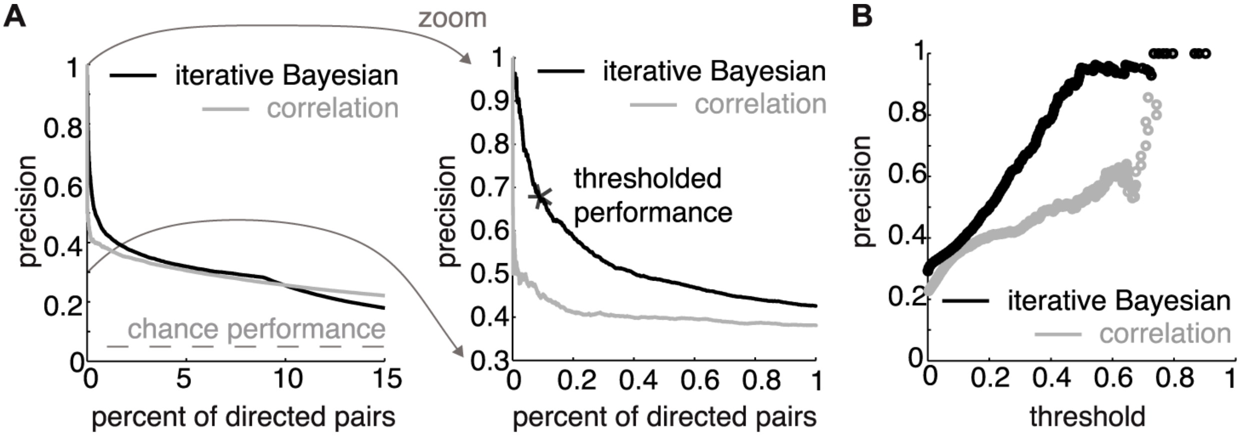

Figure 2-5. Relationship of inferred weighs to monosynaptic connectivity.

(A) Left. Sensitivity (TP / TP + FN) versus precision (TP / TP + FP) for two activity maps: iterative Bayesian inference (black) and correlation (grey). Right. Iterative Bayesian inference substantially outperforms correlation in a narrow high-precision regime. Thresholded performance (asterisk). Chance performance reflects the likelihood of identifying a recruiting synapse from a directed pair selected at random. Recruiting connections are a smaller target than synaptic connections at large. (B) Threshold on functional weight versus resulting precision, for iterative Bayesian inference (black) versus correlation (grey). Functional weights reflect probability of connectedness.

Influence of external input on inference

For evaluating inferred connectivity, we focused on recordings of sustained recurrent activity, after cessation of external Poisson input. To complement those findings, we also quantified performance for simulated activity during the input period, in the presence of ongoing external drive. For this comparison, 10 s samples of activity were recorded under conditions where three hundred excitatory neurons were visible. In the absence of external drive, precision reached 62 +/- 3.5% (mean +/- std) and sensitivity reached 0.47 +/- 0.025%. During external drive under the same conditions, precision reached 60 +/- 8.3% and sensitivity reached 0.37 +/- 0.065%. Between the two conditions, precision did not differ significantly (Wilcoxon rank-sum, p=0.84, n=5) but sensitivity was higher in the absence of external inputs (Wilcoxon rank-sum, p=0.032, n=5).

Diversity of activity between different input contexts

Primary sensory cortices exhibit diverse multineuronal activity patterns[40]. Spatially diverse input projections were used to recapitulate these experimental findings. Distinct input projections activated overlapping populations of model neurons (Fig. 6A). However, common neurons were recruited by different presynaptic neighbors and recruited different postsynaptic neighbors in turn, consistent with experimental data[41]. Thus, pairwise relationships differed even where single neuron activity was similar. Single neurons were commonly active after many input contexts, but single recruiting connections were more often unique (Fig. 6B). Cross-

---

[40] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry"; Sadovsky and MacLean, "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity."

[41] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry."

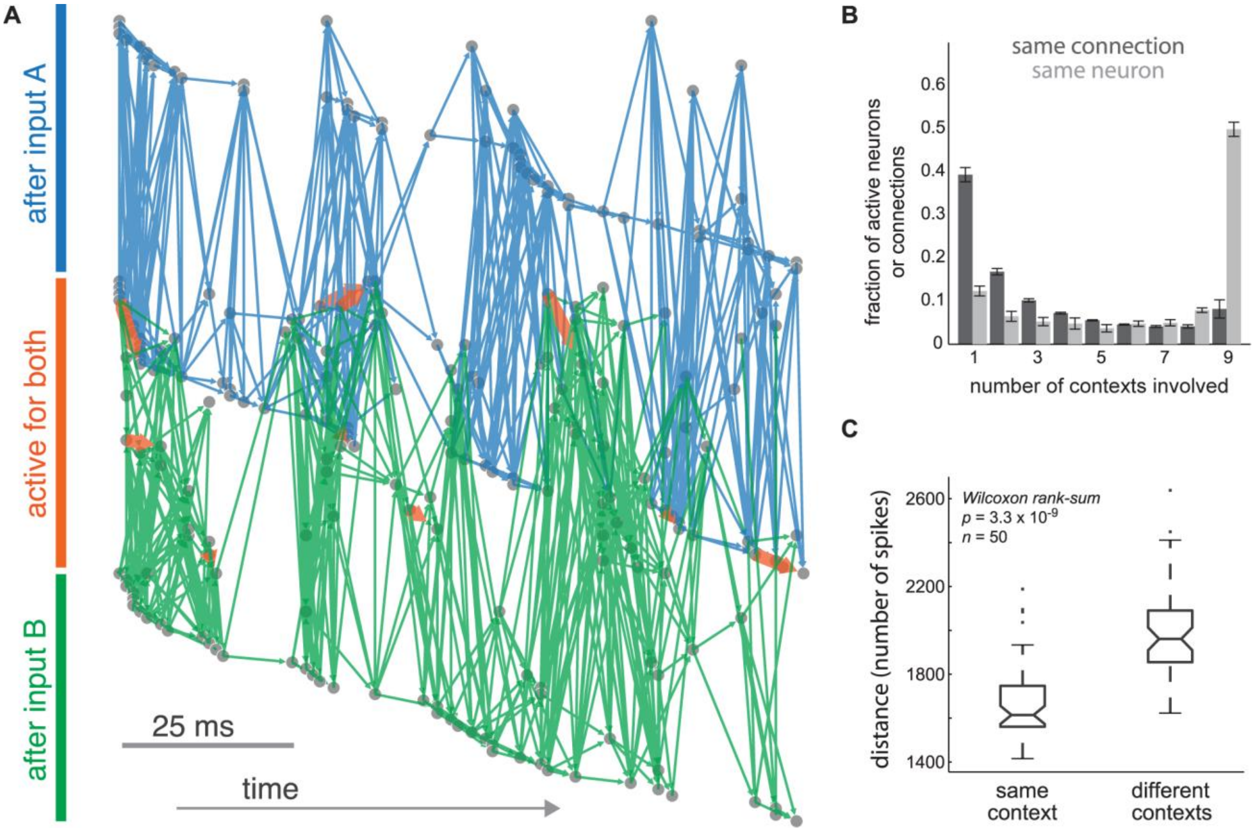

Figure 2-6. Activity flow in different input contexts.

(A) Patterns of excitatory recruitment in single trials after input A (blue connections) and after input B (green connections). Rarely, pairwise recruitment reccurred across both trials (orange connections). Hidden excitatory and inhibitory neurons not shown. Rows indicate cell identity. Although pairwise recruitment was largely unique between trials, many single neurons were active in both trials (middle rows, orange bar).

(B) Single neurons are commonly active at least once in every input context. In contrast, pairwise recruitment relationships tended to occur in just one or a few input contexts.

(C) Cross-validation verifies that activity across different input contexts genuinely differs, quantified with L1 distance (city-block distance) for activity vectors within the same input context versus between different input contexts.

validation confirmed that 5 s periods of activity differed substantially more between than within input contexts, in terms of active frames per neuron (Fig. 6C) ($p = 3.25 \times 10^{-9}$, $n = 50$, *Wilcoxon rank-sum*).

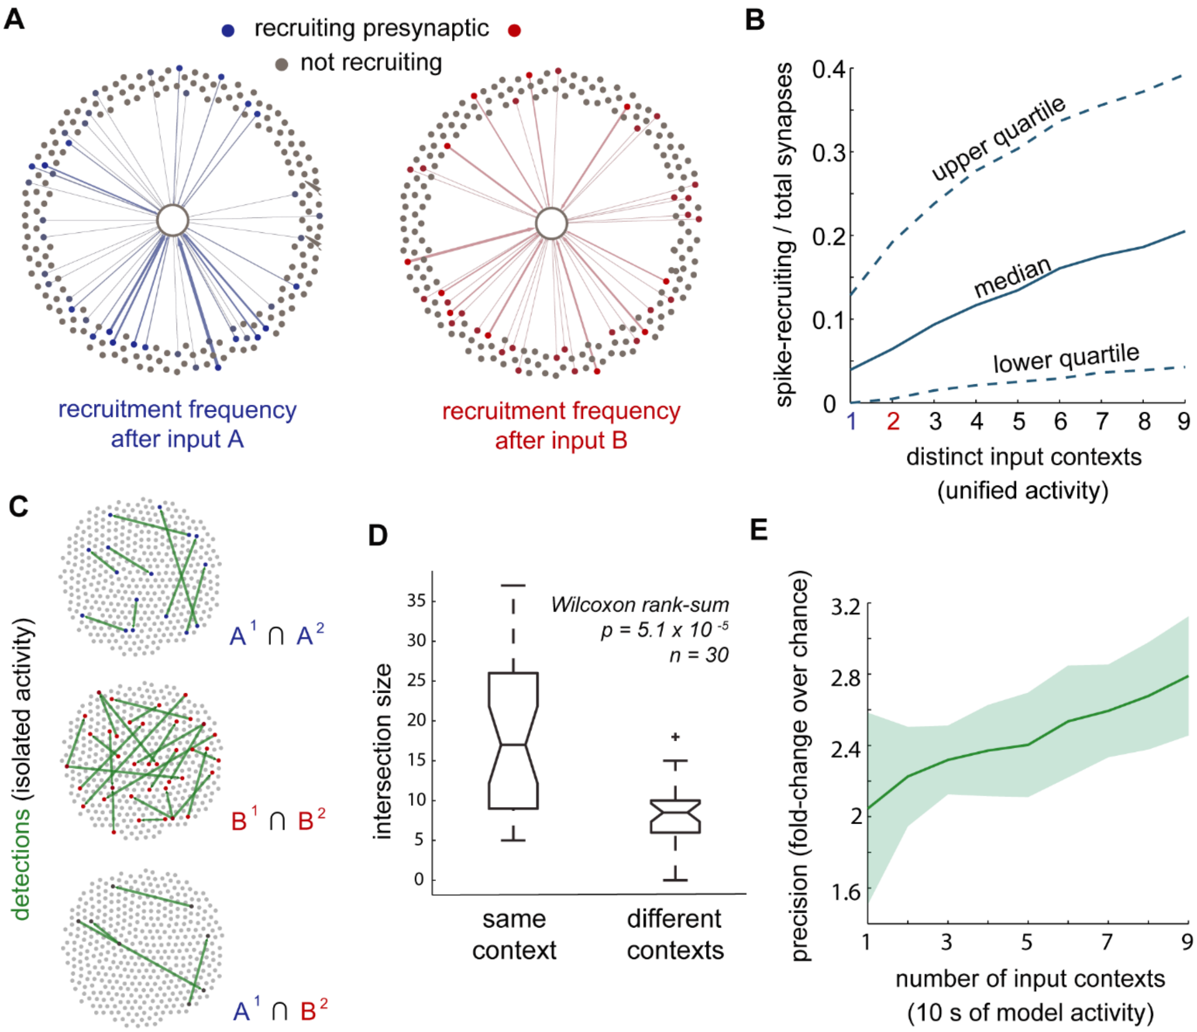

Figure 2-7. Mapping diverse activity.

(A) Presynaptic recruitment frequency for a reference postsynaptic neuron after distinct inputs. Within a single context, the majority of presynaptic neurons never recruit the postsynaptic reference. Edge weight proportional to frequency of recruitment.

(B) Input diversity and recording time increases the fraction of presynaptic neurons involved in postsynaptic recruitment.

(C) A single comparison in the second cross-validation procedure: overlap in detected synapses within shared input contexts (top, middle) and between different input contexts (bottom). Green lines indicate true positives.

(D) Pooling over comparisons, cross-validation verified that different-context overlap in true positives is small, whereas same-context overlap is large. By identifying regularities in recruitment, iterative Bayesian activity maps were linked to the structure of their inputs.

(E) Simulated data was subsampled to establish the role of diversity, for inference while controlling for sample duration. In each case the data analyzed totaled 10 s of recording time. The number of contexts was the divisor. For inference in the recording period, after inputs were silenced, precision (TP / TP+FP) increases with increasing numbers of diverse inputs.

Distinct patterns of recruitment detected after different inputs

On the pairwise level, recruitment patterns varied from one input pattern to another (Fig. 7A). As a consequence, with increasing numbers of distinct input projections and recording time, increasing numbers of recruitment relationships manifested (Fig. 7B). After a single input context, 8.0±0.21% of connections were involved in recruitment. Across all input contexts, the number of recruiting connections climbed to 24%. Diverse activity exposed additional patterns of activity flow.

To verify that distinct inputs gave rise to different patterns of detection, we performed a second cross-validation procedure. Spiking activity was initiated from two distinct populations of input units, *input context A* and *input context B*. For each context, simulated activity was divided into two halves, each of which was analyzed separately using iterative Bayesian inference, yielding four functional topologies. Identities of detected synapses were compared within and across contexts (Fig. 7C). Despite variability between individual inference trials, similar detections recurred within same-context activity. Functional relationships were shared between dissimilar contexts less often. Detected synapses were significantly more similar for activity generated from a common input context than for activity between input contexts ($p = 5.1 \times 10^{-5}$, $n = 30$, *Wilcoxon rank-sum*; Fig. 7D). During the period of sustained activity after cessation of inputs, population activity differed depending on input history. These differentiated patterns of spiking gave rise to distinct functional relationships.

To control for the effect of sample duration on inference, we fixed the total amount of time considered by the algorithm to 10 s of activity regardless of the number of different contexts (Fig. 7E). When only one context was analyzed the single context was chosen randomly for each

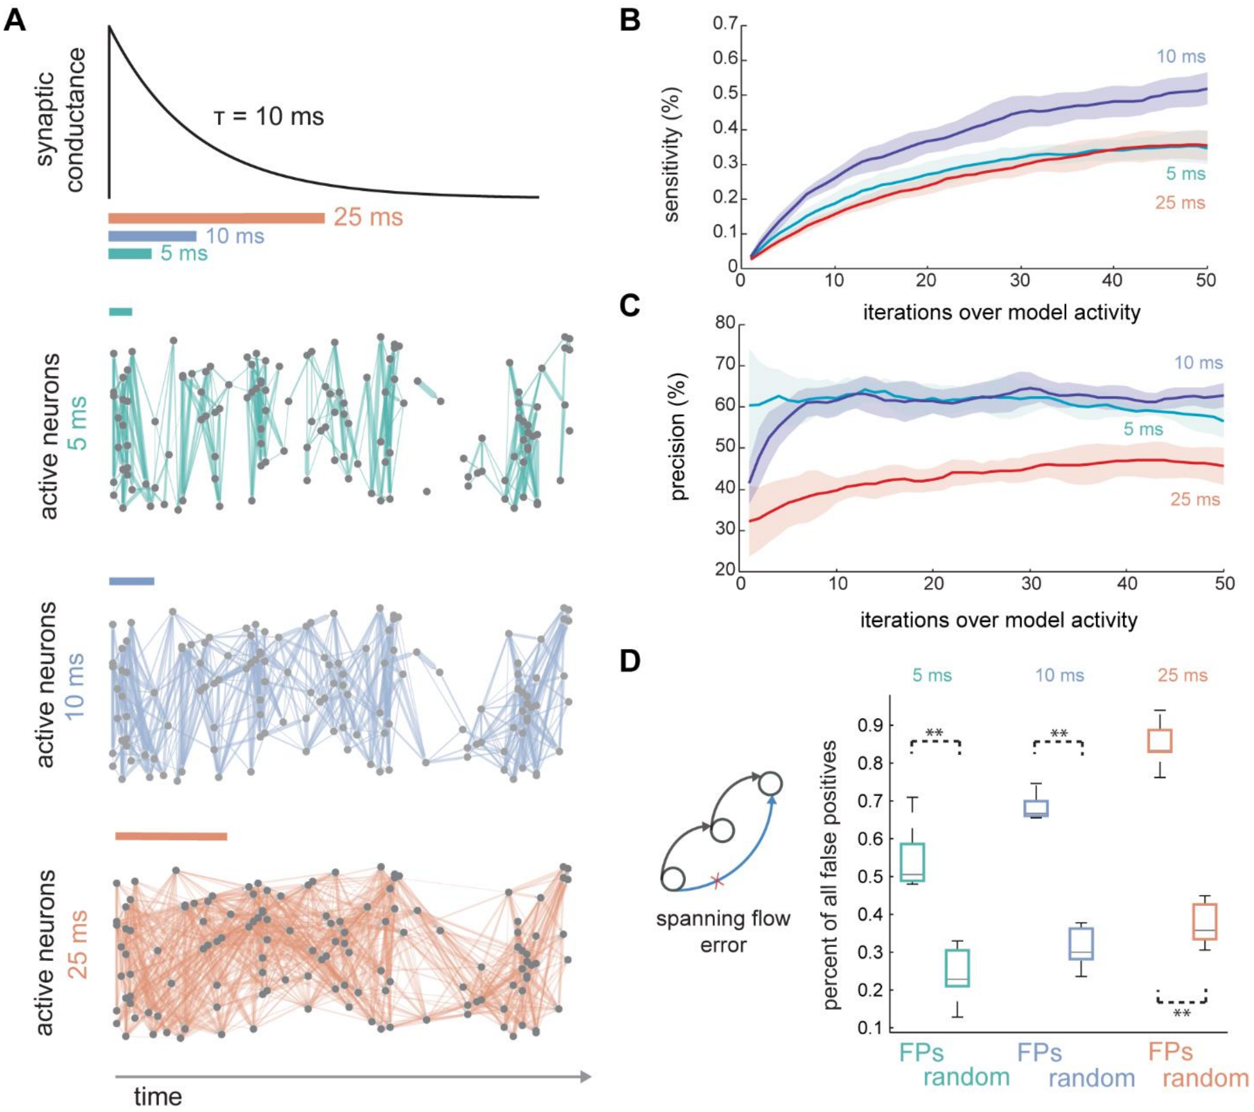

Figure 2-8. Detection depends on definition of recruitment window.

(A) Excitatory conductances in the model compared to three definitions of recruitment window. Synaptic interactions appear sparse for short interaction windows and dense for large lag windows.
(B) Sensitivity improved fastest and most extensively when synaptic interactions were mapped with 10 ms recruitment windows (5 ms, blue; 10 ms, purple; 25 ms, red).
(C) Precision improved fastest for 5 ms recruitment windows, while 10 ms windowing resulted in slower increases and a modestly higher plateau (5ms, blue; 10ms, purple; 25ms, red).
(D) Conditioned on false positive detection, the probability of spanning a chain of strong synapses became increasingly likely as recruitment window increased. Spanning errors accounted for the majority of false positives for all recruitment window definitions. Asterisks indicate $p < 0.001$.

repetition. As we increased the number of contexts considered by the algorithm, while always keeping 10s of data fixed, we randomly choose the input contexts to be analyzed. After controlling the amount of data analyzed by the algorithm (10s) inference of underlying connections was substantially better when diverse activity patterns were analyzed. In practice this meant, for example, that 2.5s of data from 4 contexts (2.5s * 4 contexts) was more effective than 10s of data from one context for the inference of connections.

Enhanced performance at the timecourse of excitatory synapses

Generally, faster framerates can be obtained at the cost of smaller population sample sizes. This provides flexibility to select temporal resolution best suited to the detection of causal connections. We considered three potential recruitment windows: 25 ms, 10 ms, and 5 ms (Fig. 8A). Using HOPS scanning *ex vivo* with single-frame recruitment windows, these frame durations correspond to approximately 300, 75 and 30 neurons scanned, respectively[42]. To investigate how choice of single-frame lag window impacts inference without the confound of changing sample size, we obscured 60% of the 1000 excitatory model cells in each trial. We found that both sensitivity and precision were maximal when a 10 ms recruitment window was used in the inference algorithm. At 10 ms, precision was 63±5.6% (median ± interquartile range) and sensitivity to recruiting synapses was 1.9±0.20%, which was significantly better than 5 ms (precision 57±1.6%: $p = 0.011$; sensitivity 1.6±0.24%: $p = 0.026$; *Wilcoxon rank-sum*, n = 7) and 25 ms (precision 45±0.19%: $p = 5.8 \times 10^{-4}$; sensitivity 1.1±0.20%: $p = 0.0023$; *Wilcoxon rank-sum*, n = 7). Although 25 ms resolution still permitted recovery of monosynaptic connectivity well

---

[42] Sadovsky et al., "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging."

above chance levels, performance at 5 ms was significantly better than for 25 ms lags (precision: $p = 5.8 \times 10^{-4}$; sensitivity: $p = 0.0041$; n = 7, *Wilcoxon rank-sum*).

For all three sampling rates, best performance was obtained by repeatedly presenting the 100 s of simulated recordings to the inference algorithm. In all cases, sensitivity improved gradually, steadily encompassing increasing numbers of monosynaptic connections (Fig. 8B). Reiteration led to particularly extensive improvement in precision at the slowest framerates, with 25 ms lag window analyses improving an additional five percentage points in precision after 5 ms and 10 ms lag window inferences had stabilized (Fig. 8C). In contrast, those faster framerates stabilized early in precision after ten to twenty iterations, with modest subsequent decline for 5 ms lag windows in conjunction with more extensive sensitivity.

Error sources and interaction with choice of recruitment window

Next, we characterized likely sources of error and their interaction with temporal resolution. For each connectivity pattern, random pairs of model neurons were used to establish chance occurrences, independent of errors in predicting synaptic connectivity. If candidate motifs occurred more frequently at sites of false positives than between random pairs, they were associated with errors in inference.

Shared input between two neurons can induce spurious correlations that appear to indicate causal connectivity when none is present[43]. Surprisingly, when synaptic connections were falsely inferred, we did not find elevated numbers of shared hidden inputs. At least under these conditions using iterative Bayesian inference, common hidden inputs did not make significant contribution to

---

[43] Gerstein, Perkel, and Subramanian, "Identification of Functionally Related Neural Assemblies"; Keshri et al., "A Shotgun Sampling Solution for the Common Input Problem in Neural Connectivity Inference."

false positive detection errors, even at slow framerates (5 ms: $p = 0.976$, $n = 567$; 10 ms: $p = 0.989$, $n = 653$; 25 ms: $p = 0.776$, $n = 910$; FPs versus random, *two-way Kolmogorov-Smirnov*).

An issue of concern is whether slow framerates lead to errors resolving connection direction. We evaluated whether the direction of a true synaptic connection was frequently opposite that of an inferred synaptic connection. At baseline, randomly selected directed pairs were connected twenty percent of the time (as can be predicted from their underlying connection probability). The probability of directionality errors was no different from that of finding a connection between two randomly selected neurons. At the slowest framerates, we did observe a possible weak association between direction errors and sampling rate (5 ms: $p = 0.883$, $n = 7$ trials; 10 ms: $p = 0.129$, $n = 7$; 25 ms: $p = 0.129$, $n = 7$; FPs versus random, *two-way Kolmogorov-Smirnov*). Nevertheless, these results suggest that slow sampling need not confound measurement of flow direction.

The majority of errors falsely attributed monosynaptic connectivity where there was actually strong polysynaptic connectivity (5 ms: $p = 5.8 \times 10^{-4}$, n = 7, *Wilcoxon rank-sum;* 10 ms: $p = 5.8 \times 10^{-4}$, n = 7, *Wilcoxon rank-sum;* 25 ms: $p = 5.8 \times 10^{-4}$, n = 7, *Wilcoxon rank-sum*; Fig. 8D). As sampling rate decreased, an increasing fraction of errors arose from functional relationships spanning polysynaptic chains. Note that random cells were more likely to be connected in chains at slower framerates simply because longer recruitment windows (temporal scale) defined more presynaptic cells as participating in recruitment, increasing the density of inferred networks.

Strong synapses are prominent in functional relationships

Synaptic weights were drawn from a lognormal distribution, matching experimental findings[44]. Strong connections from the long tail of the weight distribution occurred only infrequently. Thus, at chance-level they would be expected to participate in functional relationships only rarely.

In actuality, the effects of strong synapses were particularly detectable from patterned network activity. The distribution of detected synapses was shifted toward higher connection strength values compared to the underlying distribution used to assign synaptic conductances (Fig. 9A). The probability of identifying a recruiting connection was positively correlated with

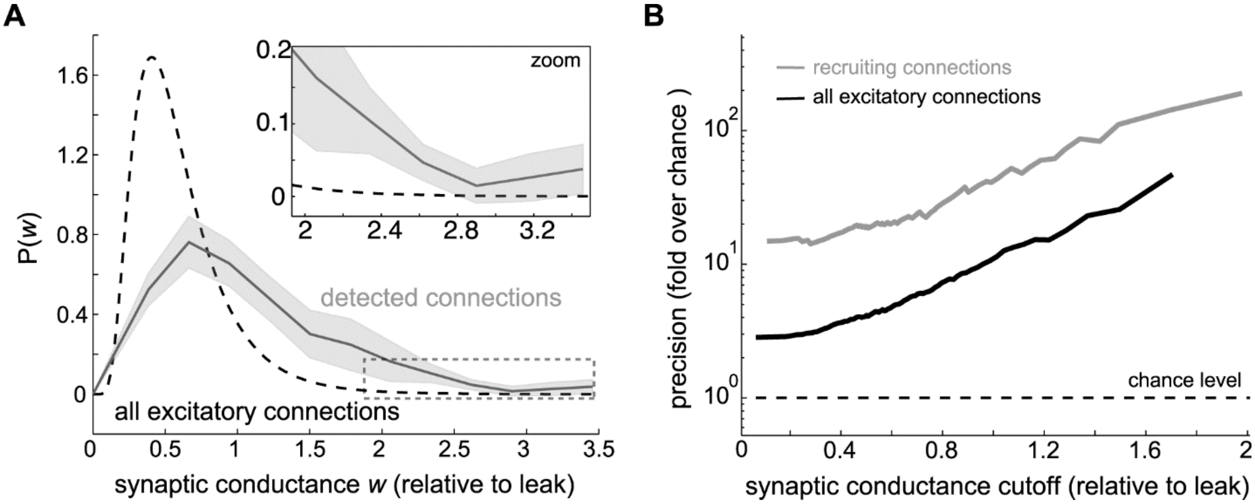

Figure 2-9. Strong connections are prominent in recruitment.

(A) Distribution of underlying excitatory synaptic conductances compared to synaptic conductances of detected connections.
(B) Precision of inferred connections increases with synaptic strength. Defining chance performance based on all excitatory synaptic connections (black). Defining chance performance more realistically, based on the subset of excitatory connections involved in recruitment (grey).

---

[44] Perin, Berger, and Markram, "A Synaptic Organizing Principle for Cortical Neuronal Groups"; Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."

connection strength (r = 0.68, p = 0.0001; Pearson correlation). Furthermore, as stronger subpopulations of structural connections were isolated, precision relative to chance rose dramatically (Fig. 9B).

DISCUSSION

The relationship between connectivity and patterned sequential firing is complex. An average synaptic connection causes only small postsynaptic depolarization, particularly when membrane conductance is high during synaptic bombardments[45]. As a result, it is difficult to predict sequential firing from synaptic weights alone. To circumvent this complexity, it is necessary to map synaptic activity in emergent contexts generated by ongoing activity.

However, misleading coincidences in spike timing have the potential to confound inference of causal synaptic connectivity[46]. We used Poisson populations, lacking interconnectivity, to establish how frequently lagged firing which would be interpreted by our inference approach to be indicative of a connection appear by chance. We found that the null, which matched rate statistics, did not reproduce the same lagged firing relationships found in network activity. Connected populations gave rise to strong and reliable pairwise timing relationships where Poisson firing did not. By applying a threshold to inferred weights, we were able to exclude many spurious relationships.

---

[45] Teramae, Tsubo, and Fukai, “Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links”; Chicharro and Panzeri, “Algorithms of Causal Inference for the Analysis of Effective Connectivity among Brain Regions”; Kumar et al., “The High-Conductance State of Cortical Networks”; Destexhe, Rudolph, and Paré, “The High-Conductance State of Neocortical Neurons in Vivo”; Troyer and Miller, “Physiological Gain Leads to High ISI Variability in a Simple Model of a Cortical Regular Spiking Cell.”

[46] Roxin, Hakim, and Brunel, “The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons.”

The modified iterative Bayesian inference algorithm applied here results in sparse connectivity matrices that relate synaptic connectivity to a small number of particularly precise timing relationships. These inferred connections map routes of activity through the synaptic network. By isolating only the most precise pairwise timing relationships, iterative Bayesian inference identifies a small number of synaptic connections with higher precision than can be achieved using correlation-based inference, a common alternative approach (e,g. [47]). For applications requiring high sensitivity but only modest precision, correlation or other inference approaches may be more appropriate; but for accurate inference of a synaptic connection with very few Type I errors, iterative Bayes should be the algorithm of choice.

A constellation of methods for inferring connectivity from imaged activity is emerging, reflecting a growing consensus that bridging function and structure is a crucial long-term goal in neuroscience (e.g. [48]). This study complements related work by emphasizing performance under the conditions of realistic optical experiments—constrained recording durations and restricted sampling[49]. We purposefully designed our simulations to pose substantial challenges to inference, including weak conductance-based synapses, random connectivity, occlusion of 60% of the neurons in the network, and limited recording times. We included these challenges in order to mimic experimental constraints.

For network simulations, input units made random connections onto the recorded population. To mimic diverse inputs, their random projections were periodically redrawn. Thus,

---

[47] Cossell et al., "Functional Organization of Excitatory Synaptic Strength in Primary Visual Cortex."

[48] Mishchencko, Vogelstein, and Paninski, "A Bayesian Approach for Inferring Neuronal Connectivity from Calcium Fluorescent Imaging Data"; Stetter et al., "Model-Free Reconstruction of Excitatory Neuronal Connectivity from Calcium Imaging Signals"; Gerhard et al., "Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone."

[49] Sadovsky et al., "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging."

between different contexts, inputs were matched in average temporal structure but differed in their vectors of projection weights. The result was a diversity of model network spike patterns, echoing our experimental observations[50]. Differences in recruitment manifested even though inputs were completely silent during the recording period, when network activity evolved in isolation. These results demonstrate a simple way that recent history impacts network activity, with context dependent recruitment situated to play an important role in neuronal computation[51].

Because each input context gave rise to new patterns of activity, detection of synaptic connections improved after diverse inputs. The importance of diverse activity suggests that mapping be carried out using statistical rich sensory stimuli such as natural movies. Spontaneous activity is also a good source of diverse activity, having been demonstrated to broadly traverse network activity patterns[52]. Independent driving stimulations are likely to generate even higher sensitivity[53], at the cost of potentially non-naturalistic patterns of emergent activity in the cortical network. We found that inference was somewhat more sensitive during the recording period, when external inputs were silent, than in the presence of external Poisson spiking. This difference may arise from the addition of common hidden inputs.

Interaction times between neurons are not dominated by fast synaptic delays, so that millisecond resolution is not necessary for resolving monosynaptic connectivity. We found it better to use a 10 ms lag window than a 5 ms lag window, measuring functional relationships which closely mirrored the timecourse of excitatory synaptic conductances in our model. Using HOPS

---

[50] Sadovsky and MacLean, “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry”; Sadovsky and MacLean, “Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity.”

[51] Buonomano and Maass, “State-Dependent Computations”; Kumar, Rotter, and Aertsen, “Spiking Activity Propagation in Neuronal Networks.”

[52] Luczak et al., “Sequential Structure of Neocortical Spontaneous Activity in Vivo”; MacLean et al., “Internal Dynamics Determine the Cortical Response to Thalamic Stimulation.”

[53] Van Bussel, Kriener, and Timme, “Inferring Synaptic Connectivity from Spatio-Temporal Spike Patterns.”

scanning[54] a 10 ms frame duration corresponds to approximately 75 neurons imaged. Note that equal or better precision could be obtained by considering lags spanning multiple 5 ms frames, rather than a single 10 ms frame. But in practice, the doubled scan-rate would necessitate more than halving the imaged population size. At slower framerates, reliable timing relationships were still predictive of monosynaptic connectivity. However, they also increasingly corresponded to misidentification of a single connection where in fact there was a chain of strong connections—a spanning error. Regardless of temporal resolution, polysynaptic spanning errors may help explain why random networks are harder to reconstruct than clustered architectures[55]. In clustered networks, strong chains are more likely to be crisscrossed with additional connections, resulting in neighbors-of-neighbors arrangements[56]. These results suggest designing experiments with fast population scan rates if the aim is to infer causal monosynaptic connectivity from activity. This consideration must be balanced with the need to sufficiently sample the network numerically, in order to see enough spiking relationships to detect statistical dependencies[57]. The membrane time constant is also relevant to the interval within which integration can occur, and slow metabotropic conductances may extend integration time even farther in real cortical neurons.

Strong cortical connections are a desirable experimental target because of their proposed roles in information processing[58]. However, they are rare in the population of all connections.

---

[54] Sadovsky et al., “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.”

[55] Kobayashi and Kitano, “Impact of Network Topology on Inference of Synaptic Connectivity from Multi-Neuronal Spike Data Simulated by a Large-Scale Cortical Network Model.”

[56] Perin, Berger, and Markram, “A Synaptic Organizing Principle for Cortical Neuronal Groups.”

[57] Gururangan, Sadovsky, and MacLean, “Analysis of Graph Invariants in Functional Neocortical Circuitry Reveals Generalized Features Common to Three Areas of Sensory Cortex.”

[58] Song et al., “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits”; Litwin-Kumar et al., “Balanced Synaptic Input Shapes the Correlation between Neural Spike Trains”; Teramae, Tsubo, and Fukai, “Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links”; Gururangan, Sadovsky, and MacLean, “Analysis of Graph Invariants in Functional Neocortical Circuitry Reveals Generalized Features Common to Three Areas of Sensory Cortex.”

Consistent with their importance, we have shown that in patterned spiking the strongest synapses are particularly salient and detectable: 10- to 100-fold more so than predicted by chance detection levels. These results suggest combining imaging and inference with paired-patch experiments to study the role of strong synapses in local circuit dynamics, enabling higher success rates in searching for strongly connected pairs.

Our findings highlight a feature of structure-function relationships that is discussed surprisingly little: in realistic activity, a synaptic connection need not imply the existence of sequential firing. That is, in a given recording, a synaptic connection may contribute to depolarizing its target without ever preceding a postsynaptic action potential. Similarly, a single connection may function as a route for propagating activity in one network context, but not in another[59]. These different functional roles have important consequences for activation of projection neurons. Relatedly, propagating activity establishes the connections that can be potentiated through Hebbian mechanisms, because of proximal pre- and postsynaptic spiking. Yet for arbitrary input and network history, the problem of predicting routes of propagation from the collective interaction of many synapses is exceedingly hard. Activity-mapping approaches sidestep this daunting problem by isolating specific synaptic connections underlying multineuronal firing, potentially linking specific synaptic connections to computation in the neocortex.

---

[59] Fisher et al., “A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit.”

REFERENCES

Barth, Alison L., and James FA Poulet. “Experimental Evidence for Sparse Firing in the Neocortex.” *Trends in Neurosciences* 35, no. 6 (2012): 345–355.

Beggs, John M., and Dietmar Plenz. “Neuronal Avalanches in Neocortical Circuits.” *The Journal of Neuroscience* 23, no. 35 (2003): 11167–11177.

Brette, Romain, Michelle Rudolph, Ted Carnevale, Michael Hines, David Beeman, James M. Bower, Markus Diesmann, et al. “Simulation of Networks of Spiking Neurons: A Review of Tools and Strategies.” *Journal of Computational Neuroscience* 23, no. 3 (2007): 349–398.

Brunel, Nicolas. “Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons.” *Journal of Computational Neuroscience* 8, no. 3 (2000): 183–208.

Buonomano, Dean V., and Wolfgang Maass. “State-Dependent Computations: Spatiotemporal Processing in Cortical Networks.” *Nature Reviews Neuroscience* 10, no. 2 (February 2009): 113–25. doi:10.1038/nrn2558.

Cavallari, Stefano, Stefano Panzeri, and Alberto Mazzoni. “Comparison of the Dynamics of Neural Interactions between Current-Based and Conductance-Based Integrate-and-Fire Recurrent Networks.” *Frontiers in Neural Circuits* 8 (2014): 12.

Chicharro, Daniel, and Stefano Panzeri. “Algorithms of Causal Inference for the Analysis of Effective Connectivity among Brain Regions.” *Information-Based Methods for Neuroimaging: Analyzing Structure, Function and Dynamics*, 2015. https://books.google.com/books?hl=en&lr=&id=go5kCgAAQBAJ&oi=fnd&pg=PA148&dq=Algorithms+of+causal+inference+for+the+analysis+of+effective+connectivity+among+brain+regions&ots=mQt7U6DIAi&sig=ppFk-qTIULrRrmBHkRppvtOFaHM.

Cossell, Lee, Maria Florencia Iacaruso, Dylan R. Muir, Rachael Houlton, Elie N. Sader, Ho Ko, Sonja B. Hofer, and Thomas D. Mrsic-Flogel. “Functional Organization of Excitatory Synaptic Strength in Primary Visual Cortex.” *Nature* 518, no. 7539 (2015): 399–403.

Crochet, Sylvain, James FA Poulet, Yves Kremer, and Carl CH Petersen. “Synaptic Mechanisms Underlying Sparse Coding of Active Touch.” *Neuron* 69, no. 6 (2011): 1160–1175.

Destexhe, Alain, Michael Rudolph, and Denis Paré. “The High-Conductance State of Neocortical Neurons in Vivo.” *Nature Reviews Neuroscience* 4, no. 9 (2003): 739–751.

Dombeck, Daniel A., Michael S. Graziano, and David W. Tank. “Functional Clustering of Neurons in Motor Cortex Determined by Cellular Resolution Imaging in Awake Behaving Mice.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 29, no. 44 (November 4, 2009): 13751–60. doi:10.1523/JNEUROSCI.2985-09.2009.

Fisher, Dimitry, Itsaso Olasagasti, David W. Tank, Emre R. F. Aksay, and Mark S. Goldman. “A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit.” *Neuron* 79, no. 5 (September 4, 2013): 987–1000. doi:10.1016/j.neuron.2013.06.041.

Gerhard, Felipe, Tilman Kispersky, Gabrielle J. Gutierrez, Eve Marder, Mark Kramer, and Uri Eden. “Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone.” *PLOS Comput Biol* 9, no. 7 (July 11, 2013): e1003138. doi:10.1371/journal.pcbi.1003138.

Gerstein, G. L., and D. H. Perkel. “Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation.” *Science (New York, N.Y.)* 164, no. 3881 (May 16, 1969): 828–30.

Gerstein, George L., Donald H. Perkel, and K. N. Subramanian. “Identification of Functionally Related Neural Assemblies.” *Brain Research* 140, no. 1 (1978): 43–62.

Goodman, Dan FM, and Romain Brette. “The Brian Simulator.” *Frontiers in Neuroscience* 3, no. 2 (2009): 192.

Gururangan, Suchin S., Alexander J. Sadovsky, and Jason N. MacLean. “Analysis of Graph Invariants in Functional Neocortical Circuitry Reveals Generalized Features Common to Three Areas of Sensory Cortex.” *PLoS Comput Biol* 10, no. 7 (July 10, 2014): e1003710. doi:10.1371/journal.pcbi.1003710.

Haldeman, Clayton, and John M. Beggs. “Critical Branching Captures Activity in Living Neural Networks and Maximizes the Number of Metastable States.” *Physical Review Letters* 94, no. 5 (2005): 58101.

Harvey, Christopher D., Philip Coen, and David W. Tank. “Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task.” *Nature* 484, no. 7392 (April 5, 2012): 62–68. doi:10.1038/nature10918.

Keshri, Suraj, Eftychios Pnevmatikakis, Ari Pakman, Ben Shababo, and Liam Paninski. “A Shotgun Sampling Solution for the Common Input Problem in Neural Connectivity Inference.” *arXiv* 1309 (2013). https://pdfs.semanticscholar.org/3870/7e5043182fec9fc15ea2709e40d000590d2c.pdf.

Ko, Ho, Sonja B. Hofer, Bruno Pichler, Katherine A. Buchanan, P. Jesper Sjöström, and Thomas D. Mrsic-Flogel. “Functional Specificity of Local Synaptic Connections in Neocortical Networks.” *Nature* 473, no. 7345 (May 5, 2011): 87–91. doi:10.1038/nature09880.

Kobayashi, Ryota, and Katsunori Kitano. “Impact of Network Topology on Inference of Synaptic Connectivity from Multi-Neuronal Spike Data Simulated by a Large-Scale Cortical Network Model.” *Journal of Computational Neuroscience* 35, no. 1 (2013): 109–124.

Kruskal, Peter B., Lucy Li, and Jason N. MacLean. “Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex.” *Nature Communications* 4 (2013). http://www.nature.com/ncomms/2013/131010/ncomms3574/full/ncomms3574.html.

Kumar, Arvind, Stefan Rotter, and Ad Aertsen. “Spiking Activity Propagation in Neuronal Networks: Reconciling Different Perspectives on Neural Coding.” *Nature Reviews Neuroscience* 11, no. 9 (2010): 615–627.

Kumar, Arvind, Sven Schrader, Ad Aertsen, and Stefan Rotter. “The High-Conductance State of Cortical Networks.” *Neural Computation* 20, no. 1 (2008): 1–43.

Lefort, Sandrine, Christian Tomm, J.-C. Floyd Sarria, and Carl CH Petersen. “The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex.” *Neuron* 61, no. 2 (2009): 301–316.

Litwin-Kumar, Ashok, Anne-Marie M. Oswald, Nathaniel N. Urban, and Brent Doiron. “Balanced Synaptic Input Shapes the Correlation between Neural Spike Trains.” *PLOS Comput Biol* 7, no. 12 (December 22, 2011): e1002305. doi:10.1371/journal.pcbi.1002305.

Lobo, Jorge M., Alberto Jiménez-Valverde, and Raimundo Real. “AUC: A Misleading Measure of the Performance of Predictive Distribution Models.” *Global Ecology and Biogeography* 17, no. 2 (2008): 145–151.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. “Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations.” *Neuron* 62, no. 3 (2009): 413–425.

Luczak, Artur, Peter Barthó, Stephan L. Marguet, György Buzsáki, and Kenneth D. Harris. “Sequential Structure of Neocortical Spontaneous Activity in Vivo.” *Proceedings of the National Academy of Sciences* 104, no. 1 (January 2, 2007): 347–52. doi:10.1073/pnas.0605643104.

Lütcke, Henry, Felipe Gerhard, Friedemann Zenke, Wulfram Gerstner, and Fritjof Helmchen. “Inference of Neuronal Network Spike Dynamics and Topology from Calcium Imaging Data.” *Neural Circuits Revealed*, 2015. https://books.google.com/books?hl=en&lr=&id=Jg2qCgAAQBAJ&oi=fnd&pg=PA162&dq=.+Inference+of+neuronal+network+spike+dynamics+and+topology+from+calcium+imaging+data&ots=gh9EB-DNXJ&sig=o9tugMv38SUn76EESFOBNjsG9E8.

MacLean, Jason N., Brendon O. Watson, Gloster B. Aaron, and Rafael Yuste. “Internal Dynamics Determine the Cortical Response to Thalamic Stimulation.” *Neuron* 48, no. 5 (December 8, 2005): 811–23. doi:10.1016/j.neuron.2005.09.035.

Magee, Jeffrey C. “Dendritic Integration of Excitatory Synaptic Input.” *Nature Reviews Neuroscience* 1, no. 3 (2000): 181–190.

Miller, Jae-eun Kang, Inbal Ayzenshtat, Luis Carrillo-Reid, and Rafael Yuste. “Visual Stimuli Recruit Intrinsically Generated Cortical Ensembles.” *Proceedings of the National Academy of Sciences* 111, no. 38 (2014): E4053–E4061.

Mishchencko, Yuriy, Joshua T. Vogelstein, and Liam Paninski. “A Bayesian Approach for Inferring Neuronal Connectivity from Calcium Fluorescent Imaging Data.” *The Annals of Applied Statistics*, 2011, 1229–1261.

Neske, Garrett T., Saundra L. Patrick, and Barry W. Connors. “Contributions of Diverse Excitatory and Inhibitory Neurons to Recurrent Network Activity in Cerebral Cortex.” *The Journal of Neuroscience* 35, no. 3 (2015): 1089–1105.

Nykamp, Duane Q. “A Mathematical Framework for Inferring Connectivity in Probabilistic Neuronal Networks.” *Mathematical Biosciences* 205, no. 2 (2007): 204–251.

Nykamp, Duane Q. “Pinpointing Connectivity despite Hidden Nodes within Stimulus-Driven Networks.” *Physical Review E* 78, no. 2 (2008): 21902.

O’Connor, Daniel H., S. Andrew Hires, Zengcai V. Guo, Nuo Li, Jianing Yu, Qian-Quan Sun, Daniel Huber, and Karel Svoboda. “Neural Coding during Active Somatosensation Revealed Using Illusory Touch.” *Nature Neuroscience* 16, no. 7 (2013): 958–965.

Pajevic, Sinisa, and Dietmar Plenz. “Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches,” 2009. http://dx.plos.org/10.1371/journal.pcbi.1000271.

Perin, Rodrigo, Thomas K. Berger, and Henry Markram. “A Synaptic Organizing Principle for Cortical Neuronal Groups.” *Proceedings of the National Academy of Sciences* 108, no. 13 (2011): 5419–5424.

Roxin, Alex, Vincent Hakim, and Nicolas Brunel. “The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons.” *The Journal of Neuroscience* 28, no. 42 (2008): 10734–10745.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.” *Journal of Neurophysiology* 106, no. 3 (September 1, 2011): 1591–98. doi:10.1152/jn.00334.2011.

Sadovsky, Alexander J., and Jason N. MacLean. “Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity.” *The Journal of Neuroscience* 34, no. 23 (2014): 7769–7777.

Sadovksy, Alexander J., and Jason N. MacLean. “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry.” *The*

Journal of Neuroscience 33, no. 35 (August 28, 2013): 14048–60. doi:10.1523/JNEUROSCI.1977-13.2013.

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B. Chklovskii. “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits.” *PLoS Biol* 3, no. 3 (2005): e68.

Stetter, Olav, Demian Battaglia, Jordi Soriano, and Theo Geisel. “Model-Free Reconstruction of Excitatory Neuronal Connectivity from Calcium Imaging Signals.” *PLOS Comput Biol* 8, no. 8 (August 23, 2012): e1002653. doi:10.1371/journal.pcbi.1002653.

Teramae, Jun-nosuke, Yasuhiro Tsubo, and Tomoki Fukai. “Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links.” *Scientific Reports* 2 (2012). http://www.nature.com/srep/2012/120702/srep00485/full/srep00485.html.

Troyer, Todd W., and Kenneth D. Miller. “Physiological Gain Leads to High ISI Variability in a Simple Model of a Cortical Regular Spiking Cell.” *Neural Computation* 9, no. 5 (1997): 971–983.

Van Bussel, Frank, Birgit Kriener, and Marc Timme. “Inferring Synaptic Connectivity from Spatio-Temporal Spike Patterns.” *Frontiers in Computational Neuroscience* 5 (2011): 3.

Vogels, Tim P., and L. F. Abbott. “Signal Propagation and Logic Gating in Networks of Integrate-and-Fire Neurons.” *The Journal of Neuroscience* 25, no. 46 (November 16, 2005): 10786–95. doi:10.1523/JNEUROSCI.3508-05.2005.

Watson, Brendon O., Jason N. MacLean, and Rafael Yuste. “UP States Protect Ongoing Cortical Activity from Thalamic Inputs.” *PLoS One* 3, no. 12 (2008): e3971.

Watts, Jo, and Alex M. Thomson. “Excitatory and Inhibitory Connections Show Selectivity in the Neocortex.” *The Journal of Physiology* 562, no. 1 (2005): 89–97.

CHAPTER 3

Higher-order synaptic interactions coordinate dynamics in recurrent networks

This work is currently under peer review: Chambers B & MacLean JN (2016). Higher-order synaptic interactions coordinate dynamics in recurrent networks. *PLoS Computational Biology*, submitted March 30, 2016.

ABSTRACT

Linking synaptic connectivity to dynamics is key to understanding information processing in neocortex. Circuit dynamics emerge from complex interactions of interconnected neurons, necessitating that links between connectivity and dynamics be evaluated at the network level. Here we map propagating activity in large neuronal ensembles from mouse neocortex and compare it with a recurrent network model, where connectivity can be precisely measured and manipulated. We find that a dynamical feature dominates statistical descriptions of propagating activity for both neocortex and the model: convergent clusters comprised of fan-in triangle motifs, where two input neurons are themselves connected. Fan-in triangles coordinate the timing of presynaptic inputs during ongoing activity to effectively generate postsynaptic spiking. As a result, paradoxically, fan-in triangles dominate the statistics of spike propagation even in randomly connected recurrent networks. Interplay between higher-order synaptic connectivity and the integrative properties of neurons constrains the structure of network dynamics and shapes the routing of information in neocortex.

AUTHOR SUMMARY

Active networks of neurons exhibit beyond-pairwise dynamical features. In this work, we identify a canonical higher-order correlation in network dynamics and trace its emergence to synaptic integration. We find that temporally coordinated firing preferentially occurs at sites of fan-in triangles—a synaptic motif which coordinates presynaptic timing, leading to greater likelihood of postsynaptic spiking. The influence of fan-in clustering leads to the surprising emergence of non-random routing of spiking in random synaptic networks. When synaptic weights are made artificially stronger in simulation, so that cooperative input is less crucial, dynamics are no longer dominated by fan-in triangles but instead more closely reflect the random synaptic network. Thus, the emergence of fan-in clustering in maps of synaptic recruitment is a collective property of individually weak connections in neuronal networks. Because higher-order interactions are necessary to shape the timing of presynaptic inputs, activity does not propagate uniformly through the synaptic network. Like water finding the deepest channels as it flows downhill, spiking activity follows the path of least resistance and is routed through triplet motifs of connectivity. These results argue that clustered fan-in triangles are a canonical network motif and mechanism for spike routing in local neocortical circuitry.

INTRODUCTION

Understanding any complex system requires a mechanistic account of how dynamics arise from underlying architecture. Patterns of connections shape dynamics in diverse settings ranging from electric power grids to gene transcription networks[1]. It is critical to establish how synaptic connectivity orchestrates the dynamics of propagating activity in neocortical circuitry, since dynamics are closely tied to cortical computation. For example, trial-to-trial differences in network dynamics[2] can be used to decode sensory inputs and behavioral choice[3]. It is particularly important to understand the transformation from connectivity to activity within local populations of neurons since this is the scale at which the majority of connections arise. Locally, neocortical neurons are highly interconnected, and their connectivity schemes are characterized by the prevalence of specific motifs[4]. At the level of local populations, functional coordination has been demonstrated in diverse ways, *e.g.* on the basis of active neurons[5] and their correlation patterns[6]. Yet predicting population responses on the basis of pairwise connections alone has proven to be difficult.

Establishing a mechanistic link between connectivity and dynamics in neocortical networks is intricate and non-trivial because individual neurons themselves are complex computational

---

[1] Womelsdorf et al., "Dynamic Circuit Motifs Underlying Rhythmic Gain Control, Gating and Integration"; Motter et al., "Spontaneous Synchrony in Power-Grid Networks"; Rohden et al., "Self-Organized Synchronization in Decentralized Power Grids"; Shen-Orr et al., "Network Motifs in the Transcriptional Regulation Network of Escherichia Coli"; Barzel and Barabási, "Universality in Network Dynamics."

[2] Luczak, McNaughton, and Harris, "Packet-Based Communication in the Cortex"; Luczak, Barthó, and Harris, "Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations"; Sadtler et al., "Neural Constraints on Learning"; Sadovsky and MacLean, "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity."

[3] Harvey, Coen, and Tank, "Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task"; Sederberg, Palmer, and MacLean, "Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity."

[4] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."

[5] MacLean et al., "Internal Dynamics Determine the Cortical Response to Thalamic Stimulation"; Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry."

[6] Rothschild, Nelken, and Mizrahi, "Functional Organization and Population Dynamics in the Mouse Primary Auditory Cortex."

units[7]. Fundamentally, neurons are state dependent non-linear integrators of synaptic input[8]. When neurons in neocortex process information, they are generally subjected to numerous synaptic inputs which activate diverse receptors, and concomitant gating of voltage-dependent channels[9]. In consequence, neocortical neurons tend to operate in a high-conductance state, which lessens the impact of any one synaptic input[10]. Because inputs are weak individually, collective synaptic bombardments are necessary to depolarize a neuron to threshold for action potential generation. As a result, it is difficult to predict the flow of activity through a synaptic network based solely on knowledge of single connections, without the context of ongoing activity in the entirety of the system.

Network models are an important tool for linking synaptic connectivity to dynamics in neocortex because they enable precise measurement and manipulation of simulated connectivity. In this work, we generate networks comprised of leaky integrate-and-fire model neurons with naturalistic dynamics that mimic recordings from superficial neocortical layers. Despite random synaptic topology in the model network, we find that small-world topological organization emerges in maps of propagating activity. This paradoxical divergence of dynamics from synaptic

---

[7] Behabadi and Mel, “Mechanisms Underlying Subunit Independence in Pyramidal Neuron Dendrites”; Cazé, Humphries, and Gutkin, “Passive Dendrites Enable Single Neurons to Compute Linearly Non-Separable Functions”; Jahnke, Timme, and Memmesheimer, “Guiding Synchrony through Random Networks”; Shai et al., “Physiology of Layer 5 Pyramidal Neurons in Mouse Primary Visual Cortex”; Xu et al., “Nonlinear Dendritic Integration of Sensory and Motor Input during an Active Sensing Task.”

[8] Destexhe, Rudolph, and Paré, “The High-Conductance State of Neocortical Neurons in Vivo”; Gasparini and Magee, “State-Dependent Dendritic Computation in Hippocampal CA1 Pyramidal Neurons”; Tran-Van-Minh et al., “Contribution of Sublinear and Supralinear Dendritic Integration to Neuronal Computations.”

[9] Azouz and Gray, “Cellular Mechanisms Contributing to Response Variability of Cortical Neurons in Vivo”; Paré and Lang, “Calcium Electrogenesis in Neocortical Pyramidal Neurons in Vivo”; Sherman, “The Function of Metabotropic Glutamate Receptors in Thalamus and Cortex.”

[10] Destexhe, Rudolph, and Paré, “The High-Conductance State of Neocortical Neurons in Vivo”; Haider, Häusser, and Carandini, “Inhibition Dominates Sensory Responses in the Awake Cortex.”

connectivity is not explained by coactivity alone. Rather, among active connected pairs, recruitment preferentially occurs in a selective subset.

In the model, activity is preferentially routed through clustered fan-in triangles, despite their statistical scarcity. Because they result in coordinated presynaptic timing, fan-in triangle motifs are particularly effective for spike generation. By comparison, among neurons converging on a common target but lacking presynaptic interconnectivity, presynaptic timing is less synchronous on average, and postsynaptic recruitment is less likely. Moreover, when we decrease the need for cooperative presynaptic action, by doubling synaptic weights in network models, the fan-in triangle motif becomes significantly less prevalent. We evaluate the prediction of our model using high speed two-photon imaging of emergent network activity *ex vivo*, in somatosensory cortex. We verify that propagating activity in real neuronal networks has small-world characteristics and elevated clustering, Decomposing this clustering, we discover that neocortical circuitry also manifests propagating activity that is dominated by the fan-in triangle motif. These results suggest a mechanistic account for the widespread findings of clustered activity in neuronal populations[11]. We suggest that clustered fan-in triangles are a canonical building block for reliable cortical dynamics.

RESULTS

Representing activity and connectivity with directed graphs

---

[11] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches"; Orlandi et al., "Noise Focusing and the Emergence of Coherent Activity in Neuronal Cultures"; Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry"; Shimono and Beggs, "Functional Clusters, Hubs, and Communities in the Cortical Microconnectome"; Nigam et al., "Rich-Club Organization in Effective Connectivity among Cortical Neurons."

Multineuronal dynamics are the computational substrate for sensation and behavior, implemented by synaptic architectures. Propagating multineuronal activity arises from three main sources: the underlying connectivity itself, recent network history, and the non-linear integrative properties of individual neurons. Here, multineuronal activity was modeled using conductance-based leaky integrate-and-fire neurons, stimulated with brief periods of Poisson input and recorded during self-sustained firing (Fig 1a). Model neurons were connected with heterogeneous synaptic weights drawn from a heavy-tailed distribution, in a random arrangement

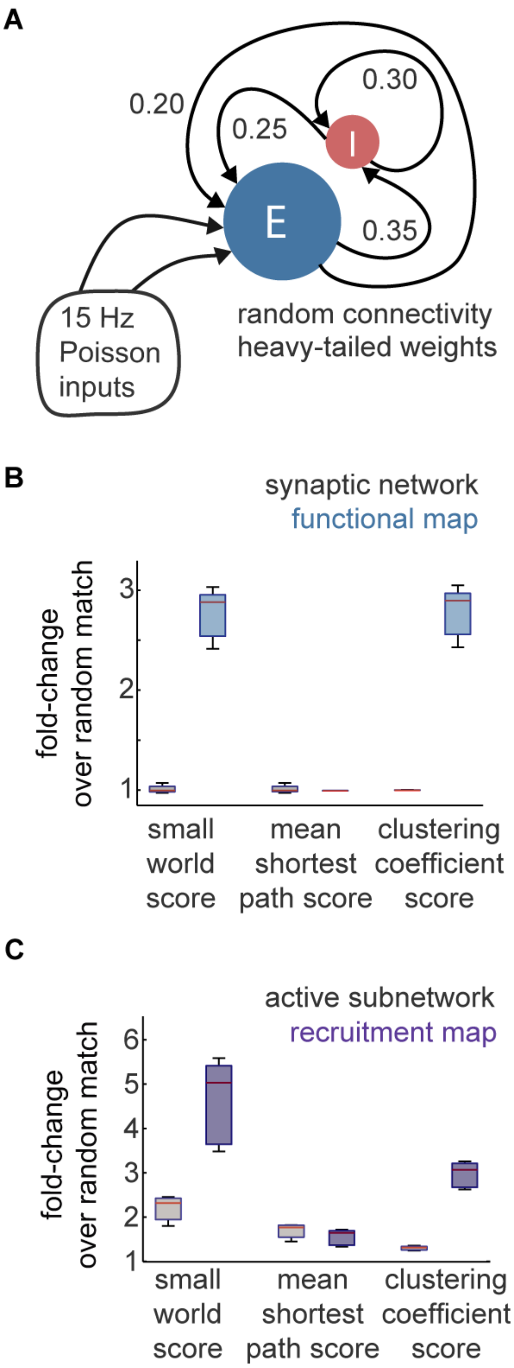

Figure 3-1. Emergent functional networks are structured despite random synaptic connectivity.

(a) Integrate-and-fire neurons with conductance-based synapses were connected randomly according to source and target class (200 inhibitory and 1000 excitatory cells). Activity was initiated with 50 ms of independent Poisson inputs.
(b) Box plots of the fold change over random for the small world score, shortest path length score, and clustering coefficient score in the synaptic network and the functional network.
(c) Box plots of the fold change over random for the small world score, shortest path length score, and clustering coefficient score in the active subnetwork and the recruitment network.

(Erdős-Rényi; p_{ee}=0.2). Simulated dynamics were asynchronous, irregular, and sparse, with critical branching (see *Methods*).

A synaptic network was constructed for each simulation, consisting of excitatory model neurons and their synaptic connectivity. For each structural iteration of the model we generated three distinct maps of activity (and in two of the cases, multiplex connectivity and activity): a *functional network*, the *active subnetwork*, and a *recruitment network* (Fig 2). Edges in the functional network summarized network dynamics and represented frequency of lagged firing between every pair of nodes (with maximum interspike interval T = 25 ms; see *Methods*). The active subnetwork was a subgraph of the synaptic network and consisted of model neurons active at least once and all their interconnections (regardless of lagged firing relationships). Finally, the recruitment network was a subgraph of the functional network defined by its intersection with the synaptic network, to map the routing of activity through synaptic interactions. In this way, non-zero edges in the recruitment network linked synaptically connected nodes that also spiked sequentially in the interval T at least once. For T = 25 ms, 10.9 ± 3.52 excitatory presynaptic input spikes immediately preceded each postsynaptic spike (mean±std).

Surprisingly, although underlying synaptic connectivity was Erdős-Rényi (i.e. random), functional activity networks were small world (Fig 1b)[12]. To judge the small world character of these networks, global clustering coefficient and characteristic path were normalized by their respective abundances in density-matched Erdős-Rényi networks and combined as a quotient[13]. Comparison with density-matches was important given that sparseness itself results in enhanced smallworldness[14].

Functional networks were marked by significantly increased small world scores (functional network: 2.8±0.23; synaptic network: 1.0±0.035; n=5, p=0.0079, *Wilcoxon rank-sum*) resulting from increased clustering (function: 2.8±0.23; synaptic network : 1.0±0.035, n=5, p=0.0079), with characteristic path lengths similar to random-matches (function: $1.0\pm6.4x10^{-4}$; synaptic network: 0.99±0.033; n=5, p=0.69). The lag interval T was chosen to encompass

---

[12] Watts and Strogatz, "Collective Dynamics of 'small-World'networks."

[13] Humphries and Gurney, "Network 'small-World-Ness.'"

[14] Hlinka, Hartman, and Paluš, "Small-World Topology of Functional Connectivity in Randomly Connected Dynamical Systems."

Network	Vertices	Edges	Description
synaptic 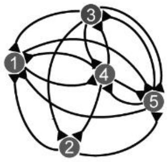	excitatory model neurons	$\mathbf{W}_e$	synaptic connections among excitatory neurons
active 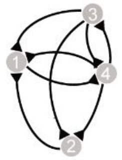	excitatory model neurons active at least once	**if** activity rate(*pre*) > THRESH **then** $\mathbf{W}_{active}(pre, post) = \mathbf{W}_e(pre, post)$ **else** $\mathbf{W}_{active}(pre, post) = 0$	active cells and their connections
functional 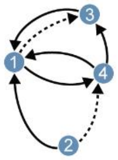	excitatory model neurons active at least once	$\mathbf{W}_{lag}(pre, post)$ = P(*pre* recently active | *post* active)	conditional lag-correlated activity
recruitment 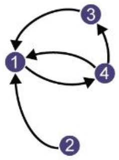	excitatory model neurons active at least once	**if** $\mathbf{W}_e(pre, post) > 0$ **then** $\mathbf{W}_{routing}(pre, post)$ = P(*pre* recently active | *post* active) **else** $\mathbf{W}_{routing}(pre, post) = 0$	intersection between synaptic and functional (routing of activity)
subthreshold voltage for window *k* 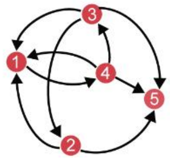	excitatory model neurons	**if** $\mathbf{W}_e(pre, post) > 0$ **then** $\mathbf{W}^k_{voltage}(pre, post) = P(\mathbf{a}_k < \mathbf{V}^{post} <= \mathbf{b}_k \mid pre \text{ recently active})$ **else** $\mathbf{W}_{voltage}(pre, post) = 0$	probability postsynaptic membrane potential lies in window *k* given presynaptic firing

Figure 3-2. Glossary of network definitons.

important network timescales for synaptic plasticity and integration[15]. We also generated functional networks using intervals of 10 and 50 ms, which showed that the emergence of non-random features does not depend strongly on choice of *T* (functional network for *T*=10ms: small world ratio 3.2±0.24, *n*=5, *p*=0.0079; functional network for *T*=50ms: small word ratio 2.6±0.22, *n*=5, *p*=0.0079).

---

[15] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints"; Kruskal, Li, and MacLean, "Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex."

Given modest sampling conditions (e.g. binning near timescales of synaptic integration), functional relationships can indicate locations of probable synaptic recruitment[16]. However, a subset of edges in functional networks are 'false positives'—they reflect polysynaptic relationships and other combined statistical dependencies rather than monosynaptic connectivity and recruitment[17]. To determine whether these measurement artifacts were responsible for the statistical differences between function and synaptic networks, we turned to recruitment networks. Pruned of false positives, recruitment networks were significantly more small world than functional networks constructed from the same activity (4.6±0.87; n=5, p=0.0079), with even shorter characteristic paths (recruitment: 0.65±0.072, n=5, p=0.0079 compared to function, *Wilcoxon rank-sum*) and a similar elevation in clustering (recruitment: 3.0±0.26; n=5, p=0.22). Thus, emergent statistical structure in the functional networks reflected coordinated timing among multiple synaptically connected neurons.

Preferential routes for propagating activity

As demonstrated by non-random recruitment, i.e. clustering in the recruitment network, activity did not propagate homogeneously through the random topology. However, it remained a possibility that the seemingly non-random routing of activity was simply the byproduct of shared activity, without being selective on the basis of connectivity. As a control, the active subnetwork establishes the role of interactions among neurons with elevated firing rates (including pairs of neurons which never recruited one another within the interval T). Compared to functional networks, the corresponding active subnetwork exhibited reduced small world ratio (active

---

[16] Chambers and MacLean, “Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints.”

[17] Ibid.; Gerstein, Perkel, and Subramanian, “Identification of Functionally Related Neural Assemblies.”

network: 2.2±0.26, *n*=5, *p*=0.0159) and reduced clustering (1.3±0.041, *p*=0.0079), despite somewhat shorter characteristic paths (0.60±0.055, *n*=5, *p*=0.0079).

If directed connections that never fired sequentially were pruned from the active subnetwork, it would attain the same binary topology as the recruitment network. Comparing the active network with the recruitment network, global clustering ratio was significantly increased (from 1.3±0.041 to 3.0±0.26, *n*=5, *p*=0.0079, *Wilcoxon rank-sum*). Thus, the select connections which were directly involved in propagation of spiking activity were more clustered than activated connections as a whole (Fig 1c).

We next evaluated whether neuronal pairs that never fired sequentially differed from those that did. Comparisons were performed between in-degree matched samples. Connected neurons that never fired in succession shared significantly fewer neighbors than those that did fire sequentially at least once (*n*=500 pairs, $p=3.1 \times 10^{-17}$, *Wilcoxon rank-sum*). In the model, activity was selectively routed through interconnected neighborhoods.

Deconstructing patterns of directed clustering

Connectivity within a triplet is the simplest way two nodes can share a common neighbor and be clustered. However, this measure fails to account for the direction of connection. Since direction is crucial in synaptic communication, we turned to a formulation which differentiates directed triangle motifs[18]. From the perspective of a reference postsynaptic neuron, clustered neighbors can be arranged into four kinds of three-edge triangle motifs: fan-in, fan-out, middleman, and cycle arrangements (Figure 3a). Taken in isolation, fan-in, middle-node, and cycle triangles are isomorphic to one another through rotation, i.e. dependent on labeling the reference

---

[18] Fagiolo, "Clustering in Complex Directed Networks."

node (which is necessary to compute local clustering). Measures of undirected clustering can be decomposed fractionally into these four components. Because the underlying model synaptic connectivity was random, none of the four triangle motifs were more prevalent than the others, and each contributed equally to synaptic clustering (Figure 3b). By contrast, in recruitment networks, fan-in triangle motifs were highly overrepresented (Figure 3c). The overrepresentation of fan-in triangle motif was also present in the functional network: for example, iterative Bayesian inference[19] was sensitive to asymmetric directed clustering in model activity (fan-in: 0.38±0.052, fan-out: 0.29±0.032, middleman: 0.19±0.016, cycle: 0.15±0.0076; mean±std, threshold at the 95th percentile).

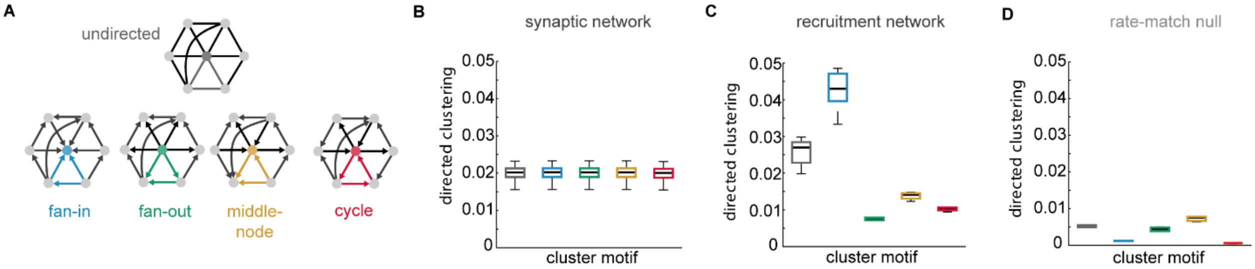

Figure 3-3. Clustered fan-in triangle motifs dominate recruitment networks.

(a) Scheme for factoring transitive clustering into constituent directed patterns.
(b) Boxplots of the prevalence of the directed patterns in the randomly connected synaptic network.
(c) Boxplots of the prevalence of the directed patterns in the recruitment network.
(d) Boxplots of the prevalence of directed patterns in nulls generated from rate-matched Poisson populations without synaptic interactions analyzed with iterative Bayesian inference,

---

[19] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

To understand whether these higher order asymmetric features emerge from chance correlations tied to firing rates, we generated Poisson populations that were rate-matched on a neuron-by-neuron and trial-by-trial basis. This resulted in an inhomogeneous distribution of firing rates across all trails. Our Poisson null populations had identical expected spike counts as model activity in each 100ms bin but no synaptic interactions and no causal propagation of activity. Undirected clustering was significantly lower in iterative Bayesian maps of uncoupled Poisson rate-matched activity compared to connected network models (Poisson rate-match:

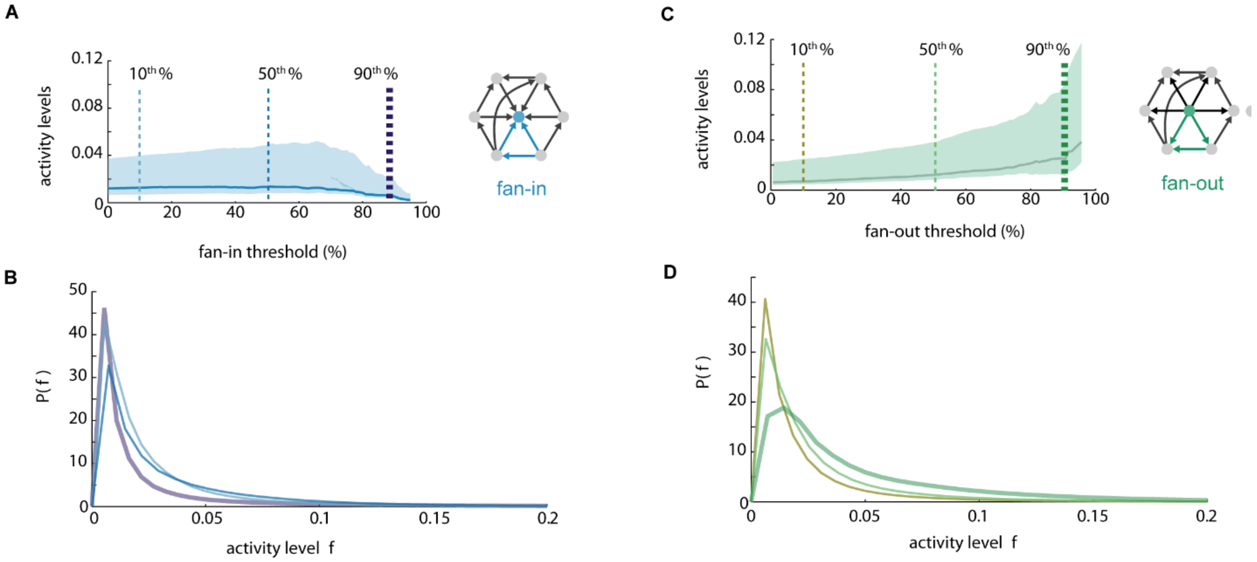

Figure 3-4. Fan-in triangle motifs are not simply the result of firing rate.

(a) Median firing rate across all fan-out clustering thresholds (middle quartiles shaded).
(b) Probability distribution of firing rates for fan out triangle motif. Firing rate distributions for subpopulations thresholded to exclude the bottom 10% (mustard), 50% (light green), and 90% (dark green) of fan-out clustered model neurons.
(c) median firing rate across all fan-in clustering thresholds (middle quartiles shaded).
(d) Probability distribution of firing rates for fan in triangle motif. Reference cells with high fan-in clustering had lower firing rates than the population as a whole: bottom10% (light blue), 50% (blue), and 90% (purple) of fan-in clustered model neurons.

$0.0052 \pm 3.6x10^{-4}$; simulated activity: 0.024±0.013; *Wilcoxon rank-sum p*=0.036; *n*=3), and the fan-in triangle motif was not elevated relative to other clustering patterns (Figure 2d). The Poisson populations demonstrated that elevated fan-in triangle motifs do not result trivially from the analysis procedure but instead are the result of synaptic interactions between neurons. Interestingly, we found that model neurons with high fan-out clustering were characterized by elevated firing rates (Figure 3a,b), but model neurons which comprised the fan-in triangle motif actually contracted towards low firing rates (Figure 3c,d). Fan-in triangles were more abundant in propagating activity than would be expected from their frequency in the synaptic network or component firing rates alone. Like undirected clustering the emergence of fan-in clustering in maps of propagating activity was robust to choice of *T*. Fan-in clustering was highly elevated in recruitment maps for *T*=10 ms (undirected: 0.0068±0.0007; fan-in 0.011±0.0017; fan-out: 0.0028±0.0001; middle-node: 0.0068±0.0007; cycle: 0.0052±0.0004; mean±std for 5 simulations) and *T*=50 ms (undirected: 0.019±0.0015; fan-in 0.027±0.0027; fan-out: 0.0077±0.0003; middle-node: 0.019±0.0013; cycle: 0.015±0.0007 ; mean±std for 5 simulations). Because of the different levels of sparseness in the numbers of connections these values should not be compared across values of T. Instead these analyses demonstrate that the over-representation of fan-in triangles is robust across a number of time scales.

Activity at fan-in triangle motifs is temporally organized

To investigate the mechanism for overrepresentation of fan-in triangles in recruitment networks, we measured spike timing at their locations. The signature of fan-in triangle motifs is convergence from interconnected presynaptic neurons, a motif that could potentially facilitate cooperative summation of synaptic inputs. Consistent with this postulate, presynaptic neurons in fan-in triangle motifs were marked by increased probability of firing in the 10 ms prior to postsynaptic spiking (Figure 5a, 5b).

We next compared differences in presynaptic timing relationships at loci of fan-in triangle motifs compared to loci of simple convergence, to assess the role of presynaptic interconnectivity. For this analysis, random samples were obtained from epochs of *coincident firing*: 50 ms windows where every neuron in a triplet was active, centered on a spike in the postsynaptic reference neuron. To avoid confounds from juxtaposing multiple motifs, neuron triplets with any additional connections, including recurrent loops, were excluded for this specific analysis alone. As a result only fan-in triangles with exactly three interconnections were analyzed in this case. We found fan-in presynaptic neurons were stereotypically ordered in a manner consistent with the direction of their interconnection, resulting in an asymmetric

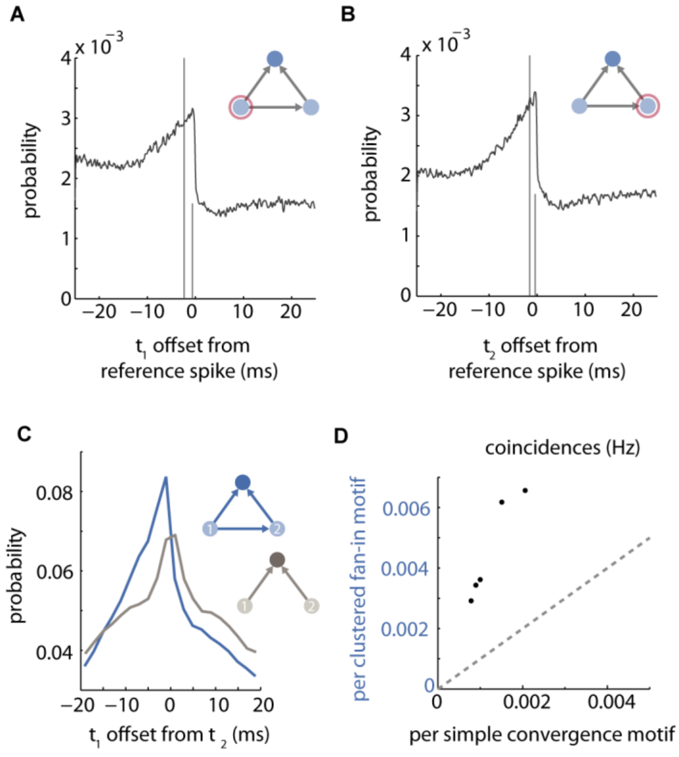

Figure 3-5. Coordinated timing among model neurons in fan-in triangle motifs.

(a,b) Probability distribution of spiking within fan-in triangle motifs. Postsynaptic spiking at t=0 (tall mark, center of mass; short mark, peak).
(c) Probability versus difference in presynaptic timing t2 – t1 during coincident epochs.
(d) Rate of observing coincident firing [50 ms] in fan-in triangle motifs

distribution of intervals between their firing (Figure 5c). In addition to the temporal structure imposed by this asymmetry, mean absolute timing difference between presynaptic neurons in clustered fan-in motifs was modestly but significantly more temporally precise than were neurons in simple convergence motifs (13.5±10.2 ms compared to 14.9±10.7 ms; *Wilcoxon rank-sum* on mean-absolute timing difference, p=0.0035, n=1000 samples).

Moreover, we found that coincidence in fan-in triangle motifs occurred nearly twice as frequently as in motifs of simple convergence (1.9 ± 0.17 times more frequent, *mean* ± *std*; *Wilcoxon rank-sum*, p=0.0079, n=5 model datasets). Accounting for expected frequency of the two connection patterns in the underlying synaptic network, coincident activity is far more common at sites of fan-in triangles than at sites of simple convergence (linear regression: slope 3.0, y-intercept 0.00075, n=5 simulations, r^2=0.91, p=0.011) (Figure 5d).

Increasing clustering among active inputs with depolarization

We postulated that clustering is efficacious for synaptic integration and examined whether the prevalence of clustering was predictive of postsynaptic membrane potentials. Pooling over all neurons and time bins, we binned the distribution of membrane voltages into segments that contained equal numbers of samples (Figure 6a). On average, because the model was active in the analyzed simulations, membrane voltages were depolarized from the resting equilibrium potential of -65 mV (median: -60.2 mV; lower quartile: -63.6 mV; upper quartile: -56.8 mV). To test our hypothesis, we generated functional networks that related recent presynaptic activity (within a 25 ms interval) to postsynaptic voltage (Figure 6b; see *Methods)*, yielding one network for each division of the voltage distribution (Figure 6c). These networks can be viewed as reverse correlograms conditioned on postsynaptic voltage, and differed in the statistics of their topologies across different voltage regimes. At more negative membrane

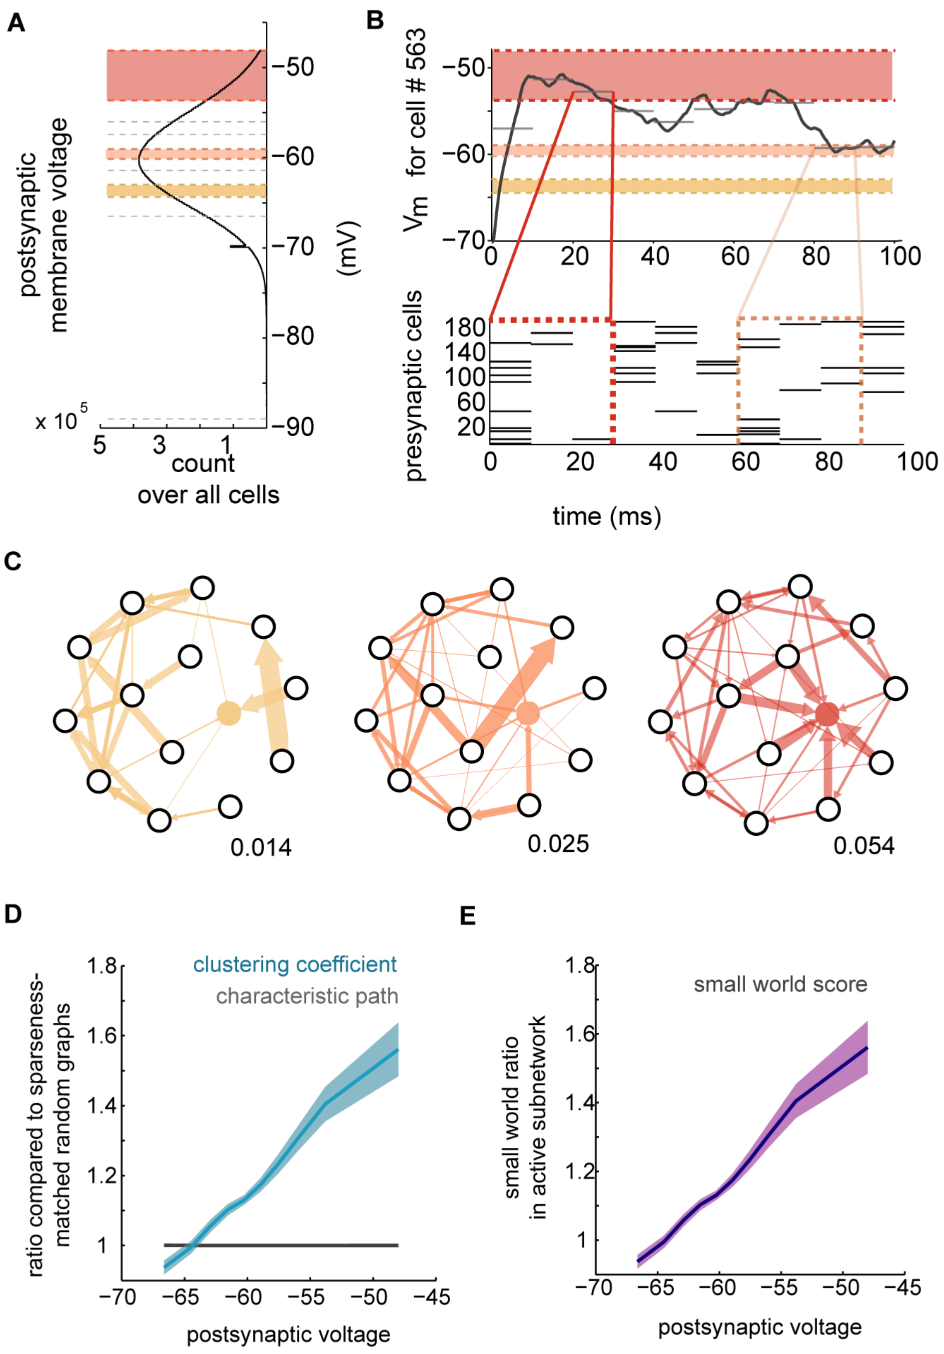

Figure 3-6. The prevalence of fan-in triangle motifs increased with post-synaptic voltage.

(a) Distribution of postsynaptic voltage. Shading corresponds to (b) and (c) and contain the same number of samples per voltage bin. (b) Postsynaptic voltage was mapped in relation to presynaptic spiking.
(c) One example weighted directed topology for each division of the voltage distribution.
(d) Ratio versus voltage (clustering coefficient: blue, characteristic path: gray, shading reflects one standard deviation).
(e) Small world ratio versus voltage (shading reflects one standard deviation).

potentials, the active neurons which connected to the postsynaptic reference neuron (and accounted for its recent excitatory synaptic drive) were only modestly more clustered than random sparseness-matched controls. As the postsynaptic neuron depolarized, the presynaptic nodes driving that depolarization became increasingly clustered, peaking at the threshold for firing (Figure 6d). Characteristic paths were similar to random graphs at all subthreshold voltages. As a result of elevated clustering during membrane depolarization, small world ratios peaked at the most depolarized voltages corresponding to threshold for action potential generation. These data support the hypothesis that activity among clustered presynaptic neurons is particularly effective for recruiting the postsynaptic neuron to spike.

Emergence of higher-order features depends on mean synaptic weight

The statistical incongruence of function and synaptic connectivity indicates that spiking activity does not flow in an egalitarian fashion through the synaptic network. Instead, patterns of local clustering influence and direct where propagating activity occurs most frequently. That is, patterns of activity are shaped by higher-order patterns in synaptic connectivity and not just pairwise couplings. To further explore the dependence of activity flow on higher order synaptic connections we evaluated postsynaptic recruitment in a network model with a modest increase in mean synaptic strength. Synaptic connections were twice as strong on average compared to the network models used throughout the remainder of this study but remained too weak to drive spiking alone (Figure 7a). The two network designs did not differ in connection density. After synaptic weights were doubled, functional networks became more similar in topology to synaptic networks (small world ratio decreased; Wilcoxon rank-sum, p=0.0079, n=5) (Figure 7b).

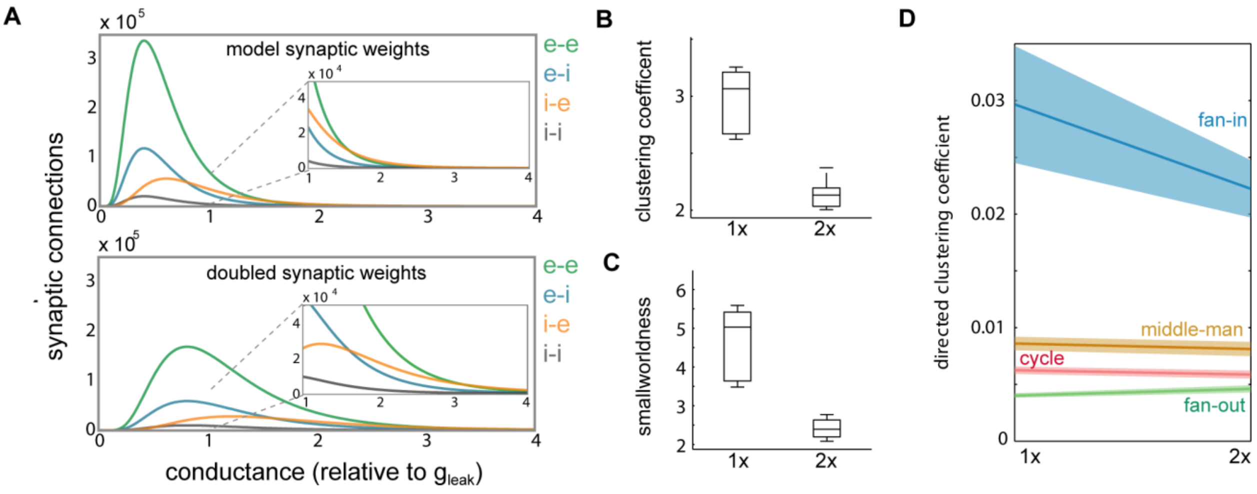

Figure 3-7. Increased synaptic weights reduced higher-order functional coordination. (a) Distribution of model synaptic weights (excitatory-excitatory: green, excitatory-inhibitory: blue, inhibitory-excitatory: orange, inhibitory-inhibitory: gray). Top: naturalistic model. Bottom: double-strength model. Inset in both cases shows zoom to better illustrate heavy tail. (b) Box plot of clustering coefficient in the two models. 2X indicated double-strength synaptic connections. (c) Box plot of small worldness in the two models. 2X indicated double-strength synaptic connections. (d) Comparison of mean directed clustering with each model iteration on either side. Each class of directed clustering is labeled in the plot.

The double-strength models were less clustered (Figure 7c) (Wilcoxon rank-sum, *p*=0.0079, *n*=5), and exhibited longer average path lengths (Wilcoxon rank-sum, *p*=0.0079, *n*=5). Directed clustering was compared across the two families of models. Recruitment networks were analyzed with binary edges to control for their distinct mean synaptic weights. In addition to their decreased overall clustering, the fan-in triangle motif was significantly rarer in double-strength recruitment networks (Figure 7d) (from 0.030±0.0051 to 0.022±0.0025, p=0.030, n=6), while the fan-out triangle motif showed a small but significant increase in abundance (from $0.0040 \pm 2.0x10^{-4}$ to $0.0046 \pm 3.2x10^{-4}$, p=0.0043, n=6). Stronger presynaptic inputs reduced the need for extensive postsynaptic integration, allowing individual presynaptic cells to have a more independent impact on their postsynaptic partners. As a result, statistics of propagating activity were more faithful to underlying pairwise connections in the models with increased synaptic strength.

Fan-in triangles characterize neocortical circuit dynamics

In model simulations, fan-in triangle motifs were abundant in maps of function and recruitment. We next evaluated whether the preponderance of fan-in triangle motifs would be robust to additional complexity in single-neurons and their connections. Unlike the simple model neurons that we used for simulation, real neurons are complex elements[20] and the connections between them are structured[21]. If clustered fan-in triangle motifs are a general feature of high-conductance nodes in a complex system, where coordinated inputs drive integration, the fan-in triangle will be overabundant in experimental dynamics. This postulate would be falsified if all directed clustering motifs were equally common in functional networks. To investigate, we analyzed high speed imaging data (20 Hz) of spontaneous circuit activity collected *ex vivo* in mouse somatosensory cortex (Figure 8a) (following [22]). We generated functional networks from the imaged experimental data using an iterative Bayesian approach which is robust to relatively small numbers of observations (33). We then measured the prevalence of fan-in motifs in the functional topology (Figure 8b). Importantly, iterative Bayesian inference was not biased toward detection of fan-in triangle motifs, as demonstrated with rate-matched Poisson spiking (see Figure 3d).

---

[20] Behabadi and Mel, “Mechanisms Underlying Subunit Independence in Pyramidal Neuron Dendrites.”

[21] Perin, Berger, and Markram, “A Synaptic Organizing Principle for Cortical Neuronal Groups”; Song et al., “Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits.”

[22] Sadovsky et al., “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.”

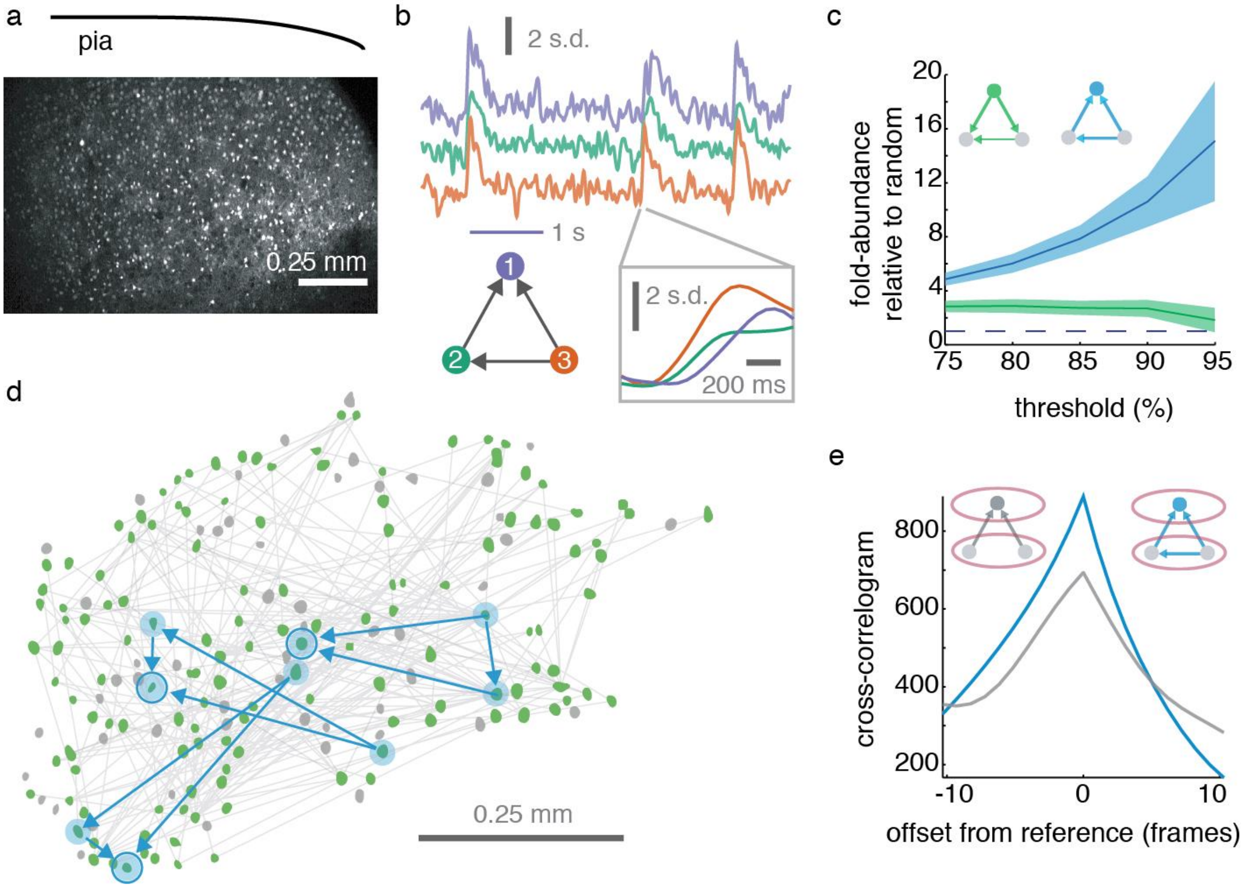

Figure 3-8. Clustering in experimentally recorded emergent cortical circuit activity was characterized by the fan-in triangle motif.

(a) Two photon image of Ca2+ indicator dye in a slice of mouse somatosensory cortex.
(b) Example *z*-scored fluorescent traces among functionally related neurons identified as members in a fan-in triangle motif.
(c) Relative abundance of both fan-in (blue) and fan-out (green) clustering relative to density-matched random graphs as a function of inclusion threshold on inferred connections. Shading reflects one standard deviation based on bootstrap resampling for a 30% false positives rate (*n*=100).
(d) Top 5% of functional edges (light grey) and corresponding neurons (green) collected from the same field of view as a. Grey neurons were also active but were not connected with edges exceeding the cut-off. Three example fan-in triangles are illustrated with directed edges (blue arrows) and blue shading (motif specific neurons). Circle indicates reference neuron in each triangle.
(e) Separate functional analysis of triplet motifs using cross-correlation. For each triplet, the product of the *z*-scored presynaptic traces were compared to the postsynaptic trace. Fan-in triangle motifs were characterized by higher levels of coordination than motifs of simple convergence.

Though imperfect indicators, functional weights probabilistically identify the likelihood of true monosynaptic excitatory connectivity[23]. As a result, expected error rate for inferred connections can be adjusted with a sliding threshold on functional weight. Stricter thresholds yield a more accurate approximation of the underlying recruitment network at the cost of restricted sampling. Using inferred recruitment networks, beginning at the top quartile of inferred weights, directed clustering was computed in five-percentile increments. Confidence intervals were obtained using bootstrap resampling under the assumption of a 30% false-positive rate. As confidence of synaptic connectivity increased, the fan-in triangle motif became increasingly abundant and fan-out triangles less so (Figure 8c). Differences between the two motifs were significant (threshold at 95th percentile, p=4.8x10^{-34}, n=100 bootstrap resampled functional networks, *Wilcoxon ranksum*).

We next measured whether strong functionally coupled neurons were more spatially proximal than random pairs. We defined strong functional connections as those exceeding a 95% threshold on non-zero weights since previous work has indicated that these particular functional connections are more likely to reflect a causal synaptic connection. We found that the median pairwise distance separating strong functionally connected cells was 249 μm, whereas randomly chosen pairs of neurons were separated by a median 263 μm (Wilcoxon-ranksum p=0.0336, $n^{\text{functional}} = 638$, $n^{\text{random}} = 10000$). We then measured triplets of neurons with functional connections that form triangles to determine whether these neurons were more spatially proximal to one another than randomly chosen triplets of neurons. To investigate, proximity was quantified as the perimeter around the triangle formed by vertices at the spatial location of each neuron. Neurons in functional triangles with mutual connectivity and at least three functional connections were inscribed by

---

[23] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

perimeters of median length 807 μm, compared to median perimeter of 823 μm for randomly selected triplets that were unconstrained by direction and number of edges (Wilcoxon rank-sum p=0.0097, $n^{triangles} = 2556$, $n^{random}=10,000$). Interestingly, triplets of neurons connected into arrangements of either simple divergence or simple convergence (i.e. neurons in wedges, lacking interconnectedness between the common neighbors), were even more distant, inscribed by a perimeter of median 839 μm (Wilcoxon rank-sum, $n^{triangles} = 2556$, $n^{wedges} = 14,882$). Thus, clustered triplets (triangles) tended to be arranged significantly more locally than simple convergent or simple divergent triplets (wedges).

We then compared measures of clustering between the model, which was comprised of random connections and the experimental data which almost certainly contained structured connectivity (Song, Perin) to evaluate how the measure of fan in and fan-out triangles depend on the underlying structural topology. To do so we used a measure of clustering propensity[24] which allowed us to make comparisons of networks which have very different connection densities. Clustering propensity ($1-\Delta C^{fan\text{-}in}$ and $1-\Delta C^{fan\text{-}out}$) results in a normalized value where 1 is extreme clustering as seen in lattices, and 0 indicates no clustering above that expected in Erdős-Rényi random networks. For the model, fan-in clustering was scored at 0.18 ± 0.019; and for the experimental data, fan-in clustering was scored at $0.20 \pm 2.0\text{x}10^{-4}$ (Wilcoxon ranksum $p=1.74\text{x}10^{-4}$, $n^{model} = 5$ simulations; $n^{data} = 100$ bootstrap samples). Thus, fan-in clustering was modestly but significantly more abundant in maps of propagating activity based on experimental recordings. We note that we compared thresholded graphs at the 80%-level (i.e. top 20% of non-zero edges) for this measure because the experimentally derived functional networks were not well-matched by regular lattices below this density.

---

[24] Muldoon, Bridgeford, and Bassett, "Small-World Propensity and Weighted Brain Networks."

Finally, we measured timing relationships among imaged active neurons. Reliable timing relationships were measured independent of other functional analyses, using cross-correlations on the normalized fluorescence traces (*Methods*). Presynaptic coactivity was assessed as the product of the two *z*-scored presynaptic traces and compared to postsynaptic fluorescence as a straightforward cross correlation. The resulting average cross-correlogram for fan-in triangles was stronger and more asymmetric than those measured from simple-convergence motifs (Figure 8d).

Thus, presynaptic activity in fan-in triangles was more predictive of postsynaptic firing than presynaptic activity in motifs of simple convergence. These results are consistent with fan-in triangles supporting coincident input and favoring reliable propagation of activity. Results from the model indicated that the fan-in triangle motif temporally coordinates presynaptic inputs, rendering them more capable of driving recipient neurons to threshold. Supporting our prediction of its fundamental importance for reliable recruitment, in acutely dissected neocortical tissue with more complex patterns of connectivity and intrinsic neuronal properties, we find a robust elevation of the same directed motif.

DISCUSSION

Using a model composed of random connections among leaky integrate-and-fire neurons with conductance-based synapses, we found that maps of propagating activity were structured and non-random. Small-world patterning in the dynamics emerged because a specific higher-order connection pattern was particularly effective for postsynaptic integration: convergence of synaptic input from connected neighbors. The synaptic connection between neighbors favored coincident timing of inputs onto their targets. This coincident activation led to efficient postsynaptic integration. As a consequence, clustering among active presynaptic cells tracked depolarization of model postsynaptic neurons. Thus, activity was preferentially routed through fan-in triangle motifs.

In experimental recordings of emergent activity in hundreds of neurons *ex vivo*, after mapping inferred recruitment patterns (33), we found that fan-in triangles were even more dramatically overrepresented than in the model. These results are contextualized by increasing recognition of non-random functional structure in networks of neurons: Rich club structure has been reported *ex vivo* and *in vivo*[25]. Clustered[26], small world functional networks[27], and nucleation of dynamics[28] have also been observed in neuronal cultures. Since cultured populations differ from neocortex in the details of their topological makeup, these findings across model systems further suggest that clustering in general and the fan-in triangle motif in particular may be a canonical feature of propagating activity among interconnected neurons. Despite differences in details of connectivity and neuronal intrinsic properties, dynamics are constrained by the requirement for coincident summation of individually weak inputs. Constraining dynamics beyond pairwise relationships can be helpful for cortical computation. Theoretical work has shown that non-uniform features of connection topology impact information transfer[29], and higher-order correlations were particularly impactful in low spike-rate regimes[30]. These complementary results from complex networks, statistical physics and network biology suggest that, by shaping feasible dynamics, the fan-in triangle motif could enhance information transfer from inputs to outputs.

---

[25] Nigam et al., "Rich-Club Organization in Effective Connectivity among Cortical Neurons."

[26] Shimono and Beggs, "Functional Clusters, Hubs, and Communities in the Cortical Microconnectome."

[27] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."

[28] Orlandi et al., "Noise Focusing and the Emergence of Coherent Activity in Neuronal Cultures."

[29] Timme et al., "High-Degree Neurons Feed Cortical Computations."

[30] Zylberberg and Shea-Brown, "Input Nonlinearities Can Shape beyond-Pairwise Correlations and Improve Information Transmission by Neural Populations."

We hypothesize that local circuits are organized around fan-in triangle motifs, promoting cooperative patterns of firing and stabilizing[31] the propagation of activity despite individually unreliable neurons. This canonical mechanism provides the coordination necessary to propagate signal despite weak synaptic connections. Indeed, reliable sequential firing was associated with number of fan-in triangles even after controlling for overall in-degree. Although clustering among fan-in triangles has not been tested directly until now, paired patch clamp recordings have shown that local neocortical circuitry is characterized structurally by abundant triplet motifs[32]. Our data and modeling suggest a functional consequence for a subset of these synaptic motifs: connected presynaptic neurons help establish coordinated timing among convergent inputs, leading to cooperative summation at the postsynaptic membrane. Such cooperativity has been shown to be one potential mechanism capable of generating spike trains that are consistent with experimental observations *in vivo*[33].

While there are certainly explicit developmental rules that govern neuron to neuron connectivity, our results suggest that higher-order connectivity need not require specification *a priori*. It could emerge autonomously if fan-in triangle motifs within a random network were stabilized and magnified during network development, *e.g.* by pruning non-recruiting connections through activity-dependent plasticity. Thus, higher-order synaptic motifs that are particularly effective for postsynaptic recruitment could potentially self-organize[34].

---

[31] Angulo, Liu, and Slotine, "Network Motifs Emerge from Interconnections That Favour Stability."

[32] Perin, Berger, and Markram, "A Synaptic Organizing Principle for Cortical Neuronal Groups"; Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."

[33] Softky and Koch, "The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs."

[34] Miner and Triesch, "Plasticity-Driven Self-Organization under Topological Constraints Accounts for Non-Random Features of Cortical Synaptic Wiring."

These results do not indicate a complete schism between synaptic connectivity and dynamics—one clearly depends on the other. However, their relationship is complicated by the integrative properties of single neurons. Synaptic integration constrains feasible dynamics, and distributed synaptic motifs route the propagation of activity. These interactions are a source of higher-order dynamical structure. The routing of information is coordinated by higher-order synaptic patterns and the context of ongoing activity because the routing of spikes is determined by relative timing and collective interactions.

MATERIALS AND METHODS

Local cortical population model

Simulations were implemented using the Brian Brain Simulator[35]. Model populations consisted of 1000 excitatory neurons, 200 inhibitory neurons and 50 Poisson input units. Connection probabilities depended on source and target identity. For example, inhibitory-excitatory connections occurred with probability 0.25 ($P_{ee} = 0.2$, $P_{ei} = 0.35$, $P_{ie} = 0.25$, $\mathrm{P}_{ii} = 0.3$).

Conductance based synaptic weights were drawn from a heavy-tailed distribution and assigned randomly[36]. Weights were drawn randomly from a lognormal distribution with mu = -0.64 and sigma = 0.51. These parameters are the mean and standard deviation of the corresponding normal curve. The resulting lognormal ensemble has expected mean of 0.60 and variance of 0.11, in multiples of the leak conductance. Connections from inhibitory to excitatory cells were scaled by a further 50% to simulate efficacious somatic contacts. A small tonic excitatory drive g_t was

---

[35] Goodman and Brette, “The Brian Simulator.”

[36] Lefort et al., “The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex”; Teramae, Tsubo, and Fukai, “Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links.”

supplied to all units to help stabilize sparse spiking. Synaptic bombardments induced exponentially shaped membrane conductances with leaky-integrate-and-fire summation. Conductance-based synapses are important for recapitulating synaptic integration in the high-conductance state[37]. We used sparse and randomly connected networks in which we did not impose any synaptic organization beyond cell-type dependent connection probabilities.

Trials began with 50 ms of activity in the input pool at 15 Hz, exciting the network via random input projections. After input units were silenced, the recording period began, and activity flowed through the network for 100 ms. Input units projecting to excitatory cells randomly and independently with probability 0.1. Every 100 trials (an epoch), new random projections were drawn from the input pool to the excitatory population, simulating a diversity of activity. Participation during a single input epoch totaled 64±0.98% of neurons (mean ± std), growing to encompass 85.5% of neurons when all sets of input projections were considered (i.e. over all epochs).

Excitatory reversal potential E_e was 0 mV, as was E_t. Inhibitory reversal potential E_i was -90 mV. Reversal potential for leak current E_{leak} was -65 mV. Firing threshold was -48 mV, and post-spike reset was -70 mV. In addition to after spike hyperpolarization induced by the reset potential, a 1 ms absolute refractory period was imposed on model neurons. Leak conductance g_{leak} was fixed at 0.20 mS. Tonic depolarizing conductance g_t was equal in magnitude to the leak conductance. Membrane time constant τ_m was 20 ms; excitatory synaptic time constant τ_e was 10 ms; and inhibitory synaptic time constant τ_i was 5 ms. Additional description can be found in[38].

---

[37] Destexhe, Rudolph, and Paré, "The High-Conductance State of Neocortical Neurons in Vivo"; Kumar et al., "The High-Conductance State of Cortical Networks."

[38] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

Quantification of simulated dynamics

Spiking dynamics were compared to *in vivo* activity according to the following criteria: asynchrony[39] was measured with spike-rate correlations, by convolving spike times with a Guassian kernel of width σ = 3 ms. Among excitatory neurons in the recording period, mean correlation coefficient was 0.0019[40]. This asynchrony emerged in the presence of heterogeneous connection strengths, raising the possibility of combining stable propagation with rich internal dynamics[41]. Irregularity was measured with interspike-intervals, which were observed to have mean squared-coefficient of variation of 0.81, consistent with other reports of irregular activity[42]. To measure inter-spike intervals, model activity was stimulated with Poisson firing for 50 ms, then allowed to evolve for 950 ms in isolation. This procedure was repeated 100 times. Excitatory spiking activity was characterized by a median branching coefficient of 1.00 (for 10 ms bins), indicating near-critical dynamics[43]. Firing rates in the excitatory population during the recording period were 1.33 ± 3.15 Hz (mean ± std) consistent with findings in awake behaving mice[44]. Collective spiking generated spike-driven conductances that dwarfed the leak conductance, in keeping with definitions of high-conductance state[45].

---

[39] Brunel, “Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons.”

[40] Kumar et al., “The High-Conductance State of Cortical Networks.”

[41] Teramae, Tsubo, and Fukai, “Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links”; Ostojic, “Two Types of Asynchronous Activity in Networks of Excitatory and Inhibitory Spiking Neurons.”

[42] Kumar, Rotter, and Aertsen, “Spiking Activity Propagation in Neuronal Networks.”

[43] Beggs and Plenz, “Neuronal Avalanches in Neocortical Circuits”; Haldeman and Beggs, “Critical Branching Captures Activity in Living Neural Networks and Maximizes the Number of Metastable States”; Millman et al., “Self-Organized Criticality Occurs in Non-Conservative Neuronal Networks during Up States”; Chialvo, “Emergent Complex Neural Dynamics.”

[44] Crochet et al., “Synaptic Mechanisms Underlying Sparse Coding of Active Touch.”

[45] Destexhe, Rudolph, and Paré, “The High-Conductance State of Neocortical Neurons in Vivo.”

Network construction

Call the directed network of synaptic connections among excitatory neurons E^{syn} and the population of excitatory cells V^e. Construct the directed graph of synaptic connections:

$$G^{structural} \equiv (V^e, E^{syn})$$

To map functional relationships using lagged firing, define *recent activity* for neuron *i* at time *t* as firing at least once in the 25 ms preceding *t.*

$$E_{ij}^{lag} \equiv P(j \text{ active} \mid i \text{ recently active})$$

$$G^{functional} \equiv (V^e, E^{lag})$$

More formally, we can define random variable S_i representing the activity of neuron *i* such that

$$s_i^t \equiv \left\{ \begin{array}{lll} 2 & if & \text{neuron } i \text{ is firing at time } t \\ 1 & if & \text{neuron } i \text{ is not firing but was active within the last 25 ms} \\ 0 & otherwise & \end{array} \right\}$$

In that case,

$$E_{ij}^{lag} \equiv P(s_j = 2 \mid s_i > 0)$$

The recruitment network encompassed synaptically connected neurons manifesting lagged patterns:

$$E_{ij}^{omniscient} \equiv \left\{ \begin{array}{ccc} E_{ij}^{lag} & if & E_{ij}^{syn} > 0 \\ 0 & otherwise & \end{array} \right\}$$

$$G^{recruitment} \equiv (V^{e}, E^{omniscient})$$

Iterative Bayesian networks were measured with a heuristic optimization procedure, described further below and in [46], following [47].

Global network statistics

Since shortest path measurements assume a cost matrix, edge weights were first inverted so strong connections were cheap and zero-weighted connections were infinitely costly. Shortest paths between all pairs were computed using Dijkstra's algorithm. Mean path length was compared to sparseness-matched Erdős-Rényi graphs analyzed in the same way. Local clustering coefficients were computed using the neighbors of neighbors formulation[48] and aggregated as the mean over all neurons. Sparseness-matched Erdős-Rényi graphs were analyzed in the same fashion. Clustering score was the ratio of the actual mean to sparseness-matched null mean. Small-world topologies can be quantified as a ratio of ratios, clustering elevation divided by mean path length reduction[49].

Transitive clustering and directed clustering

---

[46] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

[47] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."

[48] Watts and Strogatz, "Collective Dynamics of 'small-World'networks."

[49] Humphries and Gurney, "Network 'small-World-Ness.'"

Clustering was also investigated using a related definition, the number of connected undirected triangles as a fraction of all possible undirected triangles (transitivity formulation). Directed clustering was computed in the same way, using directed triangles instead of undirected[50].

To compare clustering between data and model networks, across connection densities that were very different, we followed the small-world propensity approach[51]. In that work, clustering levels ΔC are normalized as the fractional distance between density-matched lattice and random graphs. We termed this measure clustering propensity, expressing it as 1 – ΔC so that 1 signified extreme clustering and 0 signified no clustering beyond that expected at random. We made a straightforward extension to this approach to account for directed clustering, simply substituting directed triangle counts for undirected triangle counts, with appropriate normalizations[52]. Quantifications based on clustering propensity recapitulated our findings quantifying clustering as fractional abundance over random expectation.

Mapping presynaptic ensemble in relation to postsynaptic voltage

For the set of voltage bins with lower bounds **a** and upper bounds **b**, construct one network for each bin k, where edge $(i, j)^k$ is quantifying the probability model neuron j will have postsynaptic potential M_j between $\mathbf{a}_k$ and $\mathbf{b}_k$ conditioned on presynaptic model neuron i being recently active. *Recently active* was defined as firing within 25 ms relative to postsynaptic voltage measurement. A final condition was imposed: that connected pairs also share a synaptic connection, a convenience of measurement unique to simulated networks.

$$G^k \equiv (V^e, E^k)$$

---

[50] Fagiolo, "Clustering in Complex Directed Networks."

[51] Muldoon, Bridgeford, and Bassett, "Small-World Propensity and Weighted Brain Networks."

[52] Fagiolo, "Clustering in Complex Directed Networks."

$$E_j^k \equiv \begin{cases} P(\mathbf{a}_k < M_j < \mathbf{b}_k \mid i \text{ recently active}) & \mathit{if} \qquad E_{ij}^{syn} > 0 \\ 0 & \mathit{otherwise} \end{cases}$$

Scaled synaptic topologies

Functional topologies were measured for simulations having typical synaptic weight distributions (n=5) and for simulations where random draws from the synaptic weight distributions were scaled to double strength (n=6). Ratios for global clustering, characteristic path, and smallworldness were quantified following[53], as above, on the two sets of weighted, symmetrized topologies. Directed clustering was measured following[54]. The directed clustering measurements were conducted on binary topologies to control for potential differences stemming from their different underlying mean synaptic weights.

Preparation of Ca^{2+}-dye loaded slices

All procedures were performed in accordance with and approved by the Institutional Animal Care and Use Committee at the University of Chicago. One juvenile mouse (postnatal day 14, of strain C57BL/6) was anesthetized by intraperitoneal injection of ketamine-xylazine and rapidly decapitated. The brain was dissected and placed in oxygenated, ice-cold artificial cerebrospinal fluid (Cut-ACSF; contents contain the following in mM: 3 KCl, 26 NaHCO3, 1 NaH2PO4, 0.5 CaCl2, 3.5 MgSO4 25 dextrose, and 123 sucrose). The brain was then sliced coronally using a vibratome (VT1000S; Leica) into 450 μm thick slices. These slices encompassed the mouse whisker somatosensory cortex. Slices were then transferred into 35°C oxygenated

---

[53] Humphries and Gurney, "Network 'small-World-Ness.'"
[54] Fagiolo, "Clustering in Complex Directed Networks."

incubation fluid (Incu-ACSF; contents contain the following, in mM: 123 NaCl, 3 KCl, 26 NaHCO3, 1 NaH2PO4, 2 CaCl2, 6 MgSO4, 25 dextrose) for 30 min. Bulk loading of Ca^{2+} dye was then performed, via transfer of slices into a Petri dish containing ~2 ml of Incu-ACSF and an aliquot of 50 µg Fura-2AM (Product code, Invitrogen, location) dissolved in 13 µl DMSO and 2 µl of Pluronic F-127 (Code, Invitrogen, location) (following [55]).

Ca^{2+}-imaging procedure

Throughout the duration of imaging, slices were continuously perfused with a standard ACSF solution (contents contain the following, in mM: 123 NaCl, 3 KCl, 26 $NaHCO_3$, 1 NaH_2PO_4, 2 $CaCl_2$, 2 $MgSO_4$, and 25 dextrose, which was continuously aerated with 95% O_2, 5% CO_2). Visualization of Fura-2AM loaded neurons was performed via serial 5 min recordings, collected using the HOPS scanning technique (a suite of software and custom microscopy setup developed in-house, see [56]). This method allowed us to monitor action potential generation within individual neurons, by detecting contours of loaded cells from a raster image, then computing an efficient traveling salesman tour over those cell bodies. Our dwell time parameter was fixed at a value of 16 samples/cell/frame. Population framerate was 20 Hz, resulting in ~450 neurons sampled once every ~50 ms. Changes in emitted fluorescence were analyzed with a threshold-crossing approach. First, a signal-to-noise cutoff was implemented by measuring the ratio of the 99th percentile divided by the mean for the fluorescence trace of each cell. Cells exceeding 1.55 by this metric were retained for further analysis. Of the 444 sampled neurons, 189 exceeded our strict criterion on signal-to-noise (see Methods). Among these cells with clean fluorescent signals, instances of

---

[55] Sadovsky and MacLean, "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity."

[56] Sadovsky et al., "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging."

elevated firing were identified from excursions in the signal exceeding two-sigma, with inflection points more precisely identified by following these excursions backwards to the bin of their most recent median-crossing. The resulting binary vector identified high-probability periods of spiking activity across the imaged population[57].

Inferring connectivity

Recurring timing relationships can be used to identify likely synaptic connections between individual pairs, particularly lagged firing near the timescale of synaptic integration. We used an iterative Bayesian inference algorithm to parse these lagged firing patterns[58]. The inference algorithm was initialized five times, and final weights were pooled as an average. The combined network was thresholded to isolate its strongest relationships. With increasing threshold, functional relationships became more precise in indicating true monosynaptic connectivity, and also more confidently overabundant in the fan-in triangle motif.

To understand the impact of mistaken inferences from a different perspective, independent of relationships between functional weight and true connectivity, bootstrap resampling was used to estimate how errors in inferred connectivity affected estimates of directed clustering measures. For an error rate of 30% estimated from simulated experimental constraints[59], differences in directed clustering were significant even after redacting possible false positives (100 bootstrap-resampled topologies; Figure 8c,e).

---

[57] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry"; Runfeldt, Sadovsky, and MacLean, "Acetylcholine Functionally Reorganizes Neocortical Microcircuits."

[58] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints"; Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."

[59] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

In a typical simulated network, the density of the recruitment network was 0.049, meaning only about one quarter of synaptic connections were a site of propagating activity. Since only those pairs are visible in patterns of lagged firing, the density of recruiting connections was shown for an additional definition of optimal performance (one potentially more appropriate for models with sparse firing).

Validating inferred relationships with cross-correlation

Average cross-correlations were computed over a two-second sliding window using *z*-scored fluorescence traces. The first signal was computed as the product of two putative presynaptic fluorescence traces, as a simple score of their activity and/or coactivity. The second signal was the postsynaptic fluorescence trace. Their raw cross-correlation measures the timing offsets between putative presynaptic activity and postsynaptic firing. The functional relationships used to define fan-in triangle motifs versus simple convergence motifs inferred using iterative Bayesian inference, on the basis of single-frame lagged activity, measured in 50 ms bins.

REFERENCES

Angulo, Marco Tulio, Yang-Yu Liu, and Jean-Jacques Slotine. "Network Motifs Emerge from Interconnections That Favour Stability." *Nature Physics* 11, no. 10 (October 2015): 848–52. doi:10.1038/nphys3402.

Azouz, R., and C. M. Gray. "Cellular Mechanisms Contributing to Response Variability of Cortical Neurons in Vivo." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 19, no. 6 (March 15, 1999): 2209–23.

Barzel, Baruch, and Albert-László Barabási. "Universality in Network Dynamics." *Nature Physics* 9, no. 10 (October 2013): 673–81. doi:10.1038/nphys2741.

Beggs, John M., and Dietmar Plenz. "Neuronal Avalanches in Neocortical Circuits." *The Journal* *of Neuroscience* 23, no. 35 (2003): 11167–11177.

Behabadi, Bardia F., and Bartlett W. Mel. "Mechanisms Underlying Subunit Independence in Pyramidal Neuron Dendrites." *Proceedings of the National Academy of Sciences* 111, no. 1 (2014): 498–503.

Brunel, Nicolas. "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons." *Journal of Computational Neuroscience* 8, no. 3 (2000): 183–208.

Cazé, Romain Daniel, Mark Humphries, and Boris Gutkin. "Passive Dendrites Enable Single Neurons to Compute Linearly Non-Separable Functions," 2013. http://dx.plos.org/10.1371/journal.pcbi.1002867.

Chambers, Brendan, and Jason N. MacLean. "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints." *Journal of Neurophysiology* 114, no. 3 (September 1, 2015): 1837–49. doi:10.1152/jn.00429.2015.

Chialvo, Dante R. "Emergent Complex Neural Dynamics." *Nature Physics* 6, no. 10 (October 2010): 744–50. doi:10.1038/nphys1803.

Crochet, Sylvain, James FA Poulet, Yves Kremer, and Carl CH Petersen. "Synaptic Mechanisms Underlying Sparse Coding of Active Touch." *Neuron* 69, no. 6 (2011): 1160–1175.

Destexhe, Alain, Michael Rudolph, and Denis Paré. "The High-Conductance State of Neocortical Neurons in Vivo." *Nature Reviews Neuroscience* 4, no. 9 (2003): 739–751.

Fagiolo, Giorgio. "Clustering in Complex Directed Networks." *Physical Review E* 76, no. 2 (2007): 26107.

Gasparini, Sonia, and Jeffrey C. Magee. "State-Dependent Dendritic Computation in Hippocampal CA1 Pyramidal Neurons." *The Journal of Neuroscience* 26, no. 7 (2006): 2088–2100.

Gerstein, George L., Donald H. Perkel, and K. N. Subramanian. "Identification of Functionally Related Neural Assemblies." *Brain Research* 140, no. 1 (1978): 43–62.

Goodman, Dan FM, and Romain Brette. "The Brian Simulator." *Frontiers in Neuroscience* 3, no. 2 (2009): 192.

Haider, Bilal, Michael Häusser, and Matteo Carandini. "Inhibition Dominates Sensory Responses in the Awake Cortex." *Nature* 493, no. 7430 (2013): 97–100.

Haldeman, Clayton, and John M. Beggs. "Critical Branching Captures Activity in Living Neural Networks and Maximizes the Number of Metastable States." *Physical Review Letters* 94, no. 5 (2005): 58101.

Harvey, Christopher D., Philip Coen, and David W. Tank. "Choice-Specific Sequences in Parietal

Cortex during a Virtual-Navigation Decision Task." *Nature* 484, no. 7392 (April 5, 2012): 62–68. doi:10.1038/nature10918.

Hlinka, Jaroslav, David Hartman, and Milan Paluš. "Small-World Topology of Functional Connectivity in Randomly Connected Dynamical Systems." *Chaos: An Interdisciplinary Journal of Nonlinear Science* 22, no. 3 (2012): 33107.

Humphries, Mark D., and Kevin Gurney. "Network 'small-World-Ness': A Quantitative Method for Determining Canonical Network Equivalence." *PloS One* 3, no. 4 (2008): e0002051.

Jahnke, Sven, Marc Timme, and Raoul-Martin Memmesheimer. "Guiding Synchrony through Random Networks." *Physical Review X* 2, no. 4 (2012): 41016.

Kruskal, Peter B., Lucy Li, and Jason N. MacLean. "Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex." *Nature Communications* 4 (2013). http://www.nature.com/ncomms/2013/131010/ncomms3574/full/ncomms3574.html.

Kumar, Arvind, Stefan Rotter, and Ad Aertsen. "Spiking Activity Propagation in Neuronal Networks: Reconciling Different Perspectives on Neural Coding." *Nature Reviews Neuroscience* 11, no. 9 (2010): 615–627.

Kumar, Arvind, Sven Schrader, Ad Aertsen, and Stefan Rotter. "The High-Conductance State of Cortical Networks." *Neural Computation* 20, no. 1 (2008): 1–43.

Lefort, Sandrine, Christian Tomm, J.-C. Floyd Sarria, and Carl CH Petersen. "The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex." *Neuron* 61, no. 2 (2009): 301–316.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. "Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations." *Neuron* 62, no. 3 (2009): 413–425.

Luczak, Artur, Bruce L. McNaughton, and Kenneth D. Harris. "Packet-Based Communication in the Cortex." *Nature Reviews Neuroscience*, 2015. http://www.nature.com/nrn/journal/vaop/ncurrent/full/nrn4026.html.

MacLean, Jason N., Brendon O. Watson, Gloster B. Aaron, and Rafael Yuste. "Internal Dynamics Determine the Cortical Response to Thalamic Stimulation." *Neuron* 48, no. 5 (December 8, 2005): 811–23. doi:10.1016/j.neuron.2005.09.035.

Millman, Daniel, Stefan Mihalas, Alfredo Kirkwood, and Ernst Niebur. "Self-Organized Criticality Occurs in Non-Conservative Neuronal Networks during Up States." *Nature Physics* 6, no. 10 (October 2010): 801–5. doi:10.1038/nphys1757.

Miner, Daniel, and Jochen Triesch. "Plasticity-Driven Self-Organization under Topological Constraints Accounts for Non-Random Features of Cortical Synaptic Wiring." *PLOS*

Comput Biol 12, no. 2 (February 11, 2016): e1004759. doi:10.1371/journal.pcbi.1004759.

Motter, Adilson E., Seth A. Myers, Marian Anghel, and Takashi Nishikawa. “Spontaneous Synchrony in Power-Grid Networks.” *Nature Physics* 9, no. 3 (2013): 191–197.

Muldoon, Sarah Feldt, Eric W. Bridgeford, and Danielle S. Bassett. “Small-World Propensity and Weighted Brain Networks.” *Scientific Reports* 6 (February 25, 2016). doi:10.1038/srep22057.

Nigam, Sunny, Masanori Shimono, Shinya Ito, Fang-Chin Yeh, Nicholas Timme, Maxym Myroshnychenko, Christopher C. Lapish, et al. “Rich-Club Organization in Effective Connectivity among Cortical Neurons.” *The Journal of Neuroscience* 36, no. 3 (January 20, 2016): 670–84. doi:10.1523/JNEUROSCI.2177-15.2016.

Orlandi, Javier G., Jordi Soriano, Enrique Alvarez-Lacalle, Sara Teller, and Jaume Casademunt. “Noise Focusing and the Emergence of Coherent Activity in Neuronal Cultures.” *Nature Physics* 9, no. 9 (2013): 582–590.

Ostojic, Srdjan. “Two Types of Asynchronous Activity in Networks of Excitatory and Inhibitory Spiking Neurons.” *Nature Neuroscience* 17, no. 4 (2014): 594–600.

Pajevic, Sinisa, and Dietmar Plenz. “Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches,” 2009. http://dx.plos.org/10.1371/journal.pcbi.1000271.

Paré, D., and E. J. Lang. “Calcium Electrogenesis in Neocortical Pyramidal Neurons in Vivo.” *The European Journal of Neuroscience* 10, no. 10 (October 1998): 3164–70.

Perin, Rodrigo, Thomas K. Berger, and Henry Markram. “A Synaptic Organizing Principle for Cortical Neuronal Groups.” *Proceedings of the National Academy of Sciences* 108, no. 13 (2011): 5419–5424.

Rohden, Martin, Andreas Sorge, Marc Timme, and Dirk Witthaut. “Self-Organized Synchronization in Decentralized Power Grids.” *Physical Review Letters* 109, no. 6 (2012): 64101.

Rothschild, Gideon, Israel Nelken, and Adi Mizrahi. “Functional Organization and Population Dynamics in the Mouse Primary Auditory Cortex.” *Nature Neuroscience* 13, no. 3 (March 2010): 353–60. doi:10.1038/nn.2484.

Runfeldt, Melissa J., Alexander J. Sadovsky, and Jason N. MacLean. “Acetylcholine Functionally Reorganizes Neocortical Microcircuits.” *Journal of Neurophysiology* 112, no. 5 (September 1, 2014): 1205–16. doi:10.1152/jn.00071.2014.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. “Heuristically Optimal Path Scanning for High-Speed

Multiphoton Circuit Imaging." *Journal of Neurophysiology* 106, no. 3 (September 1, 2011): 1591–98. doi:10.1152/jn.00334.2011.

Sadovsky, Alexander J., and Jason N. MacLean. "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity." *The Journal of Neuroscience* 34, no. 23 (2014): 7769–7777.

Sadovsky, Alexander J., and Jason N. MacLean. "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry." *The Journal of Neuroscience* 33, no. 35 (August 28, 2013): 14048–60. doi:10.1523/JNEUROSCI.1977-13.2013.

Sadtler, Patrick T., Kristin M. Quick, Matthew D. Golub, Steven M. Chase, Stephen I. Ryu, Elizabeth C. Tyler-Kabara, Byron M. Yu, and Aaron P. Batista. "Neural Constraints on Learning." *Nature* 512, no. 7515 (August 28, 2014): 423–26. doi:10.1038/nature13665.

Sederberg, Audrey J., Stephanie E. Palmer, and Jason N. MacLean. "Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity." *Journal of Neurophysiology*, February 18, 2015, jn.00885.2014. doi:10.1152/jn.00885.2014.

Shai, Adam S., Costas A. Anastassiou, Matthew E. Larkum, and Christof Koch. "Physiology of Layer 5 Pyramidal Neurons in Mouse Primary Visual Cortex: Coincidence Detection through Bursting." *PLoS Computational Biology* 11, no. 3 (2015): Art–No.

Shen-Orr, Shai S., Ron Milo, Shmoolik Mangan, and Uri Alon. "Network Motifs in the Transcriptional Regulation Network of Escherichia Coli." *Nature Genetics* 31, no. 1 (2002): 64–68.

Sherman, S. Murray. "The Function of Metabotropic Glutamate Receptors in Thalamus and Cortex." *The Neuroscientist*, 2013, 1073858413478490.

Shimono, Masanori, and John M. Beggs. "Functional Clusters, Hubs, and Communities in the Cortical Microconnectome." *Cerebral Cortex* 25, no. 10 (2015): 3743–3757.

Softky, W. R., and C. Koch. "The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs." *The Journal of Neuroscience* 13, no. 1 (January 1, 1993): 334–50.

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B. Chklovskii. "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits." *PLoS Biol* 3, no. 3 (2005): e68.

Teramae, Jun-nosuke, Yasuhiro Tsubo, and Tomoki Fukai. "Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links." *Scientific Reports* 2 (2012). http://www.nature.com/srep/2012/120702/srep00485/full/srep00485.html.

Timme, Nicholas M., Shinya Ito, Maxym Myroshnychenko, Sunny Nigam, Masanori Shimono, Fang-Chin Yeh, Pawel Hottowy, Alan M. Litke, and John M. Beggs. “High-Degree Neurons Feed Cortical Computations.” *PLOS Comput Biol* 12, no. 5 (May 9, 2016): e1004858. doi:10.1371/journal.pcbi.1004858.

Tran-Van-Minh, Alexandra, Romain D. Cazé, Therése Abrahamsson, Laurence Cathala, Boris S. Gutkin, and David A. DiGregorio. “Contribution of Sublinear and Supralinear Dendritic Integration to Neuronal Computations.” *Frontiers in Cellular Neuroscience* 9 (2015): 67. doi:10.3389/fncel.2015.00067.

Watts, Duncan J., and Steven H. Strogatz. “Collective Dynamics of ‘small-World’networks.” *Nature* 393, no. 6684 (1998): 440–442.

Womelsdorf, Thilo, Taufik A. Valiante, Ned T. Sahin, Kai J. Miller, and Paul Tiesinga. “Dynamic Circuit Motifs Underlying Rhythmic Gain Control, Gating and Integration.” *Nature Neuroscience* 17, no. 8 (2014): 1031–1039.

Xu, Ning-long, Mark T. Harnett, Stephen R. Williams, Daniel Huber, Daniel H. O’Connor, Karel Svoboda, and Jeffrey C. Magee. “Nonlinear Dendritic Integration of Sensory and Motor Input during an Active Sensing Task.” *Nature* 492, no. 7428 (December 13, 2012): 247–51. doi:10.1038/nature11601.

Zylberberg, Joel, and Eric Shea-Brown. “Input Nonlinearities Can Shape beyond-Pairwise Correlations and Improve Information Transmission by Neural Populations.” *Physical Review E* 92, no. 6 (December 9, 2015): 62707. doi:10.1103/PhysRevE.92.062707.

CHAPTER 4

General Discussion

The network paradigm is a source of growing interest in the neuroscience community. The key to progress in this area is attention to relationships between the biophysics of individual neurons and emergent routing patterns between them. It is being increasingly recognized that higher-order interactions in synaptic networks are particularly important. These changing priorities are reflected by review pieces in high-profile settings[1]. Chiefly, these articles are outlining challenges and open questions. The results presented in these chapters go beyond these calls to action, furthering our understanding of how synaptic connectivity gives rise to dynamic firing patterns in isolated neocortical tissue. Chiefly, predicting population responses is difficult because individual synaptic connections are weak while postsynaptic integration is non-linear and state-dependent.

One important area for future work is the relationship between short term plasticity and dynamic patterns of functional connectivity. Short-term plasticity imbues synaptic responses with history-dependent efficacies. Glutamatergic connections tend to fall into multiple classes: Class 1A connections with proximal connections, short-term depression, and all-or-none responses, Class 1B connections with short-term depression and graded responses, and Class 2 connections

---

[1] Roudi, Dunn, and Hertz, "Multi-Neuronal Activity and Functional Connectivity in Cell Assemblies"; Benson, Gleich, and Leskovec, "Higher-Order Organization of Complex Networks"; Shiffrin, "Drawing Causal Inference from Big Data"; Hawrylycz et al., "Inferring Cortical Function in the Mouse Visual System through Large-Scale Systems Neuroscience."

characterized by distal dendritic targets and short-term facilitation[2]. Similar issues apply to inhibitory neurons[3]. It is crucial to characterize the distribution of these varying electrophysiological properties in typical cortical populations. In time, it may be possible to predict activity directly on the basis of these diverse, complex lower-level features.

In the meantime, activity mapping is the only tractable method for understanding the consequences of these myriad synaptic mechanisms. Because the brain is complex at multiple scales, it must be investigated at multiple levels of abstraction. Bridging explanations about synaptic mechanisms to their emergent consequences for activity can only be accomplished by studying active populations. Recordings of population activity are caused by numerous diverse interactions, including history-dependent synaptic effects, but the product of those interactions can be surprisingly tractable. An activity-mapping approach can circumvent the need to explicitly deduce response-patterns from their many complex lower-level features.

Should we care about patterns of population responses that evolve over tens of milliseconds? We must, because temporal features at these timescales are impactful for postsynaptic integration. The activity of cortical output neurons, such as those controlling movement through spinal interneurons, is sensitive to temporal features of activity upstream.

ACTIVITY PATTERNS IDENTIFY SELECTIVE SYNAPTIC CONNECTIONS

In Chapter 2, I advanced a methodology for investigating spatiotemporal structure arising in local cortical populations. I focused on recurring lagged firing between pairs of neurons. At the

---

[2] Pasquale and Sherman, "Synaptic Properties of Corticocortical Connections between the Primary and Secondary Visual Cortical Areas in the Mouse"; Covic and Sherman, "Synaptic Properties of Connections between the Primary and Secondary Auditory Cortices in Mice."

[3] Gupta, Wang, and Markram, "Organizing Principles for a Diversity of GABAergic Interneurons and Synapses in the Neocortex."

time that work was begun, it was unknown whether lagged-correlations in activity across pairs of neurons could offer a window into synaptic interactions. In particular, it was feared that sequential firing could give the misleading appearance of synaptic connectivity. In large population recordings, even random firing could give rise to apparent structure, especially for short experimental durations. More optimistically, synaptic connectivity has long been thought to impose lagged correlation structure on pairs of connected neurons[4].

I tested the concern that high-firing rates in a large population might reveal only pseudo-structure rather than genuine synaptic interactions. Neurons in a naturalistic simulation were matched to independent Poisson units; the only difference between these two models was the presence versus absence of synaptic connections. For each trial, firing rate was measured in the simulated neurons. Each Poisson trial was then generated so that Poisson units attained the same distribution of expected firing rates. Thus, rates were matched over the entire 100 s of simulated recording, as well as within each 100 ms trial. This comparison was likely overly conservative, since individual trial firing rates are almost certainly modulated in part by interconnectivity and causal influences. Pairwise lagged-correlations between the Poisson units were significantly weaker and sparser than their counterparts in the simulated dataset. Even after accounting for current limits on experimental duration, measurements of correlation structure reflect causal synaptic interactions.

A second crucial question for interpreting these functional maps: are individual functional relationships indicative of likely synaptic connectivity? Measurements of correlation-structure were thresholded to isolate only the strongest relationships. Functional networks were constructed with edge-weights capturing pairwise measurements. Direct correspondences between functional

---

[4] Gerstein and Perkel, "Simultaneously Recorded Trains of Action Potentials."

relationships and underlying connections were abundant. Stronger functional relationships were more likely to reveal genuine monosynaptic connectivity, rather than some multi-synaptic statistical relationship or chance coupling. Among the synaptic connections revealed through activity mapping with iterative Bayesian inference, strong synaptic connections were disproportionately discovered compared to weaker connections. Strong connections are likely most salient because they can impact postsynaptic integration more independently, where weaker synaptic connections may rely relatively more on distributed coordination, resulting in mutual dependencies not accounted for in the forward model.

False positives certainly remain among the true positive inferred synaptic connections. High correlations can exist in the absence of underlying connectivity, where instead of direct connection, multi-synaptic patterns or input timing gave rise to the measured statistical regularity. For example, divergence is thought to impose misleading common-input correlations[5]. Therefore, inferred synaptic networks must be interpreted cautiously. One day, wide field in vivo recordings will be frozen and processed for electron microscopy—but 1 mm^3 volumes are wildly beyond technical feasibility for the time being. As a tractable complement to anatomical mapping, probabilistic connection mapping is an extremely valuable tool.

Following up on the observation that diverse activity increased the number of detected synaptic connections, I investigated patterns of activity following different input conditions. Inputs were not active during the mapping epoch, but they established initial conditions for the network. Even without considering time, activity vectors differed from one set of random inputs to another. Yet more so than activity vectors alone, the location of pairs cooperating in strong functional

---

[5] Moore et al., “Statistical Signs of Synaptic Interaction in Neurons”; Perkel et al., “Nerve-Impulse Patterns”; Kulkarni and Paninski, “Common-Input Models for Multiple Neural Spike-Train Data”; Stevenson et al., “Inferring Functional Connections between Neurons.”

relationships differed across input conditions. For a given presynaptic neuron, different subsets of its presynaptic field were active across inputs, and to different degrees. Decoder neurons are sensitive to reliable coordination, so dynamic functional relationships would seem to be both informative and salient from the perspective of downstream targets.

Input context modulated functional connectivity, reorganizing patterns of cooperativity in the network. Pairs of neurons have been known to exhibit dynamic correlations of this type for several decades[6]. It is now feasible to extend those findings to larger populations of interacting neurons. In addition to changes in input stimulation, neuromodulatory tone[7] and changes in arousal[8] can also reorganize network responses. Therefore, it may be fruitful to analyze long recordings after segmenting based on context.

HIGHER-ORDER INTERACTIONS COORDINATE RECURRENT DYNAMICS

Using activity mapping as a bridge between emergent responses and underlying synaptic mechanisms, it is important to determine how distributed connections give rise to patterns of activity. Are there patterns of connections and cell identities that function collectively as modules? Functional motifs have the potential to simplify the problem of predicting activity patterns, the presumptive substrate for neocortical computation. Motif counting approaches to understanding complex networks have proven useful for understanding interactions among transcription factors[9], a system similarly characterized by context dependence and non-linearity. Interaction motifs among transcription factors can predict the temporal course of transcript expression, including

---

[6] Aertsen et al., "Dynamics of Neuronal Firing Correlation"; Aertsen and Gerstein, "Dynamic Aspects of Neuronal Cooperativity."

[7] Runfeldt, Sadovsky, and MacLean, "Acetylcholine Functionally Reorganizes Neocortical Microcircuits."

[8] Ecker et al., "State Dependence of Noise Correlations in Macaque Primary Visual Cortex."

[9] Milo et al., "Network Motifs."

non-monotonic changes in expression[10]. Of course, structure-function relationships based on gene regulatory networks should not be interpreted to apply directly to neocortical populations, to the collective mechanisms of synaptic integration.

Neuroscientists have often focused on the functional motif of *convergence*. This definition is usually agnostic to the presence of connections between the presynaptic sources. In Chapter 3, I argued that shared connections among presynaptic neurons converging on a common target generate different dynamics than in the absence of shared connections. Additional connections act to coordinate collective firing, increasing the likelihood of postsynaptic recruitment.

Since shared connections increase the probability of postsynaptic recruitment, t may be instructive to investigate whether presynaptic neurons involved in convergence also share connections among themselves. Trans-synaptic retrograde methods such as rabies pseudo-virus[11] would be well-suited to addressing clustering in the presynaptic field.

This thesis emphasized local connection patterns, within interconnected populations at 1 mm^3 scales. By virtue of distance-dependent connectivity rules, these scales are likely to exhibit the highest levels of multi-neuron coordination. But long-range connections are still subject to constraints on synaptic integration, and long-range inputs are likely to be more effective when they share clustered connections upstream. Neocortical areas known to be anatomically connected exhibit dynamic coupling and uncoupling[12] in the form of varying coherence [13]. These phenomena are likely to reflect the success versus failure of long-range inputs to recruit their non-local

---

[10] Kaplan et al., "The Incoherent Feed-Forward Loop Can Generate Non-Monotonic Input Functions for Genes."

[11] Wall et al., "Monosynaptic Circuit Tracing in Vivo through Cre-Dependent Targeting and Complementation of Modified Rabies Virus"; Callaway and Luo, "Monosynaptic Circuit Tracing with Glycoprotein-Deleted Rabies Viruses."

[12] Gregoriou et al., "High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention"; Battaglia et al., "Dynamic Effective Connectivity of Inter-Areal Brain Circuits."

[13] Ioannides, "Dynamic Functional Connectivity."

neighbors. Activation of clustered long-range projection neurons is likely to be more efficacious than non-clustered projection neurons, so that the same individual connection could participate in coherent versus incoherent activity with respect to its long-range target.

Because clustered fan-in triangles are characterized by coordinated activity and low firing rates, they meet exactly the criteria useful for a reader/decoder neuron downstream. The relative synchronization across the fan-in cluster is salient from the perspective of the postsynaptic membrane; moreover, the stability of the timing relationships among fan-in clustered neurons could potentially be exploited by active processing in dendritic arbors to increase the reliability of decoding. I found that reference neurons characterized by high fan-in clustering also tended to have low firing rates, requiring coordinated activity across the clustered neighborhood for their recruitment. Because they fire so sparsely and reliably reflect activity upstream, these neurons are likely to be very informative from an information theoretic perspective.

INTERCONNECTED NEURONS ARE A COMPLEX SYSTEM

One reason neocortical computation has been intractable is that patterns of firing are hard to predict from connectivity alone—especially incomplete connectivity. Firing patterns emerge from numerous interactions among individually weak synaptic connections and depend further on myriad context-dependent cellular responses. Traditional investigations have attempted to manage this complexity by studying single connections while the rest of the network is quiet—but firing is qualitatively different during naturalistic modes of operation. Single cells and single connections are not the appropriate level of abstraction for understanding how cellular populations implement computation and behavior.

Rather than try to predict firing from the bottom-up, I took the approach of mapping propagation from the top down—in an active network. This is the layer of abstraction appropriate for understanding cortical computation: reliable relationships between input activity and output firing, separated in time by the propagation of activity through numerous intermediaries.

Sampling activity near the temporal resolution of synaptic integration provided me with a different kind of dataset—one with rich information about relative timing. Patterns of lagged firing times were abundant in recorded data relative to chance—out of the unpredictable complexity of synaptic interactions, organized activity emerged. After mapping stereotyped sequential firing between pairs of neurons, activity in the neuronal population could be seen as a *functional network* whose nodes were neurons and whose edge weights quantified frequent lagged firing. This functional network is informative of the underlying synaptic network.

But patterns of propagating activity are not identical copies of synaptic networks. In particular, traffic on the synaptic network can be selective and context dependent. This is an important issue for interpreting experimental data, be they connectomic reconstructions or population recordings. To understand how connectivity and traffic patterns differ, it is necessary to know ground-truth features of synaptic interactions. This omniscient perspective is best achieved through network simulation.

Simulated synaptic networks were completely random, yet emergent statistical structure was found in traffic patterns. In particular, maps of propagating activity were characterized by small-world characteristics[14]. This result echoed other reports of small-world functional structure[15] i.e. elevated neighbor-of-neighbor interactions (global clustering coefficient) and short mean path

---

[14] Watts and Strogatz, "Collective Dynamics of 'small-World'networks."

[15] Pajevic and Plenz, "The Organization of Strong Links in Complex Networks"; Achard et al., "A Resilient, Low-Frequency, Small-World Human Brain Functional Network with Highly Connected Association Cortical Hubs."

lengths. I trace this small-world structure to a dominant underlying interaction at the mechanistic level of synaptic integration: clustered fan-in triangles[16] (Figure 2A). Despite occurring only at chance levels in the synaptic network (Figure 2B), fan-in triangles were abundant in patterns of traffic (Figure 2C). Selective, asymmetric elevation of this motif within propagating activity reflects higher-than-pairwise synaptic interactions.

I show that fan-in clustering was associated with emergent coordination of firing times (Figure 3A), supplying strong coincident input (Figure 3B), leading to effective postsynaptic integration (Figure 3C). Compared to other connection schemes, clustered fan-in triangles actively synchronize input timing. This synchronous input benefits from a non-linear increase in efficacy because of the biophysics of individual neurons. I discuss how fan-in clustered networks could self-organize by stabilizing local features of a less structured (or even random) developing synaptic network.

I went on to measure directed clustering in empirically recorded populations (Figure 4A), using the experimental approach described previously. Fan-in clustering was strongly over-represented in these patterns of firing (Figure 4B,4C,4D). Neurons in fan-in triangles tended to be more temporally coordinated than neurons in simple-convergence (those lacking neighbor-connectivity) (Figure 4E). It is important to note that neocortical networks are known to be highly non-random15,16, underscoring the generality of fan-in clustering for active networks of neurons. I have suggested that clustered fan-in triangles are a generic neocortical primitive for implementing reliable mapping of inputs onto outputs20.

The synaptic networks linking neurons in the mammalian neocortex are the epitome of a complex system, but they do differ from some well-studied networks in an important way: their

---

[16] Fagiolo, “Clustering in Complex Directed Networks.”

nodes are complex information processing units themselves with rich internal workings and state-dependent properties. Thus, results for e.g. electrical circuits of resistors will usually not generalize to neuronal systems. Effectively utilizing network analyses in neuroscience benefits from recognizing that nodes are non-linear integrators of input, sensitive to synchrony and recent history.

A second counter-intuitive feature of cortical networks is that synaptic integration is often slow compared to synaptic transmission. Most (excitatory) neocortical neurons fire sparsely, requiring either sustained or coordinated input to become active. Neurons typically integrate inputs from the last 3 – 25 ms, depending on the neurotransmitter milieu. Yet traversing local connections is very fast, requiring ~1 ms for neurotransmitter release and passive diffusion across the synaptic cleft, plus ~0.5 ms of electrical conduction time in the axon, to traverse the 1 mm field of view described in the experiments above. As a result, interpretations based on traversing edges without accounting for active processing at constituent nodes should be regarded cautiously.

On the relationship between synaptic networks and the structured propagating activity that flows through them: The presence of a single connection can have myriad consequences for spike transmission, depending on the context of other recent activity and other patterns of connections. Out of these many intricate components, organized firing patterns emerge.

FUTURE DIRECTIONS

Feedforward models of propagating activity

Regularities in connectivity across *layers* and between *cortical areas* define hierarchies and recurrent loops of synaptic influence. These features underlie definitions of *input* and *output* neurons and describe features at the scale of millimeters to centimeters. They have proven insufficient for predicting structured activity at the 50 um – 1 mm scale, the scale of the biophysics governing synaptic recruitment, because synaptic interactions are sparse, weak and variable individually [17]. Understanding the choreography of synaptic recruitment is necessary for identifying the mapping functions linking input activity, internal state, and output activity. Neocortical computation is implemented through these mapping functions. My PhD work has focused on isolated cortical networks, investigating their patterned dynamics in the absence of structured external input. How are these intrinsic network responses modulated and controlled by input activity?

Synfire chains have inspired thinking on this question [18], but they organize unrealistic levels of synchronized activity and millisecond temporal precision not reproduced in experimental data. Given their biological implausibility, synfire chains are not a realistic model for propagating activity in cortical networks. Synfire chains are also not realistic in their connection topologies, connected via feedforward connections only. Updated models known as Feed-Forward Networks (FFNs) are the subject of ongoing study [19]. Like synfire chains, FFNs can produce unrealistically high levels of synchrony, even from asynchronous input [20]. In an in vitro model of propagating activity, trained into the network using electrical stimulation, rate-coded inputs were transformed into synchrony-coded outputs [21]. It is unknown whether these response features characterize in vivo activity. Problems of excessive synchronization can be lessened by the presence of membrane

---

[17] Softky and Koch, “The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs.”

[18] Abeles, *Corticonics*.

[19] Diesmann, Gewaltig, and Aertsen, “Stable Propagation of Synchronous Spiking in Cortical Neural Networks.”

[20] Litvak et al., “On the Transmission of Rate Code in Long Feedforward Networks with Excitatory–Inhibitory Balance.”

[21] Reyes, “Synchrony-Dependent Propagation of Firing Rate in Iteratively Constructed Networks in Vitro.”

fluctuations [22], illustrating how function in active networks can be fundamentally different than for quiescent networks. Ultimately, low firing rates and asynchronous background can support dual propagation of rate and synchrony [23].

Altering the synaptic structure of these feedforward networks has consequential impacts for their resulting dynamics. For example, strength of connections are intimately related to stability of rate, synchrony, or both dynamic features in propagating activity [24]. Strength of connectivity also impacts asynchronous dynamics [25], which may underlie its different consequences for propagating activity. Similarly, it is natural to expect that connection topology will have major consequences for propagation of activity. Recurrent connectivity can be viewed as embedded with feedforward networks, providing a bridge from FFNs to the work in this thesis. Study of FFNs has provided a set of analytical tools for quantifying temporal features in cascades of activity. However, this approach has underemphasized the importance of connection patterns. In future work, I aim to extend these tools to study how non-random connection patterns impact the dynamics of propagating activity in recurrent networks.

Population dynamics in non-random network topologies

Understanding how heterogeneous connectivity features affect the propagation of activity is an important future direction for computational neuroscience. I am approaching this problem through simulation, with the goal of generating diverse balanced networks with naturalistic

---

[22] Rossum, Turrigiano, and Nelson, “Fast Propagation of Firing Rates through Layered Networks of Noisy Neurons”; Vogels and Abbott, “Signal Propagation and Logic Gating in Networks of Integrate-and-Fire Neurons”; Moldakarimov, Bazhenov, and Sejnowski, “Feedback Stabilizes Propagation of Synchronous Spiking in Cortical Neural Networks.”

[23] Kumar, Rotter, and Aertsen, “Conditions for Propagating Synchronous Spiking and Asynchronous Firing Rates in a Cortical Network Model.”

[24] Kumar, Rotter, and Aertsen, “Spiking Activity Propagation in Neuronal Networks.”

[25] Ostojic, “Two Types of Asynchronous Activity in Networks of Excitatory and Inhibitory Spiking Neurons.”

dynamics. By comparing topological features across this family of models, I hope to characterize constraints on connectivity necessary for asynchronous irregular firing.

This problem has remained unsolved because the technical challenges involved in designing network models with multiple degrees of freedom are significant. Automatic generation of network models remains difficult, with most successful naturalistic models being tuned laboriously and artfully by computational researchers. This approach has led to important progress, but it is slow. More significantly, comparisons between models with different parameters are difficult to interpret in isolation. To contextualize differences in feasible parameters across models, I aim to generate a large ensemble of feasible network architectures consistent with experimental evidence about network connection density and synaptic strength.

Framing simulation design as an optimization problem is attractive because it offers an empirical judgment for unconstrained design choices, rather than relying solely on scientific intuition. Other design choices are simpler, guided by the statistics of available data. For the difficult choices, stochastic optimization offers an efficient search strategy for relating design parameters to emergent dynamics. Unlike traditional optimization methods, gradients are not explicitly calculated, so objective functions need not be differentiable. This is advantageous for working with complex neuron models, where there is no closed-form analytical solution to predict their collective dynamics[26]. Indeed, that is the very reason why simulation is necessary for these problems. This approach has previously been applied successfully to single-neuron models with multiple compartments[27].

---

[26] Baptista et al., "Chaotic, Informational and Synchronous Behaviour of Multiplex Networks."

[27] Vanier and Bower, "A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models."

Over uncertain parameters, a traditional method for evaluating design choices was the *sensitivity test*. For small numbers of parameters, small perturbations could sometimes be conducted in every degree of freedom. Models with higher dimensionality are not amenable to this brute force methodology. An alternative approach to sensitivity testing has been to explore the effect of small perturbations along directions defined by eigenvectors[28]. However, for parameter landscapes with strong non-linear interdependencies, this approach may only be informative about extremely local regions, before the assumption of linearity leads to large errors in predicted outcomes. Stochastic optimization approaches offer a more comprehensive analysis of sensitivity, combining breadth-wise sampling over the entirety of parameter with depth-wise local searches, parallelized over multiple agents.

Multiple powerful stochastic optimization approaches are being actively investigated, and new insights are being gained for choosing among them. For discrete optimization problems, e.g. where design parameters are limited to integers, evolutionary algorithms excel[29]. Evolutionary algorithms encode parameter values in string-form, for a population of putative models. At each generation, putative models replicate with mutation and sexual selection, and only the best performing designs survive to the next generation. Evolutionary algorithms were inspired by biology, and they are a remarkable example of how computation in nature can inspire human engineering.

---

[28] Fisher et al., "A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit."

[29] Deb et al., "A Fast and Elitist Multiobjective Genetic Algorithm."

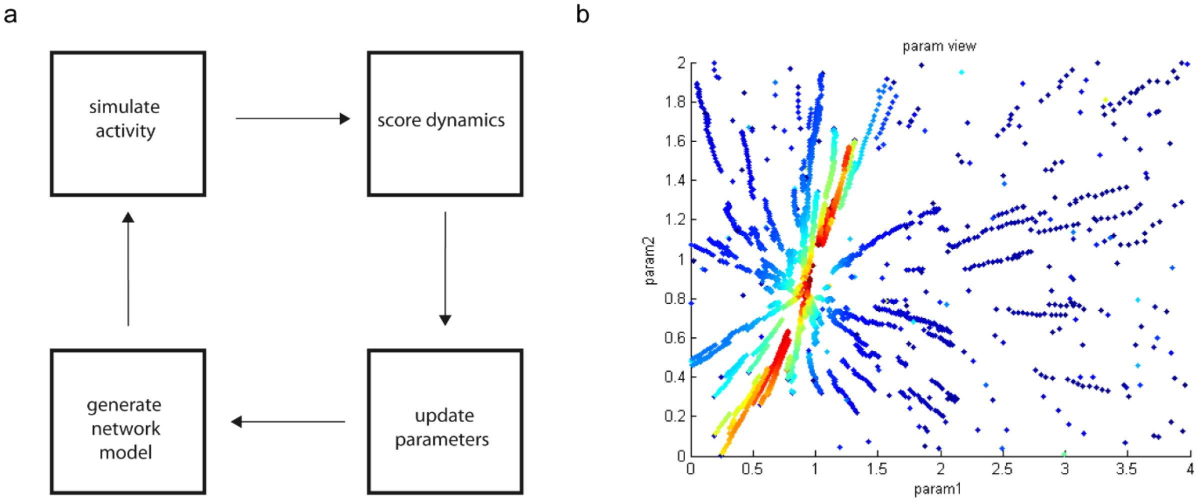

Figure 4-1. Optimization with stochastic Firefly search.

(a) Activity is simulated for a population of initial models, and scored based on dynamical realism. The population competes for influence. A new generation of models are generated with increasing bias towards good performers from the previous generation.
(b) Performance of the modified Firefly algorithm on a benchmark problem. A population of units with x and y parameters search a multi-dimensional objective space defined by the Rosenbrock function. The global optimum is successfully found, and high-scoring regions of parameter space are better explored than low-scoring regions.

Other stochastic optimization algorithms also tend to be based loosely on natural phenomena. In simulated annealing, parameter space is seeded randomly and subsequent thermal fluctuations modulate local search radius[30]. Particle swarm optimization operates similarly, with a velocity vector added to each putative model design[31]. On each generation, particles adjust their velocity towards the most successful model design. Particle swarm can outperform evolutionary algorithms in terms of convergence speed, but they are best suited for continuous optimization problems. Building on particle swarm approach, the *firefly algorithm* imagines that a given particle

---

[30] Hwang, "Simulated Annealing."
[31] Bonyadi and Michalewicz, "Particle Swarm Optimization for Single Objective Continuous Space Problems"; Elhossini, Areibi, and Dony, "Strength Pareto Particle Swarm Optimization and Hybrid EA-PSO for Multi-Objective Optimization."

is attracted not just to the single best design, but to all successful designs in its sightline[32]. The firefly algorithm is not a good model for real firefly mating behavior, but it is an excellent abstraction for high performance optimization, and firefly searches can outperform particle swarm[33].

Timescales of information processing

Appropriate choice of spatial and temporal granularity in neuroscience can differ from one scientific question to another. Then what is the appropriate temporal scale to analyze causation in populations of neurons? This is an important question, because fast sampling has a cost. Experimental realities impose constrained choices like the trade-off between large sample sizes and high temporal resolution[34]. Discussions about appropriate temporal resolution are also relevant for designing indicators of spiking activity, because fast onset of fluorescent response can be in conflict with large changes in signal magnitude, as for the genetically encoded Ca^{2+}-indicators GCaMP6f and GCaMP6s[35].

First intuition might suggest that synaptic delays dominate the statistics of propagating activity. In some dynamic regimes and model systems, including hyperactive cultures, epilepsy models, and the synfire chain framework, delays between synaptically connected neurons do tend to be dominated by the biophysics of vesicle fusion and transmitter diffusion. However, wakeful activity in vivo is sparse, asynchronous, and irregular, and inter-spike intervals are highly variable. Delays between synaptically connected neurons are much longer on average than the 1-2 ms

---

[32] Fister et al., “A Comprehensive Review of Firefly Algorithms.”

[33] Yang, “Firefly Algorithms for Multimodal Optimization.”

[34] Sadovsky et al., “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.”

[35] Chen et al., “Ultrasensitive Fluorescent Proteins for Imaging Neuronal Activity.”

required for synaptic function. Computation in single neurons maps input patterns and recent history onto output spiking, and sparse firing is necessary for informative spike trains. Populations dominated by synaptic delays are too much like transmission wires, and synaptic weights that are unrealistically strong underestimate the importance of synaptic integration. This is an important issue for designing realistic network models. On the other hand, rate codes integrating spike-counts over 100 - 1000 ms are far too slow to meet the single-trial computational needs of cortical circuits. These integration times are also precluded by the leaky membrane of real cortical neurons.

The most powerful timescale for understanding network computation is set by the timescale of synaptic integration. Integration in neurons tends to operate at timescales between 3 and 20 ms, depending on conductance state and presynaptic firing rate, as well as probable interactions with neuromodulatory tone, modulatory glutamatergic inputs, and ion channel expression profile. That is not to say that membrane time constants necessarily dominate the distribution of lags between connected neurons, either: timing interrelationships emerge through the interplay of membrane dynamics with impinging currents. Causal influences will operate over different timescales for neurons communicating via metabotropic glutamate receptors than for neurons communicating via ionotropic glutamate. Of course, during normal function, these multiple receptor-systems are active as a collective. Choice of sampling frequency should follow the interplay between currents of interest and the biophysics of the neuronal membrane. These are the considerations that need to inform experimental design, when temporal resolution is a factor of interest. For the case of information processing with synaptic networks, it is important that researchers attend to these timescales, recognizing that they are central to neuronal targets downstream.

REFERENCES

Abeles, M. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1991.

Achard, Sophie, Raymond Salvador, Brandon Whitcher, John Suckling, and E. D. Bullmore. "A Resilient, Low-Frequency, Small-World Human Brain Functional Network with Highly Connected Association Cortical Hubs." *The Journal of Neuroscience* 26, no. 1 (2006): 63–72.

Aertsen, A. M., G. L. Gerstein, M. K. Habib, and G. Palm. "Dynamics of Neuronal Firing Correlation: Modulation Of 'effective Connectivity.'" *Journal of Neurophysiology* 61, no. 5 (May 1989): 900–917.

Aertsen, Ad M. H. J., and George L. Gerstein. "Dynamic Aspects of Neuronal Cooperativity: Fast Stimulus-Locked Modulations of Effective Connectivity." In *Neuronal Cooperativity*, edited by Professor Dr Jürgen Krüger, 52–67. Springer Series in Synergetics 49. Springer Berlin Heidelberg, 1991. http://link.springer.com/chapter/10.1007/978-3-642-84301-3_4.

Baptista, M. S., R. M. Szmoski, R. F. Pereira, and S. E. de Souza Pinto. "Chaotic, Informational and Synchronous Behaviour of Multiplex Networks." *Scientific Reports* 6 (March 4, 2016). doi:10.1038/srep22617.

Battaglia, Demian, Annette Witt, Fred Wolf, and Theo Geisel. "Dynamic Effective Connectivity of Inter-Areal Brain Circuits." *PLoS Comput Biol* 8, no. 3 (2012): e1002438.

Benson, Austin R., David F. Gleich, and Jure Leskovec. "Higher-Order Organization of Complex Networks." *Science* 353, no. 6295 (July 8, 2016): 163–66. doi:10.1126/science.aad9029.

Bonyadi, Mohammad Reza, and Zbigniew Michalewicz. "Particle Swarm Optimization for Single Objective Continuous Space Problems: A Review." *Evolutionary Computation*, 2016. http://www.mitpressjournals.org/doi/abs/10.1162/EVCO_r_00180.

Callaway, Edward M., and Liqun Luo. "Monosynaptic Circuit Tracing with Glycoprotein-Deleted Rabies Viruses." *The Journal of Neuroscience* 35, no. 24 (June 17, 2015): 8979–85. doi:10.1523/JNEUROSCI.0409-15.2015.

Chen, Tsai-Wen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan, Eric R. Schreiter, et al. "Ultrasensitive Fluorescent Proteins for Imaging Neuronal Activity." *Nature* 499, no. 7458 (July 18, 2013): 295–300. doi:10.1038/nature12354.

Covic, Elise N., and S. Murray Sherman. “Synaptic Properties of Connections between the Primary and Secondary Auditory Cortices in Mice.” *Cerebral Cortex*, March 8, 2011, bhr029. doi:10.1093/cercor/bhr029.

Deb, K., A. Pratap, S. Agarwal, and T. Meyarivan. “A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II.” *IEEE Transactions on Evolutionary Computation* 6, no. 2 (April 2002): 182–97. doi:10.1109/4235.996017.

Diesmann, Markus, Marc-Oliver Gewaltig, and Ad Aertsen. “Stable Propagation of Synchronous Spiking in Cortical Neural Networks.” *Nature* 402, no. 6761 (December 2, 1999): 529–33. doi:10.1038/990101.

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. “State Dependence of Noise Correlations in Macaque Primary Visual Cortex.” *Neuron* 82, no. 1 (April 2, 2014): 235–48. doi:10.1016/j.neuron.2014.02.006.

Elhossini, Ahmed, Shawki Areibi, and Robert Dony. “Strength Pareto Particle Swarm Optimization and Hybrid EA-PSO for Multi-Objective Optimization.” *Evolutionary Computation* 18, no. 1 (2010): 127–156.

Fagiolo, Giorgio. “Clustering in Complex Directed Networks.” *Physical Review E* 76, no. 2 (2007): 26107.

Fisher, Dimitry, Itsaso Olasagasti, David W. Tank, Emre R. F. Aksay, and Mark S. Goldman. “A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit.” *Neuron* 79, no. 5 (September 4, 2013): 987–1000. doi:10.1016/j.neuron.2013.06.041.

Fister, Iztok, Iztok Fister Jr., Xin-She Yang, and Janez Brest. “A Comprehensive Review of Firefly Algorithms.” *Swarm and Evolutionary Computation* 13 (December 2013): 34–46. doi:10.1016/j.swevo.2013.06.001.

Gerstein, G. L., and D. H. Perkel. “Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation.” *Science (New York, N.Y.)* 164, no. 3881 (May 16, 1969): 828–30.

Gregoriou, Georgia G., Stephen J. Gotts, Huihui Zhou, and Robert Desimone. “High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention.” *Science* 324, no. 5931 (May 29, 2009): 1207–10. doi:10.1126/science.1171402.

Gupta, Anirudh, Yun Wang, and Henry Markram. “Organizing Principles for a Diversity of GABAergic Interneurons and Synapses in the Neocortex.” *Science* 287, no. 5451 (January 14, 2000): 273–78. doi:10.1126/science.287.5451.273.

Hawrylycz, Michael, Costas Anastassiou, Anton Arkhipov, Jim Berg, Michael Buice, Nicholas Cain, Nathan W. Gouwens, et al. “Inferring Cortical Function in the Mouse Visual System through Large-Scale Systems Neuroscience.” *Proceedings of the National Academy of Sciences* 113, no. 27 (July 5, 2016): 7337–44. doi:10.1073/pnas.1512901113.

Hwang, Chii-Ruey. “Simulated Annealing: Theory and Applications.” *Acta Applicandae Mathematica* 12, no. 1 (1988): 108–11. doi:10.1007/BF00047572.

Ioannides, Andreas A. “Dynamic Functional Connectivity.” *Current Opinion in Neurobiology*, Cognitive neuroscience, 17, no. 2 (April 2007): 161–70. doi:10.1016/j.conb.2007.03.008.

Kaplan, Shai, Anat Bren, Erez Dekel, and Uri Alon. “The Incoherent Feed-Forward Loop Can Generate Non-Monotonic Input Functions for Genes.” *Molecular Systems Biology* 4 (July 15, 2008). doi:10.1038/msb.2008.43.

Kulkarni, Jayant E., and Liam Paninski. “Common-Input Models for Multiple Neural Spike-Train Data.” *Network: Computation in Neural Systems* 18, no. 4 (January 1, 2007): 375–407. doi:10.1080/09548980701625173.

Kumar, Arvind, Stefan Rotter, and Ad Aertsen. “Conditions for Propagating Synchronous Spiking and Asynchronous Firing Rates in a Cortical Network Model.” *The Journal of Neuroscience* 28, no. 20 (May 14, 2008): 5268–80. doi:10.1523/JNEUROSCI.2542-07.2008.

Kumar, Arvind, Stefan Rotter, and Ad Aertsen. “Spiking Activity Propagation in Neuronal Networks: Reconciling Different Perspectives on Neural Coding.” *Nature Reviews Neuroscience* 11, no. 9 (2010): 615–627.

Litvak, Vladimir, Haim Sompolinsky, Idan Segev, and Moshe Abeles. “On the Transmission of Rate Code in Long Feedforward Networks with Excitatory–Inhibitory Balance.” *The Journal of Neuroscience* 23, no. 7 (April 1, 2003): 3006–15.

Milo, Ron, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. “Network Motifs: Simple Building Blocks of Complex Networks.” *Science* 298, no. 5594 (2002): 824–827.

Moldakarimov, Samat, Maxim Bazhenov, and Terrence J. Sejnowski. “Feedback Stabilizes Propagation of Synchronous Spiking in Cortical Neural Networks.” *Proceedings of the National Academy of Sciences* 112, no. 8 (February 24, 2015): 2545–50. doi:10.1073/pnas.1500643112.

Moore, George P., Jose P. Segundo, Donald H. Perkel, and Herbert Levitan. "Statistical Signs of Synaptic Interaction in Neurons." *Biophysical Journal* 10, no. 9 (September 1970): 876–900.

Ostojic, Srdjan. "Two Types of Asynchronous Activity in Networks of Excitatory and Inhibitory Spiking Neurons." *Nature Neuroscience* 17, no. 4 (2014): 594–600.

Pajevic, Sinisa, and Dietmar Plenz. "The Organization of Strong Links in Complex Networks." *Nature Physics* 8, no. 5 (2012): 429–436.

Pasquale, Roberto De, and S. Murray Sherman. "Synaptic Properties of Corticocortical Connections between the Primary and Secondary Visual Cortical Areas in the Mouse." *The Journal of Neuroscience* 31, no. 46 (November 16, 2011): 16494–506. doi:10.1523/JNEUROSCI.3664-11.2011.

Perkel, Donald H., George L. Gerstein, Mark S. Smith, and William G. Tatton. "Nerve-Impulse Patterns: A Quantitative Display Technique for Three Neurons." *Brain Research* 100, no. 2 (December 19, 1975): 271–96. doi:10.1016/0006-8993(75)90483-7.

Reyes, Alex D. "Synchrony-Dependent Propagation of Firing Rate in Iteratively Constructed Networks in Vitro." *Nature Neuroscience* 6, no. 6 (June 2003): 593–99. doi:10.1038/nn1056.

Rossum, Mark C. W. van, Gina G. Turrigiano, and Sacha B. Nelson. "Fast Propagation of Firing Rates through Layered Networks of Noisy Neurons." *The Journal of Neuroscience* 22, no. 5 (March 1, 2002): 1956–66.

Roudi, Yasser, Benjamin Dunn, and John Hertz. "Multi-Neuronal Activity and Functional Connectivity in Cell Assemblies." *Current Opinion in Neurobiology*, Large-Scale Recording Technology (32), 32 (June 2015): 38–44. doi:10.1016/j.conb.2014.10.011.

Runfeldt, Melissa J., Alexander J. Sadovsky, and Jason N. MacLean. "Acetylcholine Functionally Reorganizes Neocortical Microcircuits." *Journal of Neurophysiology* 112, no. 5 (September 1, 2014): 1205–16. doi:10.1152/jn.00071.2014.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging." *Journal of Neurophysiology* 106, no. 3 (September 1, 2011): 1591–98. doi:10.1152/jn.00334.2011.

Shiffrin, Richard M. "Drawing Causal Inference from Big Data." *Proceedings of the National Academy of Sciences* 113, no. 27 (July 5, 2016): 7308–9. doi:10.1073/pnas.1608845113.

Softky, W. R., and C. Koch. “The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs.” *The Journal of Neuroscience* 13, no. 1 (January 1, 1993): 334–50.

Stevenson, Ian H, James M Rebesco, Lee E Miller, and Konrad P Körding. “Inferring Functional Connections between Neurons.” *Current Opinion in Neurobiology* 18, no. 6 (December 2008): 582–88. doi:10.1016/j.conb.2008.11.005.

Vanier, Michael C., and James M. Bower. “A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models.” *Journal of Computational Neuroscience* 7, no. 2 (1999): 149–71. doi:10.1023/A:1008972005316.

Vogels, Tim P., and L. F. Abbott. “Signal Propagation and Logic Gating in Networks of Integrate-and-Fire Neurons.” *The Journal of Neuroscience* 25, no. 46 (November 16, 2005): 10786–95. doi:10.1523/JNEUROSCI.3508-05.2005.

Wall, Nicholas R., Ian R. Wickersham, Ali Cetin, Mauricio De La Parra, and Edward M. Callaway. “Monosynaptic Circuit Tracing in Vivo through Cre-Dependent Targeting and Complementation of Modified Rabies Virus.” *Proceedings of the National Academy of Sciences* 107, no. 50 (December 14, 2010): 21848–53. doi:10.1073/pnas.1011756107.

Watts, Duncan J., and Steven H. Strogatz. “Collective Dynamics of ‘small-World’networks.” *Nature* 393, no. 6684 (1998): 440–442.

Yang, Xin-She. “Firefly Algorithms for Multimodal Optimization.” In *Stochastic Algorithms: Foundations and Applications*, edited by Osamu Watanabe and Thomas Zeugmann, 169–78. Lecture Notes in Computer Science 5792. Springer Berlin Heidelberg, 2009. http://link.springer.com/chapter/10.1007/978-3-642-04944-6_14.