THE UNIVERSITY OF CHICAGO

NETWORK REPRESENTATION OF STIMULI IN MURINE VISUAL CORTEX

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

MAAYAN LEVY

CHICAGO, ILLINOIS

AUGUST 2021

COPYRIGHT

TABLE OF CONTENTS

List of figures v

List of tables vii

Acknowledgements viii

Abstract ix

Introduction 1

Chapter I – Ensemble stacking mitigates biases in inference of synaptic connectivity

1.1 Abstract 25

1.2 Introduction 26

1.3 Methods 29

1.4 Results 38

1.5 Discussion 59

1.6 References 67

Chapter II – Network analysis of murine cortical dynamics implicates untuned neurons in visual stimulus coding

2.1 Abstract 73

2.2 Introduction 74

2.3 Materials and Methods 76

2.4 Results 85

2.5 Discussion 104

2.6 References 108

2.7 Supplementary figures 113

Chapter III – Temporal pairwise correlations reveal sparse subsets of highly variable spikes the code for the visual stimulus

3.1 Abstract 117

3.2 Introduction 118

3.3 Methods 121

3.4 Results 128

3.5 Discussion 139

3.6 References 143

3.7 Supplementary figures 149

Discussion 151

References 167

LIST OF FIGURES

1.1 Inferring synaptic connectivity from pairwise spike timing 34

1.2 Leveraging anti-correlations to isolate excitatory connections 40

1.3 Removal of mean background timing relationships improves detection of synaptic pairs 44

1.4 Rescaling neuron-wise residual score variance 45

1.5 Survival curves for $norm\_residual(count)_{ij}$, $norm\_residual(cMI)_{ij}$ and $norm\_residual(TE1)_{ij}$ for multiple time resolutions and accuracy rates 47

1.6 Comparing gains in the regularization pipeline 48

1.7 Inference algorithms differ in their statistical preferences 54

1.8 Pooling individual measures to generate Ensemble predictions 56

1.9 Ensemble weightings generalize across simulated datasets and out-perform individual measures 58

2.1 Overlap between the most active cells in different stimuli 87

2.2 Constructing a network comprising two functional classes 89

2.3 Functional networks inferred from fluorescence contain stimulus features 93

2.4 Functional networks are stimulus-specific 95

2.5 Rich club structure in functional networks 97

2.6 Random walks on stimulus networks tend to converge onto untuned neurons 100

2.7 Two-phase model of V1 decodes accurately from both tuned and untuned neurons 103

2.S1 Tuned and untuned neurons activate with different latencies 113

2.S2 In- and out- degrees of neurons in sub-FNs are similar for stimulus and gray epochs 114

2.S3 Tuned neurons show slightly larger relative rankings in gray FNs 114

2.S4 Permuting in-edges of tuned neurons causes a decrease in untuned neurons ranking 115

2.S5 Training the decoder only with one functional sub-class results in similar degradation in performance 116

3.1 Functional networks (FNs) can be divided to 4 sub-FNs based on edge statistics 130

3.2 Sub-FNs have different hallmarks of real-time edge expression 133

3.3 Sub-FN consistent spikes are ultra-sparse and highly variable 135

3.4 Information quantity in subsets of sub-FN consistent spikes 137

3.5 Decoding the direction of drifting gratings from sub-FN consistent spikes 138

3.S1 Unique sub-types display less reciprocal edges 149

3.S2 Unique sub-FNs do not include more tuned-tuned edges 150

3.S3 Recurrent and feedforward pattern counts over time 150

4.1 F shared sub-FNs are enriched in triplet motifs 161

LIST OF TABLES

1.1 Parameters of the spiking network model 30

1.2 Percent of connections retained in the recruitment network 35

1.3 Number of spikes used for inference 39

1.4 Ensemble weights 52

ACKNOWLEDGMENTS

First and foremost I would like to thank (and thank again) my mentor Prof. Jason MacLean. Jason has not only been my advisor, but also my friend and role-model, and has helped me grow and evolve both as a scientist and a person. Through active coaching and by setting an example, he has encouraged broad reading and scholarship, rigorous scientific research, creativity and team-work. I will be forever grateful for his confidence in me and his dedication to my professional development. The joy of research and scientific exploration as well as hours of thought-provoking and inspiring discussions will always be happy memories for me.

I would also like to thank past and present MacLean lab members, who have been my lab-family during the past 6 years. Their expertise and feedback have contributed immensely to this work. I am privileged to have them as my good friends.

Many thanks to my dissertation committee members for their invaluable advice. Prof. Stephanie Palmer has taught me that there are no assumption-free approaches. Furthermore, she encouraged me to turn my assumptions into theories and then testable hypotheses. Prof. John Maunsell has inspired me to consider all the small details that matter to the model organism, a key lesson I have learnt during my thesis research. Prof. Matt Kaufman has encouraged me to place my work within the context of the field and build bridges with other subfields.

I am deeply grateful for my family, especially my parents who have instilled in me the love of nature and books. I also count myself fortunate to have incredible friends with whom I have shared chats, trips, laughs and tears. Last but not least, I would like to thank my partner Micah for his love and support, and for being a bright and warm light during difficult moments.

ABSTRACT

To meet the demands of survival, the central nervous system has to simultaneously encode the external world and internal states in the spiking activity of neurons. Populations of neurons are connected by non-random synaptic wiring, shaped by previous experience, and in turn give rise to variable yet correlated spiking activity. This works attempts to relate the structure of spiking activity to its underlying, interconnected substrate on the one hand, and to the external variables they presumably encode on the other. To do so, statistical dependencies in the activity of neurons are summarized as functional networks (FNs), where neurons are nodes and the statistical regularities between them are edges. In this dissertation, FNs are utilized in encoding, decoding and both generative and discriminative models to gain insights into the circuit level representation of visual stimuli.

As networks, FNs can be readily compared to anatomical and synaptic connectivity in neural network model. This comparison reveals that the structure of statistical dependencies depends on the timescale at which they are computed, and that neural activity has a more clustered structure than the synaptic wiring used in the model. In mouse primary visual cortex, this increased clustering is found to be characteristic of neurons that do not have a clear stimulus preference (i.e. untuned neurons). Moreover, regardless of single cell selectivity, the correlations between neurons are stimulus-specific and hence informative of the visual stimulus. Subsequently, a sparse subset of stimulus-specific pairwise interactions is identified. These correlations reliably manifest as pairs of coordinated spikes on a time-point by time-point basis, building up spatio-temporal sequences. When only these spikes are considered, single neurons still display high trial-to-trial variability. Nonetheless, these spikes have enhanced readability by hypothetical downstream elements.

INTRODUCTION

The subject of study of this dissertation is one of colossal complexity; the human brain contains approximately the same number of neurons as stars in our galaxy (not to mention glia). At the single cell level, each of these neurons is comprised of multiple elements, or compartments, for example dendritic spines. Each of these neurons acts as an integrator of signals, with the manner of integration still a subject of vigorous study. At the population level, neurons are diverse in their cell type with differing cellular properties, and are interwoven into a network. Network architecture, in turn, is determined by numerous factors including, but not limited to, cell type, anatomical location such as area and layer or nucleus, function, and past experience. Indeed, the complexity of the nervous system is further compounded by the ability to respond to changes, known as plasticity, induced by experience. This still elusive, blackbox-like system gives rise to well-conserved, almost universal behaviors like sleep, as well as to gordian thoughts, including self-referential ontological musings and research. The goal of understanding the brain thus seems (perhaps a tad) lofty for a mere Doctor of Philosophy Dissertation. Although research at the single cell resolution has provided us with techniques, experimental protocols and multiple insights into the nervous system, this work's focus is on the population level. Specifically, statistical dependencies between neurons are organized into a network, which is then examined for network structure and network behavior. This work thus promotes a complex systems framework to study a complex system.

Questions of implementational and algorithmic nature

One can invoke Marr's three levels of analysis[1] to discern and formulate the questions a network perspective of the brain should address. Starting from the bottom, at the implementational level, how does network architecture spawn dynamics and what features of the architecture underlie the spiking statistics seen in the brain? Rephrasing these questions with correlations in mind, this work asks to what extent synaptic and functional (summarizing spiking dependencies) networks (FNs) are congruent (chapter I) and to what extent neurons that are functionally different, that is, tuned and untuned neurons in mouse primary visual cortex (V1) have specific positions in a functional network (chapter II).

The first level of analysis then poses structure-function questions, yet 'structure' has several meanings. Structure might refer to anatomy, where one neuron projects to another. Structure might refer to synaptic connectivity, which is not static as synaptic efficacies as well as the internal environment of the neurons change with time. Finally, a correlational structure, or the structure of activity is often studied in relation to how well it predicts other spiking statistics. Chapters I and II follow the first and third definitions presented here, respectively.

Experimentally, obtaining appropriate data to study structure-function hypotheses where the structure in question is anatomical/synaptic is difficult due to the need to simultaneously measure anatomy and activity. Such simultaneous data is lacking even for *Caenorhabditis elegans*, where the connectome is fully mapped. Hence, ample literature have employed network

[1] Marr and Poggio, "From Understanding Computation to Understanding Neural Circuitry."

simulations where the connectivity and simulation parameters are predetermined and the resulting dynamics are being studied[2,3,4,5,6], and chapter I follows their footsteps.

The second (algorithmic) level of analysis pertains to how network activity encodes, manipulates and decodes information about external and internal variables. To answer this, one must specify which neuron and/or population spikes are correlated with anything and everything meaningful, be it sensory stimuli, executed choices and behaviors, and operationalized internal states. In the case of decoding, the 'code' can be further dissected to the identity of neurons or populations participating in the code, to the scheme of spikes that the readout element considers (e.g. rate vs. temporal pattern) and to the readout function (e.g. linear vs. nonlinear) and readout identity. This neural code question is probed in chapter II, investigating whether pairwise correlations banded together into functional networks are stimulus-specific, i.e. correlate with the direction of drifting gratings and can be decoded from. It is further the main focus of chapter III, which asks whether a subset of spikes that are correlated on average carry more information about a stimulus on a time-point by time-point basis, and sharpening the definition of a neural ensemble as a coding unit.

The neural code is the holy grail of Neuroscience, and several coding scheme have been proposed. The gold standard in coding has been the firing rate of a single unit, or more recently, of a small population of neurons. Increased firing rates have been observed in neurons early in the

[2] Doiron et al., "The Mechanics of State-Dependent Neural Correlations."
[3] Brunel, "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons."
[4] Ostojic, Brunel, and Hakim, "How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains."
[5] Roxin, Hakim, and Brunel, "The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons."
[6] Chambers, "Emergent Cooperation Supports Synaptic Integration, Selecting Routes for Activity in Cortical Networks."

sensory pathway such as retinal ganglion cells presented with light in their receptive field[7,8] or spiral ganglion cells in the inner ear in response to their preferred frequency[9]. In cortex, selectivity (tuning) to visual features[10] and sound frequency[11], is based on elevated firing rate in response to specific stimuli. Although still widely applied, the rate code suffers from several caveats and is inherently limited in its representational capacity; Firing rates show sizable variability across repetitions of the same stimulus despite perceptual stability of a stimulus[12,13,14,15]. Most neurons are biophysically unable to modulate firing rate over a large dynamic range, hence requiring unrealistically long integration times to detect an increase in firing rate[16]. Indeed, decoding attempts from firing rates in cortex have resulted in limited accuracy[17,18], and even in retina firing rates of single units do not capture information present in the stimulus[19] at a level explains perception. Coding schemes pooling over firing rates of several tens of neurons were shown to do slightly better[20,21], assuming independence between units, and optimally when correlations between neurons are taken into account[22].

An alternative single unit code, the temporal code, considers the pattern of spikes and silences and has been demonstrated to have larger information capacity and carry information

[7] Kuffler, "Discharge Patterns and Functional Organization of Mammalian Retina."
[8] Barlow, "Summation and Inhibition in the Frog's Retina."
[9] Evans, "The Frequency Response and Other Properties of Single Fibres in the Guinea-Pig Cochlear Nerve."
[10] Hubel and Wiesel, "Receptive Fields and Functional Architecture of Monkey Striate Cortex."
[11] Phillips and Irvine, "Responses of Single Neurons in Physiologically Defined Primary Auditory Cortex (AI) of the Cat."
[12] Ventura, "Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons."
[13] Tolhurst, Movshon, and Dean, "The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex."
[14] Gur, Beylin, and Snodderly, "Response Variability of Neurons in Primary Visual Cortex (V1) of Alert Monkeys."
[15] Deweese and Zador, "Shared and Private Variability in the Auditory Cortex."
[16] Shadlen and Newsome, "The Variable Discharge of Cortical Neurons."
[17] Gawne, "The Simultaneous Coding of Orientation and Contrast in the Responses of V1 Complex Cells."
[18] Bradley et al., "Visual Orientation and Spatial Frequency Discrimination."
[19] Rullen and Thorpe, "Rate Coding Versus Temporal Order Coding."
[20] Graf, "Decoding the Activity of Neuronal Populations in Macaque Primary Visual Cortex."
[21] Berens et al., "A Fast and Simple Population Code for Orientation in Primate V1."
[22] Chen, Geisler, and Seidemann, "Optimal Decoding of Correlated Neural Population Responses in the Primate Visual Cortex."

about orientations in V1[23] as well as the temporal structure of bird song[24]. However, temporal coding schemes still suffer from pitfalls of trial-to-trial variability and the long integration times required to read out a specific pattern of spikes in downstream elements.

At the circuit and population level, another approach over rate, is the combinatorial code. Combinatorial codes postulate that many neurons within the same time-bin create binary words, circumventing the long integration times required to estimate firing rates or temporal patterns in single neurons. Combinatorial codes have been shown to carry visual information about position[25] and stimulus identity[26] in retina, and about movement direction in area MT[27]. An extension of a combinatorial coding scheme from spatial to spatio-temporal, which considers several time bins or frames, yields what has been referred to as a Hebbian assembly[28,29] and Packet coding[30]. Results presented in chapter III agree with this coding scheme, and provide further details about which spatio-temporal sequences, or 'packets' carry information about the stimulus. As these coding schemes depend on the timing of spikes relative to other spikes, investigating the statistical dependencies between neurons may shed light on their organization, capacity and limitations.

Varieties of pairwise statistical dependencies

Neurons in the brain are not statistically independent. A large body of research has attempted to measure pairwise correlations, elucidate their substrate or origin[31], explore their contribution to

[23] Victor and Purpura, "Nature and Precision of Temporal Coding in Visual Cortex."
[24] Chi and Margoliash, "Temporal Precision and Temporal Drift in Brain and Behavior of Zebra Finch Song."
[25] Palmer et al., "Predictive Information in a Sensory Population."
[26] Ganmor, Segev, and Schneidman, "A Thesaurus for a Neural Population Code."
[27] Osborne et al., "The Neural Basis for Combinatorial Coding in a Cortical Population Response."
[28] Harris, "Neural Signatures of Cell Assembly Organization."
[29] Hebb, *The Organization of Behavior*.
[30] Luczak, McNaughton, and Harris, "Packet-Based Communication in the Cortex."
[31] Trong and Rieke, "Origin of Correlated Activity between Parasol Retinal Ganglion Cells."

single cell and population spiking statistics[32] and calculate the extent to which their presence increases or limits the coding capacity of the network[33]. While a comprehensive discussion of correlations is beyond the scope of this work, it is important to point out that the controversy surrounding pairwise dependencies is perhaps partly due to the range of time-scales and trials chosen for their calculation. In other words, what aspect of which spikes is being correlated needs to be clearly stated.

Briefly, signal correlations quantify the extent to which neurons in a pair respond in a similar way to different stimuli[34]. Another way to think about signal correlations in cortical areas where some neurons are selective to a stimulus feature is how similar are the tuning curves for the neurons in the pair. Signal correlations are typically computed from firing rates over the course of a trial and averaged across trials of the same stimulus feature, and are thus a rather crude measure for statistical dependencies in spike trains.

Noise correlations quantify how much neurons in the pair co-vary across trials of repeated stimuli or the same behaviors[35,36]. What can be learned from noise correlations depends heavily on the time-scale and bin resolution at which they are computed, which is commonly quite large to accommodate estimation of the firing rate. Some studies have deemed this term the overall correlation, and factor out signal correlations to obtain a metric independent of the stimulus.

Curiously, from a structure perspective, both types of correlations depend on distance between the pair[37]. From a functional perspective, both are larger for neurons with similar response

[32] Pillow et al., "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population."
[33] Averbeck, Latham, and Pouget, "Neural Correlations, Population Coding and Computation."
[34] Cohen and Kohn, "Measuring and Interpreting Neuronal Correlations."
[35] Cohen and Maunsell, "Attention Improves Performance Primarily by Reducing Interneuronal Correlations."
[36] Cohen and Kohn, "Measuring and Interpreting Neuronal Correlations."
[37] Ecker et al., "The Effect of Noise Correlations in Populations of Diversely Tuned Neurons."

profiles, i.e. selectivity[38], and from a coding perspective, both types were argued to enhance or limit information about the stimulus[39,40,41], albeit for different reasons. Shared variability (noise) cannot be averaged out by downstream elements yet provides a better estimation of the noise entropy of the population[42]. Increased firing rate in several units (signal) is advantageous from a pooling perspective and a winner-take-all scheme, but is highly redundant and thus considered inefficient[43]. The lesson here, again, is that dynamics and structure within those dynamics cannot be easily disentangled from the coding scheme. Whether any correlations are a meaningful measurement or have any effect on coding depends on the assumptions one makes about the code. Along those lines, both signal and noise correlations, being rate-based, are only relevant if downstream elements are sensitive to rate.

A third type of correlations is statistical dependencies in spike timing, i.e. correlations of the spike trains, where bin widths are short and typically do not exceed 30ms. Calculation then takes place across time and not trials. These can be further broken down into synchronous correlations (called synchrony in many studies and coherent oscillations in others[44]), and lagged correlations in which one neuron depends on the activity of the other in the previous time bin. Here a pair of neurons with correlated spike trains is termed functionally correlated / connected, and a collection of these functional connections across all pairs in the population is called a functional network (FN).

---

[38] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[39] Dehaqani et al., "Selective Changes in Noise Correlations Contribute to an Enhanced Representation of Saccadic Targets in Prefrontal Neuronal Ensembles."
[40] Mendels and Shamir, "Relating the Structure of Noise Correlations in Macaque Primary Visual Cortex to Decoder Performance."
[41] Averbeck and Lee, "Effects of Noise Correlations on Information Encoding and Decoding."
[42] Zylberberg, "Untuned but Not Irrelevant: The Role of Untuned Neurons in Sensory Information Coding."
[43] Reich, Mechler, and Victor, "Independent and Redundant Information in Nearby Cortical Neurons."
[44] Averbeck and Lee, "Coding and Transmission of Information by Neural Ensembles."

The author would like to emphasize that it is not the case that functional connections can be studied in an assumption-free manner - far from it. However, an extensive body of research has put forward compelling theoretical arguments and empirical evidence to support the relevance of coordinated spiking activity on short timescales to the activity in a downstream neuron[45,46], reviewed below. For these reasons the author feels more comfortable with the assumptions made in the functional correlation case. Measurement of FNs, their network structure and their correspondence to synaptic networks are discussed in chapter I.

Cortical activity is correlated

Direct measurement of pairwise spiking regularities in cortex only became available in the late 1960's. In a groundbreaking paper, Gerstein and Perkel[47] recorded pairs of neurons in cat auditory cortex in response to sounds and during spontaneous activity and introduced a method to visually inspect the dependency between the two spike trains. They show that functional correlations do not depend on the presence of a stimulus and are not explained by the firing rates of the neurons alone. Numerous subsequent studies in sensory and motor cortices across species and conditions have found that spikes are temporally coordinated across the population[48,49,50,51,52,to name a few]. Intuitively, as Gerstein and Perkel argued, a statistical dependency implies the neurons in the pair might be synaptically connected. However, since neurons do not fire in response to a single input,

[45] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."
[46] deCharms and Merzenich, "Primary Cortical Representation of Sounds by the Coordination of Action-Potential Timing."
[47] Gerstein and Perkel, "Simultaneously Recorded Trains of Action Potentials."
[48] Dan et al., "Coding of Visual Information by Precisely Correlated Spikes in the Lateral Geniculate Nucleus."
[49] Farooq et al., "Strengthened Temporal Coordination within Pre-Existing Sequential Cell Assemblies Supports Trajectory Replay."
[50] Lankarany et al., "Differentially Synchronized Spiking Enables Multiplexed Neural Coding."
[51] Womelsdorf et al., "Modulation of Neuronal Interactions through Neuronal Synchronization."
[52] Gutkin et al., "Turning On and Off with Excitation."

but rather receive multiple excitatory and inhibitory inputs that interact in a non-linear manner, this proved to be more complex and a full mechanistic explanation of which synapses bring about coincident spiking is still lacking.

By definition, a statistical dependency between two neurons means that knowledge of the activity of one neuron would enable better prediction of the activity the other. The validity and extent of such a prediction was still debated, especially since correlation can be caused by a third contributing factor, such as global population fluctuations[53,54], common input[55], or spatial structure in the stimulus[56]. Nevertheless, pairwise functional correlations were used to accurately model the moment-to-moment single-neuron and single-trial activity of retinal ganglion cells[57,58] (for a review in retina see[59]), pyramidal cells in layer 2/3 of primary visual cortex[60,61] (V1) and place and non-place cells in the hippocampus[62]. Furthermore, several studies have demonstrated that pairwise functional interactions can predict the statistics of the whole population; network simulations of coupled Poisson neurons captures the number of active neurons and their temporal sequences in cat and mouse primary visual cortex[63], and a network of leaky-integrate-and-fire neurons with wiring parameters inferred from data reproduces the number of active neurons and the avalanche duration seen in spontaneous activity in mouse somatosensory and auditory

[53] Ecker et al., "State Dependence of Noise Correlations in Macaque Primary Visual Cortex."
[54] Goris, Movshon, and Simoncelli, "Partitioning Neuronal Variability."
[55] Trong and Rieke, "Origin of Correlated Activity between Parasol Retinal Ganglion Cells."
[56] Rikhye and Sur, "Spatial Correlations in Natural Scenes Modulate Response Reliability in Mouse Visual Cortex."
[57] Granot-Atedgi et al., "Stimulus-Dependent Maximum Entropy Models of Neural Population Codes."
[58] Pillow et al., "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population."
[59] Shlens, Rieke, and Chichilnisky, "Synchronized Firing in the Retina."
[60] Kotekal and MacLean, "Recurrent Interactions Can Explain the Variance in Single Trial Responses."
[61] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[62] Meshulam et al., "Collective Behavior of Place and Non-Place Neurons in the Hippocampal Network."
[63] Roxin, Hakim, and Brunel, "The Statistics of Repeating Patterns of Cortical Activity Can Be Reproduced by a Model Network of Stochastic Binary Neurons."

cortices[64] (note that Chapter II takes a similar approach but preserves connectivity on the local scale and employs simpler model neurons).

A different approach that does not require using functional correlations to generate surrogate activity but rather estimates the strength of pairwise interactions that best fit the data is maximum entropy models (MaxEnt). MaxEnt models learn a probability distribution of activity patterns under assumption of (in)dependence between neurons. Models of ganglion cells in salamander retina are able to predict the number of active neurons at any time point in response to a natural movie[65] accounting for pairwise correlations, but performed less well for spatio-temporal patterns composed of multiple neurons across 10 time points. To capture the occurrences of these patterns, higher-order interactions spanning more than a pair needed to be included in the model. A model of smaller populations in cat primary visual cortex[66] have reported similar results, where pairwise correlations accounted for synchrony but failed to model patterns spreading over longer timescales. A third study in macaque V1 found that including pairwise interactions in the model improves prediction of multineuronal firing patterns over a model without such interactions, but the gains for spatially adjacent neurons and for natural stimuli were modest and the authors concluded that beyond-pairwise interactions are needed to fully explain population statistics[67]. Higher-order interactions are revisited in the general discussion section.

Why should one care about prediction of dynamics? Firstly, individual neurons show large trial-to-trial variability despite identical repetition of the stimulus[68,69]. This indicates that the

[64] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry."
[65] Ganmor, Segev, and Schneidman, "Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code."
[66] Köster et al., "Modeling Higher-Order Correlations within Cortical Microcolumns."
[67] Ohiorhenuan et al., "Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks."
[68] Gur, Beylin, and Snodderly, "Response Variability of Neurons in Primary Visual Cortex (V1) of Alert Monkeys."
[69] Deweese and Zador, "Shared and Private Variability in the Auditory Cortex."

stimulus or even the internal state is not the sole contributor (and in fact has very little effect[70]) to the firing statistics of neurons and populations. Elucidating all the factors that contribute to neuronal activity, including anatomical and functional network position, is hence of the utmost importance for our understanding of how spikes carry information or to what extent they are noise. Secondly, assuming one models activity in a biologically plausible manner, such as using an integrate-and-fire neuron model with realistic membrane constant, prediction suggests that the neurons themselves have access to and use short-timescale correlations of neighboring neurons (for a review, see[71]). Intuitively, the neuron is more likely to reach threshold if multiple inputs are adjacent in time, especially given the low firing rates observed in cortex. Direct evidence of this was provided by Chambers and MacLean[72], working with simulations and in mouse somatosensory cortex, where the synaptic structure was unknown: a downstream neurons was tightly locked to the spike times of its two inputs when those two inputs were correlated and neurons which received more clustered inputs had larger depolarizations. Prediction in this sense then provides a mechanistic explanation of how neurons might use regularities in spike timing. Taken together, pairwise correlations are accessible to and orchestrate activity in downstream elements, and are thus in an optimal position to take part in sensory coding.

Pairwise correlations and coding

Several studies across levels of the visual pathway have demonstrated that more information (as Shannon information in bits) can be extracted about the stimulus when taking pairwise correlations into account: in macaque retina, including pairwise correlations in an encoding model and then

[70] Kotekal and MacLean, "Recurrent Interactions Can Explain the Variance in Single Trial Responses."
[71] Salinas and Sejnowski, "Correlated Neuronal Activity and the Flow of Neural Information."
[72] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."

decoding from generated spikes results in increased information about the luminance of pixels in a white-noise stimulus[73]. In salamander retina, where a specific wiring supports selectivity for motion direction in ganglion cells, additional information beyond independent cells (i.e. synergy), depends on the time-scale of correlations and the selectivity of the cells in the pair, with short time-scales and pairs with different preference resulting in synergy[74]. Similar results were reported in macaque V1, where spatial proximity can serve as a proxy for similarity in response properties: Ohiorhenuan et al[75] reported no gains in information for adjacent pairs but additional information carried by correlations for pairs separated by larger distances. Decomposing the information into contributions of stimulus-independent vs. stimulus-dependent spike-train correlations on short-time scales, it was demonstrated that it is specifically stimulus-dependent synchrony or near-synchronous spikes that carry more information about the direction of drifting gratings[76]. The authors stress that they attempted to choose pairs where both neurons were driven by more than one stimulus, crucial to the quantification of stimulus-dependent correlations. Notably, this finding suggests that pairwise correlations are stimulus specific (see also[77]).

The stimulus-specificity of pairwise interactions was subsequently shown in macaque area MT, accompanied by gains in information about the direction of motion. In other words, pairwise correlations themselves are tuned, and their tuning is more stable than single neurons'[78]. Indeed, learning probability distributions across spiking patterns that are conditioned on the stimulus results in different probability distributions explained by a different set of pairwise

[73] Pillow et al., "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population."
[74] Kühn and Gollisch, "Activity Correlations between Direction-Selective Retinal Ganglion Cells Synergistically Enhance Motion Decoding from Complex Visual Scenes."
[75] Ohiorhenuan et al., "Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks."
[76] Montani and Schultz, "Information-Theoretic Analysis of the Role of Correlations in Neural Spike Trains."
[77] Josić et al., "Stimulus-Dependent Correlations and Population Codes."
[78] Ponce-Alvarez et al., "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons."

correlations[79,80]. These studies provide another line of evidence for the stimulus-specificity of correlations. It remains unclear what is the relationship between single cell tuning – a classic concept thought to underlie stimulus encoding (and decoding) since Hubel and Wiesel[81] and stimulus-specific correlations. This question is explored in Chapter II.

Subsequently, and in line with the increased information carried by correlations, numerous studies have shown that taking into account the correlation structure of the population leads accurate decoding of stimulus identity: in theoretical studies and simulations[82,83], retina[84,85], primate primary visual cortex and area MT[86,87,88], mouse somatosensory[89] and primary visual cortex[90], and in hippocampus[91]. However, many of the studies, especially those in cortex, have used noise correlations, i.e. correlations in rate across trials. As with information estimation, faster correlations (i.e. spike train correlations) are more accessible to downstream elements in real time. Chapter III takes this idea a step further, and examines how correlations that are computed across time are instantiated on a time frame by time frame basis. Some studies have shown that information-theoretic measures and decoders that explicitly incorporate higher-order correlations,

[79] Granot-Atedgi et al., "Stimulus-Dependent Maximum Entropy Models of Neural Population Codes."
[80] Ganmor, Segev, and Schneidman, "A Thesaurus for a Neural Population Code."
[81] Hubel and Wiesel, "Receptive Fields of Single Neurones in the Cat's Striate Cortex."
[82] Schaub and Schultz, "The Ising Decoder."
[83] Shamir and Sompolinsky, "Implications of Neuronal Diversity on Population Coding."
[84] Franke et al., "Structures of Neural Correlation and How They Favor Coding."
[85] Ruda, Zylberberg, and Field, "Ignoring Correlated Activity Causes a Failure of Retinal Population Codes."
[86] Zavitz et al., "Correlated Variability in the Neurons With the Strongest Tuning Improves Direction Coding."
[87] Chen, Geisler, and Seidemann, "Optimal Decoding of Correlated Neural Population Responses in the Primate Visual Cortex."
[88] Graf, "Decoding the Activity of Neuronal Populations in Macaque Primary Visual Cortex."
[89] Sederberg, Palmer, and MacLean, "Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity."
[90] Montijn, Vinck, and Pennartz, "Population Coding in Mouse Visual Cortex."
[91] Meshulam et al., "Collective Behavior of Place and Non-Place Neurons in the Hippocampal Network."

beyond pairwise[92,93], have enhanced explanatory power, motivating transition into functional networks.

A network science approach to the brain

A complex system such as the brain calls for an analytic framework that is equipped to deal with multiple interacting elements. An added intricacy in networks of neurons over other complex system is that for the former, the structure-function and coding questions posed above are closely tied and difficult to disentangle. In other words, the 'code' is inseparable from the synaptic mechanisms that generate spiking.

Complex systems with multiple interconnected elements interacting over time have been long studied in physics[94] and genetics[95], and more recently in social sciences[96]. Indeed, the frameworks of complex network analysis and graph theory date back to Euler[97]. In neuroscience, a graph theory inspired paradigm conceptualizes synaptic networks as a weighted and directed graph with neurons constituting vertices and synaptic connections serving as edges. Practically, the graph (also referred to as topology, interchangeably) can be represented as an adjacency matrix. To summarize correlated activity in a functional network (FN), edge weights stand for the strength of the statistical dependency, and the direction is often set by the lag. Consequently, circuit dynamics, represented as a graph, become mathematically tractable.

[92] Köster et al., "Modeling Higher-Order Correlations within Cortical Microcolumns."
[93] Ganmor, Segev, and Schneidman, "Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code."
[94] Dorogovtsev, Goltsev, and Mendes, "Ising Model on Networks with an Arbitrary Distribution of Connections."
[95] Akutsu, Miyano, and Kuhara, "Inferring Qualitative Relations in Genetic Networks and Metabolic Pathways."
[96] Sekara, Stopczynski, and Lehmann, "Fundamental Structures of Dynamic Social Networks."
[97] Euler, "Solutio Problematis Ad Geometriam Situs Pertinentis."

Graph theory methods have been employed to describe the structure of cortical circuits[98,99] exposing network properties supporting balanced dynamics[100], information transmission[101] and optimal storage capacity[102]. However, to date very few studies have attempted to directly tie together dynamics and coding with this paradigm. Shi et al[103] have included an adjacency matrix in a generalized-linear-model decoder of V1 activity, exhibiting only moderate success. Yet, their study was limited to 16 neurons in anesthetized rats, and model selection and fitting could be improved.

An FN approach has several advantages: First, it preserves the full dimensionality of dynamics and cell identities, enhancing interpretability, as opposed to transforming the dynamics to another space. It also allows for examination of the single-cell within the context of a network, which is especially useful for building upon decades of research measuring single-cell responses towards a description of population activity, as done in chapter II. As the field ventures into the anatomical and functional connectivity within and between different cell types, and all-optical simultaneous and targeted single-cell manipulation and imaging becomes more widely used, preserving cell labels will prove helpful for both experiments and analysis. Second, FNs allow for relatively easy analysis of higher-order functional interactions[104,105] and the rich graph-theory literature aids the construction of appropriate null models (unfortunately still underutilized). The

[98] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."
[99] Reimann et al., "Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function."
[100] Landau et al., "The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks."
[101] Timme et al., "High-Degree Neurons Feed Cortical Computations."
[102] Brunel, "Is Cortical Connectivity Optimized for Storing Information?"
[103] Shi, Niu, and Wan, "Effect of the Small-World Structure on Encoding Performance in the Primary Visual Cortex."
[104] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[105] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."

FN approach complements probabilistic methods for detecting higher-order interactions[106] due to faster runtimes and the ability to combine motifs into clusters.

Finally, FNs require relatively small amounts of data to compute. This enables the computation of stimulus-specific FNs, as well as recent advancements in the temporal domain, where FNs can be used to track changes in dynamics over time such as the course of a trial. Sliding window approaches in which an FN is generated for every time epoch have shown that: one, changes to the core-periphery network structure as a function of time are largely independent of global brain states and anatomical location[107]. And two, that reaches to different directions in monkey M1 are characterized by stereotyped and direction-specific temporal evolution of FNs, in line with Hebbian phase sequences (Sundiang, unpublished personal communication). Chapter III employed an intersection step, similar to Dan et al[108] to obtain instantaneous FNs, but the evolution of these FNs, i.e. how the statistics of one time point are related to the next, are left for future research.

[106] Ganmor, Segev, and Schneidman, "Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code."
[107] Pedreschi et al., "Dynamic Core-Periphery Structure of Information Sharing Networks in Entorhinal Cortex and Hippocampus."
[108] Dan et al., "Coding of Visual Information by Precisely Correlated Spikes in the Lateral Geniculate Nucleus."

REFERENCES

Akutsu, Tatsuya, Satoru Miyano, and Satoru Kuhara. "Inferring Qualitative Relations in Genetic Networks and Metabolic Pathways." *Bioinformatics* 16, no. 8 (August 1, 2000): 727–34. https://doi.org/10.1093/bioinformatics/16.8.727.

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. https://doi.org/10.1038/nrn1888.

Averbeck, Bruno B., and Daeyeol Lee. "Coding and Transmission of Information by Neural Ensembles." *Trends in Neurosciences* 27, no. 4 (April 1, 2004): 225–30. https://doi.org/10.1016/j.tins.2004.02.006.

Averbeck, Bruno B., and Daeyeol Lee. "Effects of Noise Correlations on Information Encoding and Decoding." *Journal of Neurophysiology* 95, no. 6 (June 1, 2006): 3633–44. https://doi.org/10.1152/jn.00919.2005.

Barlow, H. B. "Summation and Inhibition in the Frog's Retina." *The Journal of Physiology* 119, no. 1 (January 28, 1953): 69–88.

Berens, P., A. S. Ecker, R. J. Cotton, W. J. Ma, M. Bethge, and A. S. Tolias. "A Fast and Simple Population Code for Orientation in Primate V1." *Journal of Neuroscience* 32, no. 31 (August 1, 2012): 10618–26. https://doi.org/10.1523/JNEUROSCI.1335-12.2012.

Bradley, Arthur, Bernt C. Skottun, Izumi Ohzawa, Gary Sclar, and RALPH D. Freeman. "Visual Orientation and Spatial Frequency Discrimination: A Comparison of Single Neurons and Behavior." *Journal of Neurophysiology* 57, no. 3 (1987): 755–72.

Brunel, Nicolas. "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons." *Journal of Computational Neuroscience* 8, no. 3 (May 1, 2000): 183–208. https://doi.org/10.1023/A:1008925309027.

Brunel, Nicolas. "Is Cortical Connectivity Optimized for Storing Information?" *Nature Neuroscience* 19, no. 5 (April 11, 2016): 749–55. https://doi.org/10.1038/nn.4286.

Chambers, Brendan, and Jason N. MacLean. "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks." Edited by Jochen Triesch. *PLOS Computational Biology* 12, no. 8 (August 19, 2016): e1005078. https://doi.org/10.1371/journal.pcbi.1005078.

Chen, Yuzhi, Wilson S Geisler, and Eyal Seidemann. "Optimal Decoding of Correlated Neural Population Responses in the Primate Visual Cortex." *Nature Neuroscience* 9, no. 11 (November 2006): 1412–20. https://doi.org/10.1038/nn1792.

Chi, Zhiyi, and Daniel Margoliash. “Temporal Precision and Temporal Drift in Brain and Behavior of Zebra Finch Song.” *Neuron* 32, no. 5 (December 6, 2001): 899–910. https://doi.org/10.1016/S0896-6273(01)00524-4.

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

Cohen, Marlene R., and John H. R. Maunsell. “Attention Improves Performance Primarily by Reducing Interneuronal Correlations.” *Nature Neuroscience* 12, no. 12 (December 2009): 1594–1600. https://doi.org/10.1038/nn.2439.

Dan, Yang, Jose-Manuel Alonso, W. Martin Usrey, and R. Clay Reid. “Coding of Visual Information by Precisely Correlated Spikes in the Lateral Geniculate Nucleus.” *Nature Neuroscience* 1, no. 6 (October 1998): 501–7. https://doi.org/10.1038/2217.

deCharms, R. Christopher, and Michael M. Merzenich. “Primary Cortical Representation of Sounds by the Coordination of Action-Potential Timing.” *Nature* 381, no. 6583 (June 1996): 610–13. https://doi.org/10.1038/381610a0.

Dechery, Joseph B., and Jason N. MacLean. “Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons.” Edited by Jeff Beck. *PLOS Computational Biology* 14, no. 5 (May 4, 2018): e1006153. https://doi.org/10.1371/journal.pcbi.1006153.

Dehaqani, Mohammad-Reza A, Abdol-Hossein Vahabie, Mohammadbagher Parsa, Behrad Noudoost, and Alireza Soltani. “Selective Changes in Noise Correlations Contribute to an Enhanced Representation of Saccadic Targets in Prefrontal Neuronal Ensembles.” *Cerebral Cortex* 28, no. 8 (August 1, 2018): 3046–63. https://doi.org/10.1093/cercor/bhy141.

Deweese, Michael R., and Anthony M. Zador. “Shared and Private Variability in the Auditory Cortex.” *Journal of Neurophysiology* 92, no. 3 (September 1, 2004): 1840–55. https://doi.org/10.1152/jn.00197.2004.

Doiron, Brent, Ashok Litwin-Kumar, Robert Rosenbaum, Gabriel K. Ocker, and Krešimir Josić. “The Mechanics of State-Dependent Neural Correlations.” *Nature Neuroscience* 19, no. 3 (March 2016): 383–93. https://doi.org/10.1038/nn.4242.

Dorogovtsev, S. N., A. V. Goltsev, and J. F. F. Mendes. “Ising Model on Networks with an Arbitrary Distribution of Connections.” *Physical Review E* 66, no. 1 (July 8, 2002): 016104. https://doi.org/10.1103/PhysRevE.66.016104.

Ecker, Alexander S., Philipp Berens, Andreas S. Tolias, and Matthias Bethge. “The Effect of Noise Correlations in Populations of Diversely Tuned Neurons.” *Journal of Neuroscience* 31, no. 40 (October 5, 2011): 14272–83. https://doi.org/10.1523/JNEUROSCI.2539-11.2011.

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. https://doi.org/10.1016/j.neuron.2014.02.006.

Euler, Leonhard. "Solutio Problematis Ad Geometriam Situs Pertinentis." *Commentarii Academiae Scientiarum Petropolitanae*, January 1, 1741, 128–40.

Evans, E. F. "The Frequency Response and Other Properties of Single Fibres in the Guinea-Pig Cochlear Nerve." *The Journal of Physiology* 226, no. 1 (October 1, 1972): 263–87. https://doi.org/10.1113/jphysiol.1972.sp009984.

Farooq, Usman, Jeremie Sibille, Kefei Liu, and George Dragoi. "Strengthened Temporal Coordination within Pre-Existing Sequential Cell Assemblies Supports Trajectory Replay." *Neuron* 103, no. 4 (August 21, 2019): 719-733.e7. https://doi.org/10.1016/j.neuron.2019.05.040.

Franke, Felix, Michele Fiscella, Maksim Sevelev, Botond Roska, Andreas Hierlemann, and Rava Azeredo da Silveira. "Structures of Neural Correlation and How They Favor Coding." *Neuron* 89, no. 2 (January 20, 2016): 409–22. https://doi.org/10.1016/j.neuron.2015.12.037.

Ganmor, E., R. Segev, and E. Schneidman. "Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code." *Proceedings of the National Academy of Sciences* 108, no. 23 (June 7, 2011): 9679–84. https://doi.org/10.1073/pnas.1019641108.

Ganmor, Elad, Ronen Segev, and Elad Schneidman. "A Thesaurus for a Neural Population Code." *Elife* 4 (2015): e06134.

Gawne, Timothy J. "The Simultaneous Coding of Orientation and Contrast in the Responses of V1 Complex Cells." *Experimental Brain Research* 133, no. 3 (July 17, 2000): 293–302. https://doi.org/10.1007/s002210000381.

Gerstein, George L., and Donald H. Perkel. "Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation." *Science* 164, no. 3881 (1969): 828–30.

Goris, Robbe L. T., J. Anthony Movshon, and Eero P. Simoncelli. "Partitioning Neuronal Variability." *Nature Neuroscience* 17, no. 6 (June 2014): 858–65. https://doi.org/10.1038/nn.3711.

Graf, Arnulf B A. "Decoding the Activity of Neuronal Populations in Macaque Primary Visual Cortex." *Nature NEUROSCIENCE* 14, no. 2 (2011): 9.

Granot-Atedgi, Einat, Gašper Tkačik, Ronen Segev, and Elad Schneidman. "Stimulus-Dependent Maximum Entropy Models of Neural Population Codes." *PLOS Computational Biology* 9, no. 3 (March 14, 2013): e1002922. https://doi.org/10.1371/journal.pcbi.1002922.

Gur, M., A. Beylin, and D. M. Snodderly. "Response Variability of Neurons in Primary Visual Cortex (V1) of Alert Monkeys." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 17, no. 8 (April 15, 1997): 2914–20.

Gutkin, Boris S., Carlo R. Laing, Carol. L. Colby, Carson C. Chow, and G. Bard Ermentrout. "Turning On and Off with Excitation: The Role of Spike-Timing Asynchrony and Synchrony in Sustained Neural Activity." *Journal of Computational Neuroscience* 11, no. 2 (September 1, 2001): 121–34. https://doi.org/10.1023/A:1012837415096.

Harris, Kenneth D. "Neural Signatures of Cell Assembly Organization." *Nature Reviews Neuroscience* 6, no. 5 (May 2005): 399–407. https://doi.org/10.1038/nrn1669.

Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory*. Psychology Press, 1949.

Hubel, D. H., and T. N. Wiesel. "Receptive Fields and Functional Architecture of Monkey Striate Cortex." *The Journal of Physiology* 195, no. 1 (March 1, 1968): 215–43. https://doi.org/10.1113/jphysiol.1968.sp008455.

Hubel, D. H., and T. N. Wiesel. "Receptive Fields of Single Neurones in the Cat's Striate Cortex." *The Journal of Physiology* 148, no. 3 (October 1959): 574–91.

Josić, Krešimir, Eric Shea-Brown, Brent Doiron, and Jaime de la Rocha. "Stimulus-Dependent Correlations and Population Codes." *Neural Computation* 21, no. 10 (2009): 2774–2804.

Köster, Urs, Jascha Sohl-Dickstein, Charles M. Gray, and Bruno A. Olshausen. "Modeling Higher-Order Correlations within Cortical Microcolumns." Edited by Jakob H. Macke. *PLoS Computational Biology* 10, no. 7 (July 3, 2014): e1003684. https://doi.org/10.1371/journal.pcbi.1003684.

Kotekal, Subhodh, and Jason N. MacLean. "Recurrent Interactions Can Explain the Variance in Single Trial Responses." *PLOS Computational Biology* 16, no. 1 (January 30, 2020): e1007591. https://doi.org/10.1371/journal.pcbi.1007591.

Kuffler, Stephen W. "Discharge Patterns and Functional Organization of Mammalian Retina." *Journal of Neurophysiology* 16, no. 1 (January 1, 1953): 37–68.

Kühn, Norma Krystyna, and Tim Gollisch. "Activity Correlations between Direction-Selective Retinal Ganglion Cells Synergistically Enhance Motion Decoding from Complex Visual Scenes." *Neuron* 101, no. 5 (March 6, 2019): 963-976.e7. https://doi.org/10.1016/j.neuron.2019.01.003.

Landau, Itamar D., Robert Egger, Vincent J. Dercksen, Marcel Oberlaender, and Haim Sompolinsky. “The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks.” *Neuron* 92, no. 5 (December 2016): 1106–21. https://doi.org/10.1016/j.neuron.2016.10.027.

Lankarany, Milad, Dhekra Al-Basha, Stéphanie Ratté, and Steven A. Prescott. “Differentially Synchronized Spiking Enables Multiplexed Neural Coding.” *Proceedings of the National Academy of Sciences* 116, no. 20 (May 14, 2019): 10097–102.

Luczak, Artur, Bruce L. McNaughton, and Kenneth D. Harris. “Packet-Based Communication in the Cortex.” *Nature Reviews Neuroscience* 16 (2015): 745–55.

Marr, D., and T. Poggio. “From Understanding Computation to Understanding Neural Circuitry,” May 1, 1976. https://dspace.mit.edu/handle/1721.1/5782.

Mendels, Or P., and Maoz Shamir. “Relating the Structure of Noise Correlations in Macaque Primary Visual Cortex to Decoder Performance.” *Frontiers in Computational Neuroscience* 12 (2018). https://doi.org/10.3389/fncom.2018.00012.

Meshulam, Leenoy, Jeffrey L. Gauthier, Carlos D. Brody, David W. Tank, and William Bialek. “Collective Behavior of Place and Non-Place Neurons in the Hippocampal Network.” *Neuron* 96, no. 5 (December 6, 2017): 1178-1191.e4. https://doi.org/10.1016/j.neuron.2017.10.027.

Montani, Fernando, and Simon R. Schultz. “Information-Theoretic Analysis of the Role of Correlations in Neural Spike Trains.” In *Concepts and Recent Advances in Generalized Information Measures and Statistics*, 375–407. Accessed July 5, 2021. https://www.eurekaselect.com/118914/chapter.

Montijn, Jorrit S., Martin Vinck, and Cyriel M. A. Pennartz. “Population Coding in Mouse Visual Cortex: Response Reliability and Dissociability of Stimulus Tuning and Noise Correlation.” *Frontiers in Computational Neuroscience* 8 (June 2, 2014). https://doi.org/10.3389/fncom.2014.00058.

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. https://doi.org/10.1038/nature09178.

Osborne, L. C., S. E. Palmer, S. G. Lisberger, and W. Bialek. “The Neural Basis for Combinatorial Coding in a Cortical Population Response.” *Journal of Neuroscience* 28, no. 50 (December 10, 2008): 13522–31. https://doi.org/10.1523/JNEUROSCI.4390-08.2008.

Ostojic, Srdjan, Nicolas Brunel, and Vincent Hakim. “How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains.” *Journal of Neuroscience* 29, no. 33 (August 19, 2009): 10234–53.

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. https://doi.org/10.1073/pnas.1506855112.

Pedreschi, Nicola, Christophe Bernard, Wesley Clawson, Pascale Quilichini, Alain Barrat, and Demian Battaglia. “Dynamic Core-Periphery Structure of Information Sharing Networks in Entorhinal Cortex and Hippocampus.” *Network Neuroscience* 4, no. 3 (September 1, 2020): 946–75. https://doi.org/10.1162/netn_a_00142.

Phillips, D. P., and D. R. Irvine. “Responses of Single Neurons in Physiologically Defined Primary Auditory Cortex (AI) of the Cat: Frequency Tuning and Responses to Intensity.” *Journal of Neurophysiology* 45, no. 1 (January 1, 1981): 48–58.

Pillow, Jonathan W., Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M. Litke, E. J. Chichilnisky, and Eero P. Simoncelli. “Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population.” *Nature* 454, no. 7207 (August 21, 2008): 995–99. https://doi.org/10.1038/nature07140.

Ponce-Alvarez, A., A. Thiele, T. D. Albright, G. R. Stoner, and G. Deco. “Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons.” *Proceedings of the National Academy of Sciences* 110, no. 32 (August 6, 2013): 13162–67. https://doi.org/10.1073/pnas.1300098110.

Reich, Daniel S., Ferenc Mechler, and Jonathan D. Victor. “Independent and Redundant Information in Nearby Cortical Neurons.” *Science* 294, no. 5551 (December 21, 2001): 2566–68. https://doi.org/10.1126/science.1065839.

Reimann, Michael W., Max Nolte, Martina Scolamiero, Katharine Turner, Rodrigo Perin, Giuseppe Chindemi, Paweł Dłotko, Ran Levi, Kathryn Hess, and Henry Markram. “Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function.” *Frontiers in Computational Neuroscience* 11 (2017). https://doi.org/10.3389/fncom.2017.00048.

Rikhye, R. V., and M. Sur. “Spatial Correlations in Natural Scenes Modulate Response Reliability in Mouse Visual Cortex.” *Journal of Neuroscience* 35, no. 43 (October 28, 2015): 14661–80. https://doi.org/10.1523/JNEUROSCI.1660-15.2015.

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.” *Journal of Neuroscience* 28, no. 42 (October 15, 2008): 10734–45.

Ruda, Kiersten, Joel Zylberberg, and Greg D. Field. “Ignoring Correlated Activity Causes a Failure of Retinal Population Codes.” *Nature Communications* 11, no. 1 (September 14, 2020): 4605. https://doi.org/10.1038/s41467-020-18436-2.

Rullen, Rufin Van, and Simon J. Thorpe. “Rate Coding Versus Temporal Order Coding: What the Retinal Ganglion Cells Tell the Visual Cortex.” *Neural Computation* 13, no. 6 (June 1, 2001): 1255–83. https://doi.org/10.1162/08997660152002852.

Sadovsky, A. J., and J. N. MacLean. “Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry.” *Journal of Neuroscience* 33, no. 35 (August 28, 2013): 14048–60. https://doi.org/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. https://doi.org/10.1038/35086012.

Schaub, Michael T., and Simon R. Schultz. “The Ising Decoder: Reading out the Activity of Large Neural Ensembles.” *Journal of Computational Neuroscience* 32, no. 1 (February 2012): 101–18. https://doi.org/10.1007/s10827-011-0342-z.

Sederberg, Audrey J., Stephanie E. Palmer, and Jason N. MacLean. “Decoding Thalamic Afferent Input Using Microcircuit Spiking Activity.” *Journal of Neurophysiology* 113, no. 7 (April 2015): 2921–33. https://doi.org/10.1152/jn.00885.2014.

Sekara, Vedran, Arkadiusz Stopczynski, and Sune Lehmann. “Fundamental Structures of Dynamic Social Networks.” *Proceedings of the National Academy of Sciences* 113, no. 36 (September 6, 2016): 9977–82. https://doi.org/10.1073/pnas.1602803113.

Shadlen, M. N., and W. T. Newsome. “The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 18, no. 10 (May 15, 1998): 3870–96.

Shamir, M., and H. Sompolinsky. “Implications of Neuronal Diversity on Population Coding.” *Neural Computation* 18, no. 8 (August 2006): 1951–86. https://doi.org/10.1162/neco.2006.18.8.1951.

Shi, Li, Xiaoke Niu, and Hong Wan. “Effect of the Small-World Structure on Encoding Performance in the Primary Visual Cortex: An Electrophysiological and Modeling Analysis.” *Journal of Comparative Physiology A* 201, no. 5 (May 2015): 471–83. https://doi.org/10.1007/s00359-015-0996-5.

Shlens, Jonathon, Fred Rieke, and EJ Chichilnisky. “Synchronized Firing in the Retina.” *Current Opinion in Neurobiology*, Sensory systems, 18, no. 4 (August 1, 2008): 396–402. https://doi.org/10.1016/j.conb.2008.09.010.

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii. "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits." Edited by Karl J. Friston. *PLoS Biology* 3, no. 3 (March 1, 2005): e68. https://doi.org/10.1371/journal.pbio.0030068.

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 Computational Biology* 12, no. 5 (2016): e1004858.

Tolhurst, D. J., J. A. Movshon, and A. F. Dean. "The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex." *Vision Research* 23, no. 8 (1983): 775–85. https://doi.org/10.1016/0042-6989(83)90200-6.

Trong, Philipp Khuc, and Fred Rieke. "Origin of Correlated Activity between Parasol Retinal Ganglion Cells." *Nature Neuroscience* 11, no. 11 (November 2008): 1343–51. https://doi.org/10.1038/nn.2199.

Ventura, V. "Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons." *Journal of Neurophysiology* 94, no. 4 (April 20, 2005): 2928–39. https://doi.org/10.1152/jn.00644.2004.

Victor, J. D., and K. P. Purpura. "Nature and Precision of Temporal Coding in Visual Cortex: A Metric-Space Analysis." *Journal of Neurophysiology* 76, no. 2 (August 1, 1996): 1310–26.

Womelsdorf, Thilo, Jan-Mathijs Schoffelen, Robert Oostenveld, Wolf Singer, Robert Desimone, Andreas K. Engel, and Pascal Fries. "Modulation of Neuronal Interactions through Neuronal Synchronization." *Science* 316, no. 5831 (June 15, 2007): 1609–12. https://doi.org/10.1126/science.1139597.

Zavitz, Elizabeth, Hsin-Hao Yu, Marcello G P Rosa, and Nicholas S C Price. "Correlated Variability in the Neurons With the Strongest Tuning Improves Direction Coding." *Cerebral Cortex* 29, no. 2 (February 1, 2019): 615–26. https://doi.org/10.1093/cercor/bhx344.

Zylberberg, Joel. "Untuned but Not Irrelevant: The Role of Untuned Neurons in Sensory Information Coding." *BioRxiv*, January 1, 2017, 134379. https://doi.org/10.1101/134379.

CHAPTER I

Ensemble stacking mitigates biases in inference of synaptic connectivity

This work was previously published: Chambers, B.*, Levy, M.*, Dechery, J. B., & MacLean, J. N. (2018). Ensemble stacking mitigates biases in inference of synaptic connectivity. *Network Neuroscience*, *2*(1), 60-85. (*co-first authors / equal contribution)

ABSTRACT

A promising alternative to directly measuring the anatomical connections in a neuronal population is inferring the connections from the activity. We employ simulated spiking neuronal networks to compare and contrast commonly used inference methods that identify likely excitatory synaptic connections using statistical regularities in spike-timing. We find that simple adjustments to standard algorithms improve inference accuracy: a signing procedure improves the power of unsigned mutual-information based approaches and a correction which accounts for differences in mean and variance of background timing relationships, such as those expected to be induced by heterogeneous firing rates, increases the sensitivity of frequency-based methods. We also find that different inference methods reveal distinct subsets of the synaptic network and each method exhibits different biases in the accurate detection of reciprocity and local clustering. To correct for errors and biases specific to single inference algorithms we combine methods into an Ensemble. Ensemble predictions, generated as a linear combination of multiple inference algorithms, are more sensitive than the best individual measures alone, and are more faithful to ground-truth statistics of connectivity, mitigating biases specific to single inference methods. These weightings generalize across simulated datasets, emphasizing the potential for the broad utility of Ensemble based approaches.

INTRODUCTION

Propagation of activity within neuronal networks is largely determined by underlying synaptic connectivity[1,2,3]. This link has been demonstrated using recordings from pairs and small groups of neurons and has provided insights into plasticity processes[4,5], circuit structure[6,7,8] and noise correlations[9]. While methods such as paired patch clamp recordings or electron microscopy provide unambiguous indication of a synaptic connection they are technically limited to the examination of a small number of connections with unknown functional relationships. Consequently the statistics of circuit connectivity at the mesoscopic scale are difficult to conclude due to finite size effect errors[10]. In contrast, measures of dynamics, such as those generated by two-photon imaging of calcium fluorescence indicators[11], allow up to a 1000 neurons to be recorded but require that synaptic connections be inferred using statistical dependencies in spike timing. Because neuronal spiking in neocortical networks requires synaptic input, the causal relationship between connectivity and activity can be exploited to infer network topology in direct relation to synaptic recruitment[12]. In this framework, statistical dependencies in the spiking activity between pairs of neurons within a population are summarized as a weighted directed graph, and

[1] Gerstein and Perkel, "Simultaneously Recorded Trains of Action Potentials."
[2] Lindsey et al., "Repeated Patterns of Distributed Synchrony in Neuronal Assemblies."
[3] Kumar, Rotter, and Aertsen, "Spiking Activity Propagation in Neuronal Networks."
[4] Kruskal, Li, and MacLean, "Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex."
[5] Lalanne, Abrahamsson, and Sjöström, "Using Multiple Whole-Cell Recordings to Study Spike-Timing-Dependent Plasticity in Acute Neocortical Slices."
[6] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."
[7] Perin, Berger, and Markram, "A Synaptic Organizing Principle for Cortical Neuronal Groups."
[8] Ko et al., "Functional Specificity of Local Synaptic Connections in Neocortical Networks."
[9] Hofer et al., "Differential Connectivity and Response Dynamics of Excitatory and Inhibitory Neurons in Visual Cortex."
[10] Vegue, Perin, and Roxin, "On the Structure of Cortical Micro-Circuits Inferred from Small Sample Sizes."
[11] Sadovsky et al., "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging."
[12] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

this weight matrix is informative about likelihood of synaptic connections as well as their functional relationship. Only those synapses directly contributing to spike-time dependencies can be captured and summarized by these weight matrices (Chambers and MacLean 2015). While this lessens the number of synaptic connections that can be inferred, this subset of connections (which we have referred to as the 'recruitment' network) have particular importance for the propagation of spiking and are desirable targets for inference (Chambers & MacLean 2016). Study of the link between structural and functional connectivity has the promise to reveal mechanistic insights as to how information flow is directed across networks and the number of studies employing inference algorithms has grown rapidly. In this paper we compare and contrast the performance of a number of common inference methods, identify biases specific to individual inference methods and then combine them in an ensemble to mitigate these biases and consequently improve inference of synaptic connectivity within large networks of neurons.

The importance of bridging function and structure is highlighted by the increasing diversity of methods for predicting synaptic connectivity from spiking activity. These efforts encompass methods based on counting lagged firing events[13], lagged correlation[14], mutual information[15], and transfer entropy[16,17] sometimes also referred to as conditional mutual information[18], as well as other approaches. Crucial to the performance of these methods is the time resolution of the spike trains, or the bin size in the binning procedure[19] over relevant time scales relating spiking to

[13] Pajevic and Plenz, "The Organization of Strong Links in Complex Networks."
[14] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry," August 28, 2013.
[15] Endo et al., "Delayed Mutual Information Infers Patterns of Synaptic Connectivity in a Proprioceptive Neural Network."
[16] Ito et al., "Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model."
[17] Stetter et al., "Model-Free Reconstruction of Excitatory Neuronal Connectivity from Calcium Imaging Signals."
[18] Zhang et al., "Conditional Mutual Inclusive Information Enables Accurate Quantification of Associations in Gene Regulatory Networks."
[19] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under

synaptic connectivity and integration. Consequently we consider a range of bin widths in this work. The majority of these metrics are formulated in terms of correlations between consecutive time bins, which implies causality. However, as the number of neurons densely recorded with imaging increases, acquisition time generally increases as well, and the relevant correlations may shift from consecutive time bins towards simultaneous time bins. We thus propose variations of information theoretic measures that account for simultaneous and joint time bin correlations to reflect common experimental constraints.

Because each of the inference algorithms differentially quantify statistical features of population dynamics, it is possible that they identify non-identical sets of connections. Therefore, it is opportune to turn to progress in machine learning and bioinformatics, which have shown that combining approaches has the potential to pool over strengths and neutralize weaknesses of their constituent algorithms[20]. Collectively, these strategies are known as Ensemble methods. An early formal description of an Ensemble method was introduced to address the difficulty of optimization under conditions of multiple local minima in neural networks[21]. Ensemble learning has been shown to occur in both the frequentist and Bayesian frameworks of machine learning exemplified respectively by boosting[22] and Bayesian optimal classifiers. In diverse settings, combinations of algorithms can be stacked together to yield an aggregate that outperforms its components[23]. However, it is unclear whether an Ensemble approach can be applied to the problem of synaptic inference. We investigate this issue using simulated naturalistic spiking networks, where true underlying connectivity is known in full. Computational models are well-suited to investigating

Realistic Experimental Constraints."
[20] Marbach et al., "Wisdom of Crowds for Robust Gene Network Inference."
[21] Hansen and Salamon, "Neural Network Ensembles."
[22] Freund and Schapire, "A Desicion-Theoretic Generalization of on-Line Learning and an Application to Boosting"; Schapire, "The Strength of Weak Learnability."
[23] Fast and Jensen, "Why Stacked Models Perform Effective Collective Classification."

synaptic topology, because they present transparent access to synaptic connectivity and spike-timing. Prior work has demonstrated that random synaptic networks present the greatest challenges for inferring synaptic connections from activity[24]. Therefore, although synaptic networks in the neocortex are known to be non-random, random models are useful for benchmarking success-rates in synaptic inference applications. It is also important to consider naturalistic regimes of population activity, because network state can determine the success or failure when inferring synaptic connections[25]. To compare diverse inference methods and judge whether they may complement one another, we designed network simulations to reflect naturalistic dynamics and reflect experimental constraints. We show that an Ensemble approach reveals a more extensive subset of the synaptic network, and one that is more faithful to the true statistics of the synaptic recruitment network measured in our simulations. A host of weighted combinations improve over the best individual measures, and these weighting schemes are transferrable from one simulated data set to another.

METHODS

Network simulations

In order to obtain a ground-truth for network connectivity, we constructed six networks of conductance-based leaky-integrate-and-fire (LIF) neurons, similar to previous work[26][27]. Each simulated dataset was defined as a synaptic connectivity matrix (Fig 1C) with 1000 excitatory units and 200 inhibitory units, constructed probabilistically according to a stochastic blocked

[24] 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."
[25] Stetter et al., "Model-Free Reconstruction of Excitatory Neuronal Connectivity from Calcium Imaging Signals."
[26] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."
[27] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."

model with naturalistic connection probabilities. The excitatory subnetwork taken in isolation formed an Erdos-Renyi graph (p=0.2). Specifically, we set $P_{ee} = 0.2$, $P_{ei} = 0.35$, $P_{ie} = 0.25$ and $P_{ii} = 0.3$, with e denoting excitatory units, i denoting inhibitory neurons, and the first and second letters in the subscript standing for the pre- and post- synaptic neuron, respectively.

Each neuron's membrane potential was governed by:

(1) $$\tau_m \frac{dV}{dt} = g_e(E_e - V) + g_i(E_i - V) + g_l(E_l - V) + g_{tonic}(E_{tonic} - V)$$

(2) $$\tau_e \frac{dg_e}{dt} = -g_e$$

(3) $$\tau_i \frac{dg_i}{dt} = -g_i$$

A spike occurred every time the membrane potential crossed a threshold, set at -48 mV. Post spike, membrane potential was then reset to -70 mv, and a 1 ms refractory period imposed. Conductances and equilibrium potentials were defined for leak (l), excitatory synapses (e), inhibitory synapses (i) and a tonic input serving to stabilize spiking ($tonic$) (Table 1).

Table 1.1 – Parameters of the spiking network model.

Parameter	Value	Parameter	Value
E_i	-90 mV	g_l	0.2 mS
E_e	0 mV	g_{tonic}	0.2 mS
E_l	-65 mV	τ_m	20 ms
E_{tonic}	0 mV	τ_e	10 ms
		τ_i	5 ms

Synaptic weights were randomly sampled from a lognormal distribution with location and scale parameters $\mu = -0.64$ and $\sigma = 0.51$. The resulting weights distribution had a mean of 0.6 and variance of 0.11, relative to the scale of the leak conductance. Since an important subset of

inhibitory projections onto excitatory cells tend to synapse on the soma and proximal dendrites[28] and are thus more potent, we enhanced *I* to *E* weights by 50%. We started each simulation by initializing membrane potentials to values drawn randomly from a normal distribution with a mean of -65 mv and a standard deviation of 5 mv. A pool of 50 Poisson neurons was used as input to the network. Poisson neurons spiked at 15Hz and were independently connected to excitatory units with P = 0.1 and 0.6 synaptic weight in the units of the leak conductance. The network was driven with the input pool for 50 ms and activity was allowed to continue for 100 ms, after which the simulation was terminated. This procedure was repeated over 100 trials with 10 different inputs. All simulations were carried out using the Brian Brain Simulator[29], with Euler's method for integration and time-steps of 1 ms.

Inference measures

Spikes were binned in five time resolutions (1, 5, 10, 20, 40 and 80 ms) into time-frames containing binary values, resulting in 150000, 30000, 15000, 7500, 3750 and 1875 bins respectively. We employed seven pair-wise measures of connectivity between neurons: lag count - abbreviated as count, lag correlation - abbreviated correlation, consecutive mutual information (MI), simultaneous MI, confluent MI, first-order transfer entropy (TE (k=1)) and second-order transfer entropy (TE (k=1)). We thus consider a wide array of inference algorithms ranging in sophistication.

For each pair of neurons i,j we defined a binary variable $c_{ij}^{kl}(t)$, which evaluates to 1 if $i(t) = k$ and $j(t+1) = l$. For example:

[28] Markram et al., "Interneurons of the Neocortical Inhibitory System."
[29] Goodman and Brette, "The Brian Simulator."

(4) $$c_{ij}^{11}(t) = \begin{cases} 1 & i(t) = 1 \quad and \quad j(t+1) = 1 \\ 0 & otherwise \end{cases}$$

Lag count was then calculated as:

(5) $$count_{ij} = \sum_{t=1}^{T-1} c_{ij}^{11}(t)$$

where T is the number of time-bins.

Lag correlation between two spike-trains was calculated using the Phi coefficient:

(6) $$correlation_{ij} = \frac{\left[\sum_{t=1}^{T-1} c_{ij}^{11}(t) \cdot \sum_{t=1}^{T-1} c_{ij}^{00}(t)\right] - \left[\sum_{t=1}^{T-1} c_{ij}^{10}(t) \cdot \sum_{t=1}^{T-1} c_{ij}^{01}(t)\right]}{\sqrt{2(T-1)}}$$

We use three versions of mutual information; consecutive mutual information (cMI) between a pair of neurons i,j was calculated as:

(7) $$cMI_{ij} = \sum_{i(t)\in\{0,1\}} \sum_{j(t+1)\in\{0,1\}} p(i(t), j(t+1)) \cdot \log_2 \left[\frac{p(i(t), j(t+1))}{p(i(t)) \cdot p(j(t+1))} \right]$$

However, we note that binning spikes into longer time bins may result in the pre- and post- synaptic spikes being binned into the same bin. Thus, we also consider simultaneous mutual information (sMI):

(8) $$sMI_{ij} = \sum_{i(t)\in\{0,1\}} \sum_{j(t)\in\{0,1\}} p(i(t), j(t)) \cdot \log_2 \left[\frac{p(i(t), j(t))}{p(i(t)) \cdot p(j(t))} \right]$$

and confluent mutual information (conMI):

(9) $$conMI_{ij} = \sum_{i(t)\in\{0,1\}} \sum_{j(\hat{t})\in\{0,1\}} p(i(t), j(\hat{t})) \cdot \log_2 \left[\frac{p(i(t), j(\hat{t}))}{p(i(t)) \cdot p(j(\hat{t}))} \right]$$

$$\text{where} \quad j(\hat{t}) = \begin{cases} 1 & j(t) = 1 \quad OR \quad j(t+1) = 1 \\ 0 & otherwise \end{cases}$$

As a non-symmetric information theoretical measures we calculated transfer entropies under first- and second- order Markov models (TE1 and TE2, respectively) between every pair i,j of neurons:

$$(10) \qquad TE1_{ij} = \sum_{j(t), j(t+1), i(t) \in \{0,1\}} p(i(t), j(t+1), j(t)) \cdot \log_2 \left[\frac{p(j(t)) \cdot p(i(t), j(t+1), j(t))}{p(i(t), j(t)) \cdot p(j(t+1), j(t))} \right]$$

TE2 is similarly defined, with the information i(t) is providing about j(t+1) conditioned not only on j(t) but also on j(t-1).

Measure evaluation

In order to evaluate performance of individual measures and the combined ensemble we calculated the recruitment network for each model. The recruitment network (Fig. 1D-F) is the intersection between the connectivity matrix and those synapses that directly contribute to post-synaptic firing since these are the only synapses that can be inferred using spikes (see schematic; Chambers and MacLean 2016). We first defined the active network in a similar way to $count_{ij}$ (Equation 5), but modified $c_{ij}^{kl}(t)$ so that both consecutive and simultaneous time bins are considered:

$$(11) \qquad \hat{c}_{ij}^{11}(t) = \begin{cases} 1 & i(t) = 1 \quad and \quad [j(t) = 1 \quad OR \quad j(t+1) = 1] \\ 0 & otherwise \end{cases}$$

(12) $$active_{ij} = \sum_{t=1}^{T-1} \hat{c}_{ij}^{11}(t)$$

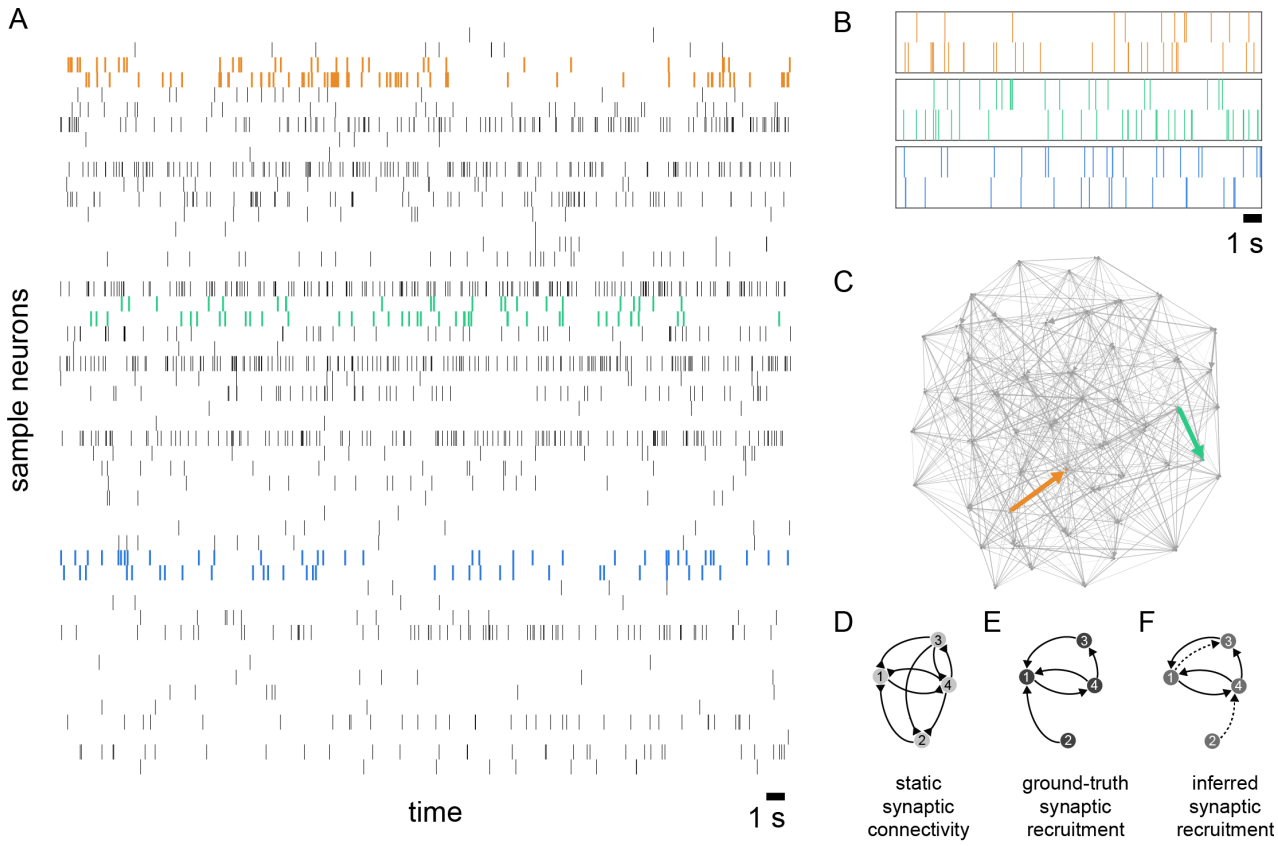

Figure 1.1 - Inferring synaptic connectivity from pairwise spike timing

A) Population spike raster for 50 random excitatory model neurons during 40 s simulated recording. Three representative pairs matched for firing rates are shown in color: strongly connected (orange), weakly connected (green), and unconnected (blue). Spikes were binarized at 20 ms time-bins. **B)** The same example pairs as in panel A during another 20 s of simulated recording. **C)** Ground truth synaptic connectivity for excitatory neurons shown in panel A. Edge width indicates weight. Arrows mark the strongly connected pair (orange) and weakly connected pair (green). Width was enhanced for visibility purposes. **D)** Schematic of a synaptic network among four active neurons. **E)** Synaptic recruitment is defined as lagged firing between pre- and postsynaptic pairs. Under the conditions of a given input, network state, and recording duration, not every synaptic connection recruits its postsynaptic partner to generate an action potential. **F)** Inferred synaptic connectivity (solid lines) mirrors the recruitment network, mapping propagating activity. Errors occur when inference algorithms fail to detect sites of synaptic recruitment (e.g. missing edge from neuron 2 to neuron 1), or assign putative connectivity (dashed lines) where there is none in truth.

The recruitment network was then computed as a binary matrix:

(13) $$recruitment_{ij} = \begin{cases} 1 & active_{ij} > 0 \quad AND \quad adj_{ij} > 0 \\ 0 & otherwise \end{cases}$$

where adj_{ij} is the adjacency matrix used to run the simulation.

The percent of connections retained in the recruitment network out of the static synaptic connectivity is described in Table 2. Previously we have reported that it is only possible to infer connections that are active and temporally proximal to an action potential in the post-synaptic neuron[30]. To reflect this fact we used the recruitment network as ground truth, and defined performance of an algorithm as the number of inferred connections at 80% true positive rate. This true positive rate was chosen based on survival curve analysis on three representative measures at all time resolutions (Fig 5). This definition provides a realistic test of performance as only those active connections that contribute to spiking in the post-synaptic neuron can be captured by an inference algorithm of any kind. A lower threshold and increased coverage under the same true positive rate means that as the refinement process progresses, inferred adjacency matrices become sparser because false alarms are being removed from those matrices.

Table 1.2 – Percent of connections retained in the recruitment network.
Data are shown for six simulated datasets binned at five time resolutions.

	5 ms	10 ms	20 ms	40 ms	80 ms
Simulated dataset 1	34.95%	40.82%	47.97%	53.49%	60.89%
Simulated dataset 2	41.41%	46.42%	53.18%	57.75%	65.22%
Simulated dataset 3	39.81%	45.63%	52.17%	56.70%	63.56%
Simulated dataset 4	45.72%	50.85%	58.57%	64.58%	72.32%
Simulated dataset 5	43.29%	49.38%	56.08%	61.04%	68.09%
Simulated dataset 6	42.32%	47.58%	54.36%	59.56%	67.14%

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

Ensemble weights and scores

We employed a simulated annealing strategy with cyclical dynamics to search over weighted linear combinations of individual measures. Before pooling, individual measures were normalized by their maximum value to provide a uniform scale. The final inputs into the simulated annealing algorithm were then computed as:

$$S(m)_{ij} = \sqrt{\left|\frac{norm\_residual(m)_{ij}}{\max(norm\_residual(m))}\right|} \cdot \mathrm{sgn}(norm\_residual(m)_{ij}) \quad (14)$$

Where $norm\_residual_{ij}$ is the final transformation of every measure, and defined in the results (Equation 20). sgn denotes the sign. m stands for the measure used, so each measure had its own $S_{ij.}$

The random search algorithm was used to train weights independently on each of six simulated networks using coverage at 80% accuracy as our objective function (Fig. 7A). This objective function proved to be pockmarked with many local maxima, motivating the use of repeated increment and cooling in the annealing process to avoid local maximum traps[31]. As step-size decreased, if no further gains were achieved, the search algorithm would occasionally jump to the last step-size that did yield improvements. If these larger jumps still failed to improve the performance, the jump size was further increased. Weights were learned across five repeats of training for each model and time-scale (Fig. 7A, 10 ms).

An ensemble score was then calculated for each pair of neurons as a linear weighted sum of $S(m)_{ij}$ (Equation 14):

[31] Kirkpatrick, Gelatt, and Vecchi, "Optimization by Simulated Annealing."

(15) $$Ensemble_{ij} = \sum_{m=1}^{M} w_m \cdot S(m)_{ij}$$

Where w_m are the weights of the measures found by the annealing algorithm, and M is the number of measures.

Comparison between inference measures

All comparisons were conducted after completing the regularization steps described in Results, i.e. on the normalized-residual adjacencies. Since summary statistics for adjacency matrices are impacted by edge-density, inferred adjacency matrices were thresholded to match sparseness before conducting any comparisons, isolating just the strongest relationships for each measure. To match sparseness, thresholding was performed at the 98th-percentile, leaving the top 2% of entries for each measure. Weights exceeding the inclusion threshold were mapped to 1 and those failing to reach inclusion threshold were mapped to 0. This pruning procedure allowed us to more clearly identify biases inherent to individual algorithms, and was not used to evaluate performance. Similarity between measures was assessed by vectorizing adjacency matrices and comparing the Euclidean distances separating each pair of measures. This comparison was performed independently for each simulated dataset. Reciprocity was quantified as the probability a randomly selected non-zero edge from neuron *i* to neuron *j* was accompanied by a non-zero edge from *j* to *i*. Local clustering for neuron *k* was quantified as the counted number of connected triangles including *k* divided by the number of possible triangles including *k*. In this formulation, a triangle must be composed of neuron *k* plus two immediate neighbors of *k*, without constraints on directionality. Thus, local clustering quantifies neighbor-of-neighbor relationships in the immediate neighborhood around *k*. Reciprocity and local clustering were aggregated by taking the mean over all edges and neurons, respectively, in the simulated dataset.

RESULTS

Simulated neuronal networks

Randomly connected recurrent networks (Fig. 1C) comprised of 1000 excitatory and 200 inhibitory LIF neurons showed persistent naturalistic activity after being driven by a sparse set of Poisson inputs for 50ms. Across six randomly connected networks, 99 ± 0.05% (mean ± SD) of neurons within the excitatory pool spiked at least once, and displayed sparse asynchronous-irregular firing (Fig. 1A and 1B). Firing rates were 1.66 ± 3.30 Hz (mean ± SD) and followed a log-normal distribution and single units showed irregular spiking consistent with experimental measures in cortical neurons[32,33], with inter-spike interval coefficient of variation 1.04 ± 0.20 (mean ± SD). The numbers of spikes used for inference in the six simulated datasets varied and depended on the size of the time bin used (Table 3). Using these simulated networks, which produced naturalistic spiking activity, we examined whether spiking activity within the network can be used to reveal underlying synaptic connectivity across the population. We employed sparse recurrent networks with random connectivity, constructed probabilistically according to naturalistic parameters (see Methods). Under some circumstances, non-random connectivity (e.g. local clustering) can facilitate inference of connections (Kobayashi & Kitano, 2013), so random synaptic connectivity may provide a more appropriate benchmark for comparing inference methods, particularly in assessing their propensity towards false positive errors. Because connectivity in neocortex is not random[34] the results presented here can be interpreted as a proof of concept of what is achievable when applying these inference approaches to experimental data. While it remains unclear how broadly results generated using randomly connected networks

[32] Softky and Koch, "The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs."
[33] Destexhe, Rudolph, and Paré, "The High-Conductance State of Neocortical Neurons in Vivo."
[34] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."

generalize, we have found that specific 'motifs' of higher order correlation found in random networks are also found in spiking data collected from neocortex (Chambers and MacLean 2016). There have been a number of studies that have examined whether correlation can be informative of causal connections, and it remains unclear how to best utilize the information provided by different inference approaches to uncover synaptic connectivity. As can be seen in Fig. 1B, synchronous or lagged events occur for both strongly and weakly connected neuron pairs, as well as by chance for unconnected neurons, rendering the task of interaction inference difficult. Here we introduce a series of refinements to traditional connectivity measures and demonstrate improvement in our ability to correctly identify monosynaptic excitatory connections using spiking across a recorded population of neurons.

Table 1.3 – Number of spikes used for inference.
Data represent mean ± SD for six simulated networks binned at 5 time resolutions.

	5 ms	10 ms	20 ms	40 ms	80 ms
Simulated dataset 1	135.13 ± 239.70	134.89 ± 238.94	130.12 ± 224.07	118.02 ± 189.30	102.41 ± 148.06
Simulated dataset 2	233.75 ± 513.30	233.34 ± 510.44	225.82 ± 471.20	205.27 ± 384.84	178.00 ± 286.54
Simulated dataset 3	264.25 ± 602.24	263.35 ± 595.56	251.24 ± 536.08	222.39 ± 425.01	187.51 ± 311.65
Simulated dataset 4	212.17 ± 340.30	212.02 ± 339.90	208.32 ± 339.47	196.14 ± 295.79	177.62 ± 245.71
Simulated dataset 5	282.00 ± 557.41	281.34 ± 554.79	269.80 ± 514.41	241.43 ± 424.08	205.72 ± 320.39
Simulated dataset 6	221.79 ± 424.11	221.58 ± 423.26	216.26 ± 404.11	200.09 ± 348.86	176.78 ± 274.40

Signed information theoretic measures - removing negative correlations

There are a number of approaches to inference. Some address the frequency of lagged or simultaneous spiking directly, such as the *count* method and the *correlation* method. For these measures, edge $e = (i,j)$ is positive only if there is the possibility that neuron i recruited neuron j to fire—that is, only if j becomes active after i at least once (Fig. 2A). In contrast, mutual information is related to lagged-firing in a more abstract sense. As a practical consequence, the entries of MI_{ij} can be large given any type of interaction between the pre- and post- synaptic neurons. Indeed, we

find that many pairs that were ranked highly under mutual information and transfer entropy had negative $correlation_{ij}$ scores (Fig. 2B and 2C). Unsurprisingly given the negative scores, the majority of those pairs were not monosynaptically connected, despite their high information theoretic scores. In order to account for interaction directionality consistent with causal synaptic interactions, we signed our mutual information and transfer entropy metrics on the basis of $correlation_{ij}$:

(16) $$signed(X)_{ij} = X_{ij} \cdot \mathrm{sgn}(correlation_{ij})$$

where X stands for the information theoretic measure considered (e.g. $signed(conMI)_{ij}$ is the signed confluent mutual information between neurons i and j).

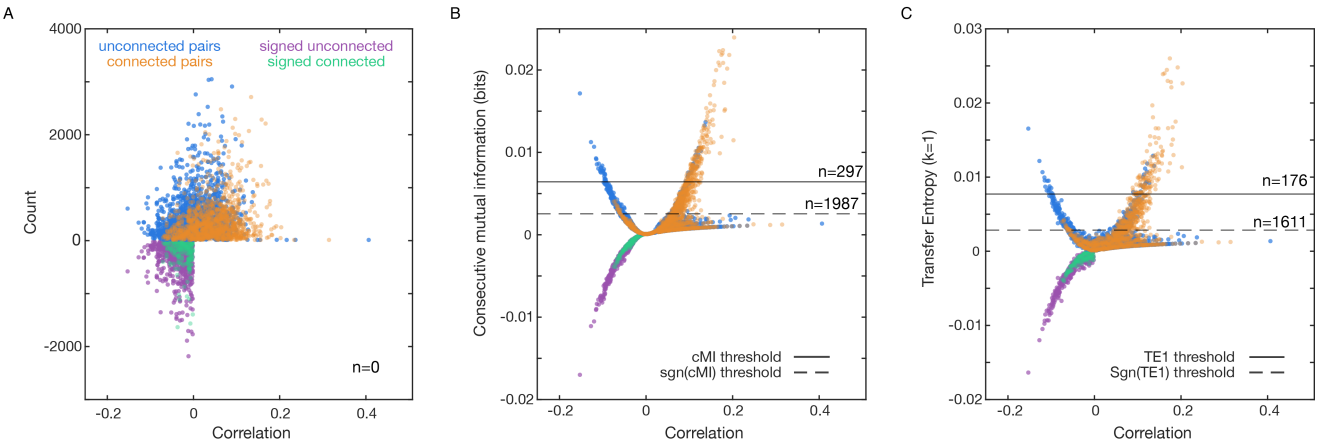

Figure 1.2 - Leveraging anti-correlations to isolate excitatory connections

Data in this figure are from a randomly chosen representative simulated network binned at 20 ms, and sub-sampled according to density for display purposes. A) Unconnected pairs often attained high lagged-count scores, but tended to exhibit negative-shifted lag-correlations compared to connected pairs. Signing lag-count scores on the basis of lag-correlation thus improved performance, although signed lag-count still fails to achieve sensitivity at the 80%-accuracy threshold. **B)** A subset of unconnected pairs exhibited high consecutive mutual information scores and strong negative lag-correlations. Signing consecutive mutual information entries on the basis of lagged correlations dramatically improved sensitivity at the 80%-accuracy threshold, increasing from 297 putative connections to 1987 putative connections in the representative model dataset. **C)** Transfer entropy is prone to the same errors, so that signing transfer entropy scores based on lag-correlation extends coverage of putative connections from 176 to 1611 directed pairs.

The signing procedure yielded gains in accuracy. These improvements are apparent when comparing the thresholds achieving 80% prediction accuracy for the raw versus signed metrics (Fig. 2B and 2C). We confirmed the intuition that negative interactions can confound the detection of excitatory connections using information theoretic methods but can be accounted for. It remains an open question whether negative information theoretic scores are indicative of inhibitory connections. Inference of inhibitory connectivity is especially challenging, because of the ambiguity in distinguishing inhibition *per se* from the absence of excitatory drive. In this work, we consider only the positive entries of $signed(X)_{ij}$, and denote those $pos(X)_{ij}$.

Removing additional spurious correlations

Removing negative correlations allowed us to identify and correct for one source of false positives. Nevertheless, it remained clear that a significant overlap between our true positive signal and false-positive background is still present. We observed that false positives sometimes appeared to span source and target nodes with high weighted out- and in-degrees. Indeed, functional interactions are known to be heterogeneous with a heavy tail[35,36,37], revealing indiscriminate patterns of spike time coordination. Since experimental evidence indicated that synaptic connectivity is sparse, we reasoned that neurons with extremely high weighted degrees reflected coordination in activity not arising from monosynaptic connections alone, but rather coordinated population dynamics. For the purpose of identifying likely synaptic pairs, these are 'background' spurious correlations. We estimated the magnitude of these background correlations, which depended on pre- and post-

[35] Shimono and Beggs, "Functional Clusters, Hubs, and Communities in the Cortical Microconnectome."
[36] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry," August 28, 2013.
[37] Nigam et al., "Rich-Club Organization in Effective Connectivity among Cortical Neurons."

synaptic identity, and removed them from inferred weights.

Previous work employing mutual information to infer protein interactions had shown that removing spurious correlations by linear regression refined inference and preferentially identified residues known to interact physically [38]. Inspired by this work, we asked whether a similar correlation existed in the neuronal interactions of our networks. Noting that positive values of information theoretic scores have highly skewed distributions, we re-expressed the measures with the exception of count due to the discrete nature of the metric. Re-expression was performed by Tukey's ladder of power [39] which finds the exponent that minimizes a distribution skewness:

$$a = \underset{a}{\arg\min}\ skewness(pos(X)^a) \tag{17}$$

Each measure was re-expressed accordingly:

$$redist(X) = pos(X)^a \tag{18}$$

Having re-expressed the scores, we next calculated the background signal for each pair of neurons. This was achieved by averaging the scores of the pre- and post- synaptic neurons when partnered with every other potential post- and pre-synaptic neuron in the network, respectively. This is equivalent to taking the mean across columns for the pre-synaptic neuron, and across rows for the post-synaptic neuron, excluding the partner neuron currently examined. We multiplied those mean scores for the pre- and post- synaptic neuron, denoted $background_{ij}$:

$$background_{ij} = \left\langle redist(X_{i,1:N-\{j\}}) \right\rangle \cdot \left\langle redist(X_{1:N-\{i\},j}) \right\rangle \tag{19}$$

[38] Little and Chen, "Identification of Coevolving Residues and Coevolution Potentials Emphasizing Structure, Bond Formation and Catalytic Coordination in Protein Evolution."
[39] Tukey, "Exploratory Data Analysis."

where $\langle ... \rangle$ is the mean.

A high average score, while potentially indicative of wide-reaching polysynaptic influence across the network, is unlikely to reveal monosynaptic connections. Indeed, $background_{ij}$ was found to be highly correlated with $count_{ij}$ (Fig. 3A), and only to a lesser extent with information theoretic measures, although correlations were still significant (Fig. 3B and 3C).

In order to remove the influence of this background signal we calculated the residual of each measure, $residual(X)_{ij}$, over $background(X)_{ij}$ by linear regression. We found improved coverage with the residual scores over the re-expressed scores, with up to 8-fold increase (93% increase on average) in connections uncovered with information theoretic measures corrected in this manner (Fig. 3D-F).

Scaling residuals to account for heteroskedasticity

Although accounting for the tendency of the neurons to participate in many interactions improved detection, pairs with high background signal were more dispersed around the regression line, especially for $count_{ij}$. In other words, the variance of $residual(X)_{ij}$ was dependent on $background(X)_{ij}$, hence displaying heteroskedasticity. This is evident from examining the residuals plotted against ranked $background(X)_{ij}$ (Fig. 4A-C). These differences in variance confound detection since pairs with high background signal are more likely to exceed selection thresholds. We found that heteroskedasticity affected connected as well as unconnected pairs, thus ruling out simple thresholding.

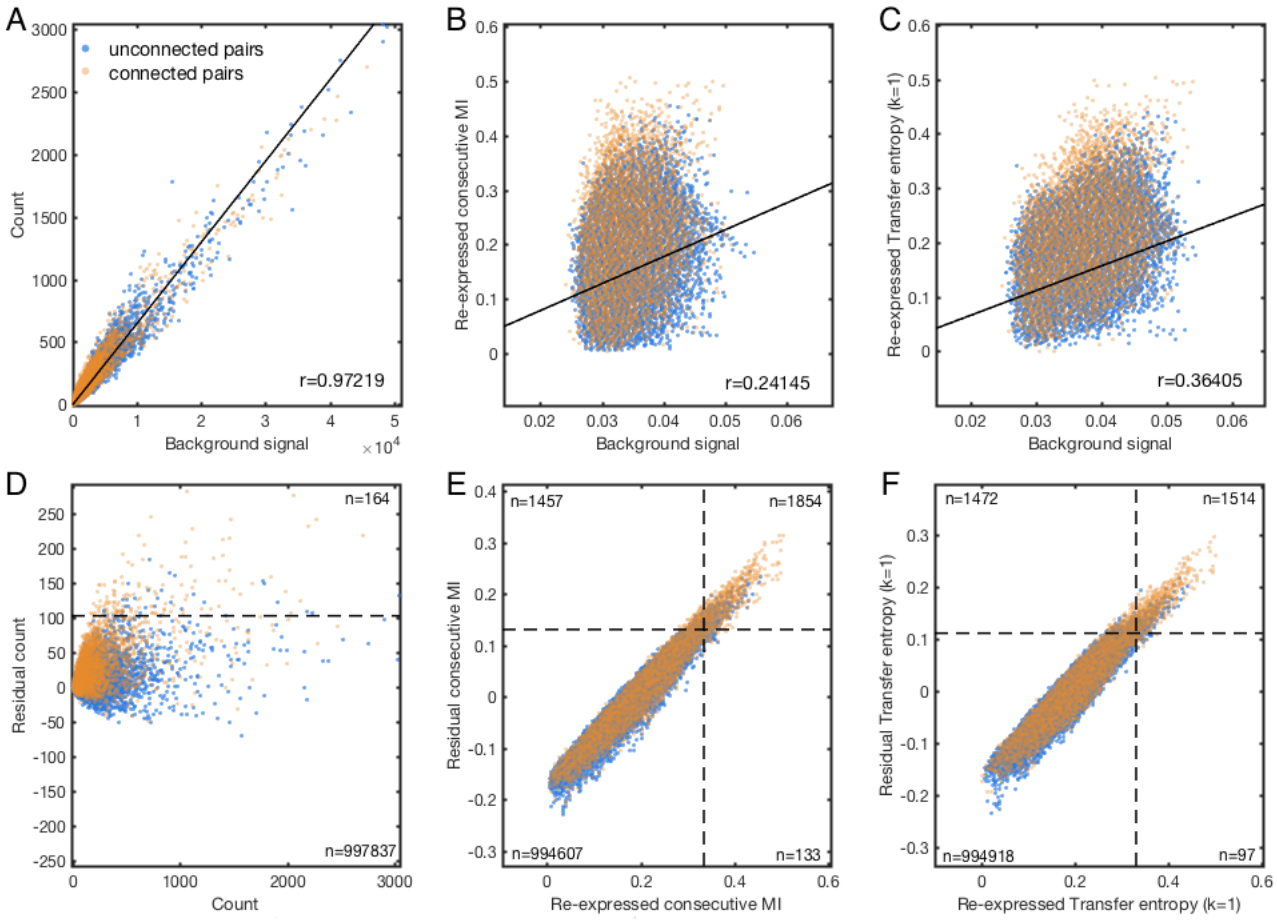

Figure 1.3 - Removal of mean background timing relationships improves detection of synaptic pairs

Data in this figure are from a randomly chosen representative simulated network binned at 20 ms, and sub-sampled according to density for display purposes. **A)** Linear regression revealed a strong background component in the C*ount* measure, reflecting a tendency for strong timing relationships to appear in tandem at select model neurons, encompassing both connected and unconnected pairs. **B)** After re-expressing global weights for each measure to approximate normality, a weaker but significant background signal was revealed for the *consecutive Mutual Information* measure. **C)** Background signal manifested somewhat more strongly for the *Transfer Entropy (k=1)* measure. **D)** Removal of mean neuron-wise background signal improved performance for the *Count* measure, but unconnected pairs with high residual scores remained a serious obstacle to high performance. **E)** For the *consecutive Mutual Information* measure, removal of background signal improved coverage at the 80% accuracy threshold from 1987 (two right quadrants) to 3311 (two top quadrants) putative pairs. **F)** For the *Transfer Entropy (k=1)* measure, removal of background improved coverage at the 80% accuracy threshold from 1611 to 2986 putative pairs.

To adjust for heteroskedasticity, we Z-normalized $residual(X)_{ij}$ by the geometric mean of the pre- and post- synaptic neurons standard deviations. To avoid inflating low variances by

dividing by small values, normalization was limited to a minimum divisor. We denote the Z-normalized scores by $norm\_residual(X)_{ij}$

(20) $$norm\_residual(X)_{ij} = \frac{residual(X)_{ij}}{\sqrt{\max(\varphi_{ij}, \varphi_{cutoff})}}$$

where:

(21) $$\varphi_{ij} = \sigma(residual(X)_{i,1:N-\{j\}}) \cdot \sigma(residual(X)_{1:N-\{i\},j})$$

and σ denotes the standard deviation.

$$\varphi_{cutoff} = median(\varphi_{ij})$$

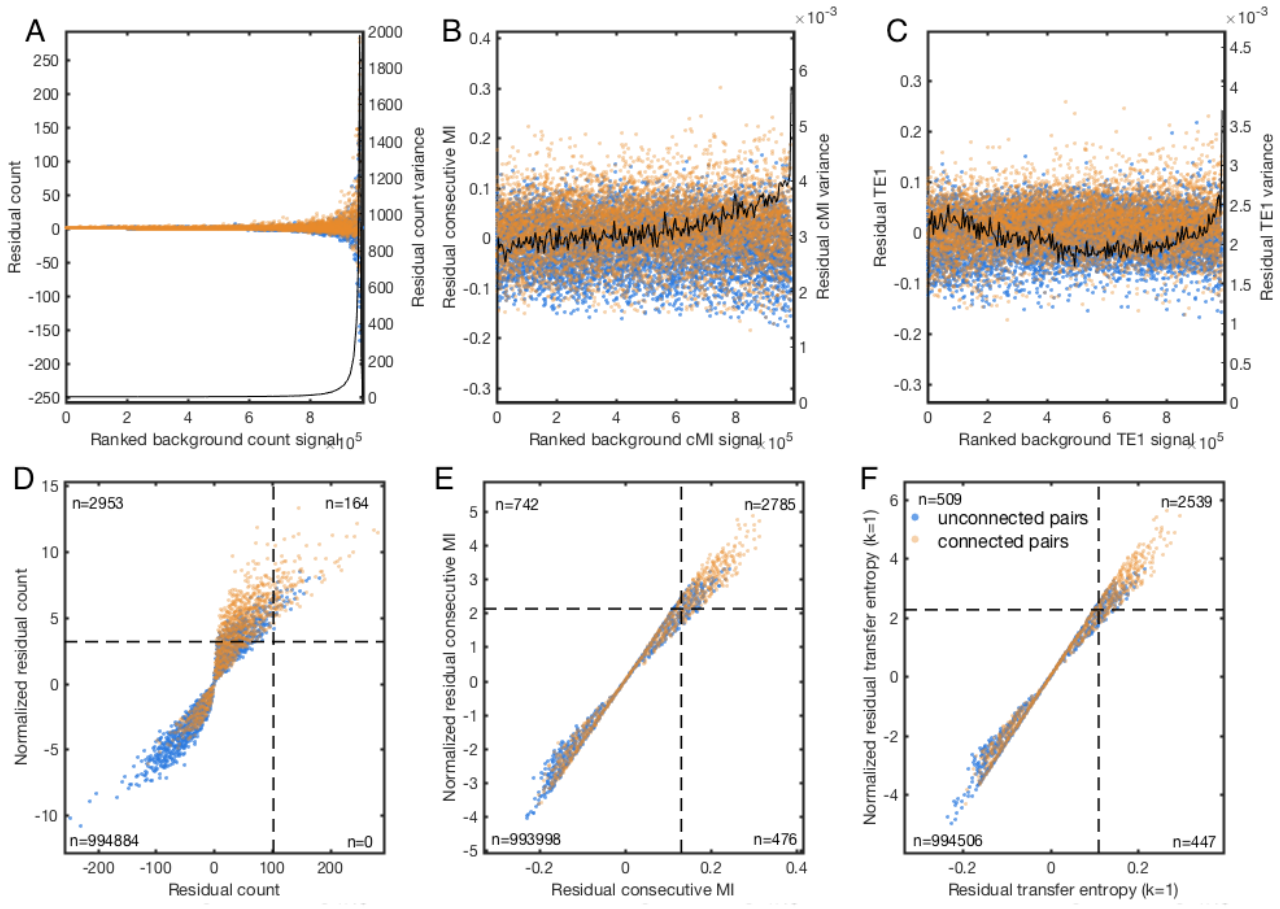

Figure 1.4 - Rescaling neuron-wise residual score variance

Data in this figure are from a randomly chosen representative simulated network binned at 20 ms, and sub-sampled according to density for display purposes. **A)** The *Count* measure exhibited strong scaling of residual variance with mean background signal, negatively impacting performance after thresholding. **B)** *Consecutive Mutual Information*

Figure 1.4 – continued

exhibited modest heteroskedasticity. **C)** Residual variance for *Transfer Entropy (k=1)* proved to be non-monotonically associated with mean background signal, with elevated variance among both the lowest and highest regularized scores. **D)** Z-normalization dramatically improved *Count* performance at the 80% accuracy threshold, increasing coverage of putative connected pairs from 164 to 3117. **E)** Z-normalization improved *consecutive Mutual Information* coverage from 3261 to 3527 putative connected pairs. **F)** *Transfer Entropy (k=1)* was not notably impacted by Z-normalization, increasing coverage from 2986 to 3048 putative pairs.

Accounting for scaled variance in background timing relationships, this refinement further improved the coverage of measures (Fig. 4D-F). The result was particularly dramatic for the count measure. Its success is particularly surprising since $count_{ij}$ was initially a poor indicator of underlying connection. We suggest that z-scored count estimates provided a simple and powerful first-pass approach to synaptic inference. The performance of this final stage of refinement across time-scales and accuracy thresholds is shown in Fig. 5A-F). As we previously reported (Chambers & MacLean, 2015), 1 ms lagged relationship between neurons is less informative compared with inference based on longer time-bins due to the time constant of synaptic integration. Owing to this finding we used time resolutions varying between 5 and 80 ms for all subsequent analyses.

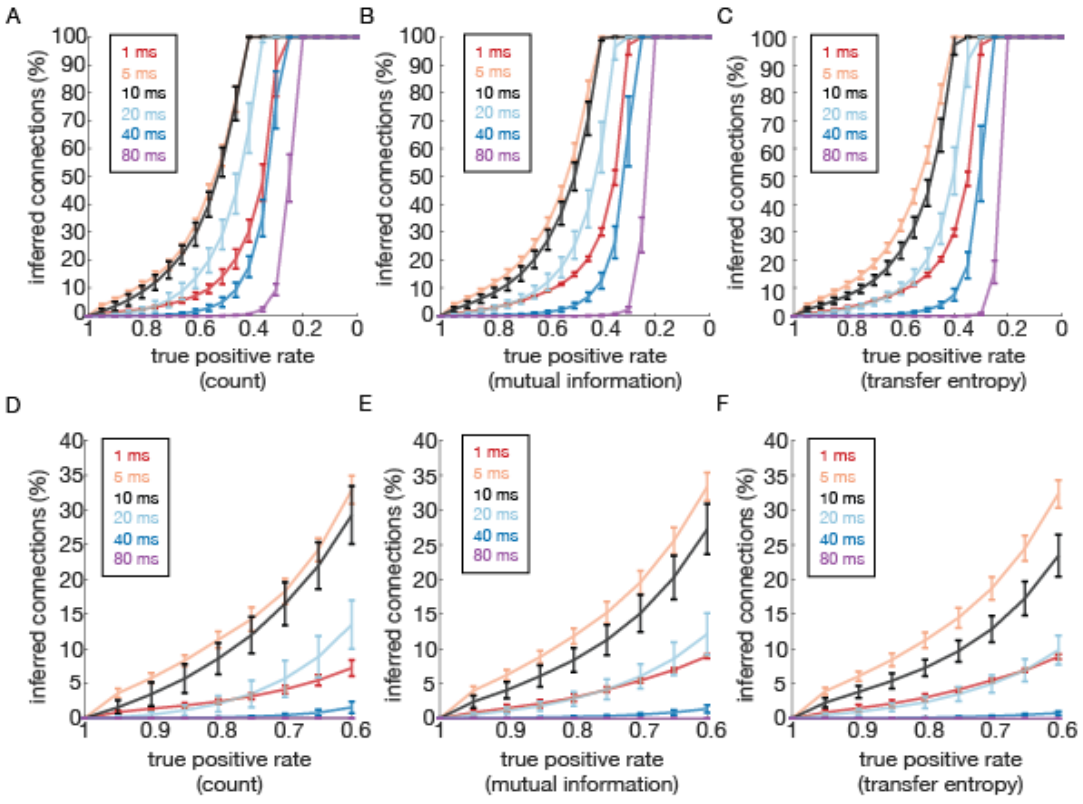

Figure 1.5 - Survival curves for *norm_residual(count)*$_{ij}$, *norm_residual(cMI)*$_{ij}$ and *norm_residual(TE1)*$_{ij}$ for multiple time resolutions and accuracy rates

A) Count perform best at 5 and 10 ms, revealing almost 50% of connections in the recruitment network. **B)** consecutive MI does similarly to count, with 5 and 10 ms time resolution achieving the best performance. **C)** Same as in A and B for Transfer Entropy (k=1). **D-F)** Zoom-in of A-C, respectively, showing true positive rates from 0.6 to 1. Note that inference algorithms calculated with 1 ms time-bins display performance in par with 20 ms time-bins. Mean and standard deviation across six simulated datasets are illustrated throughout.

Average performance gains at 80% accuracy at each stage of refinement collapsed across models for *count*$_{ij}$, *cMI*$_{ij}$ and *TE1*$_{ij}$ are shown in Fig. 6. The largest improvement to information theoretic measures resulted from accounting for the interaction sign, whereas coverage for *count*$_{ij}$ increased mainly due to Z-normalization of the residuals, bringing *norm_residual(count)*$_{ij}$ detection in par with *norm_residual(cMI)*$_{ij}$ and *norm_residual(TE1)*$_{ij}$. We next investigated whether statistical differences in the collections of predicted synaptic pairs persisted after signing, adjusting for background timing relationships, and selection by thresholding.

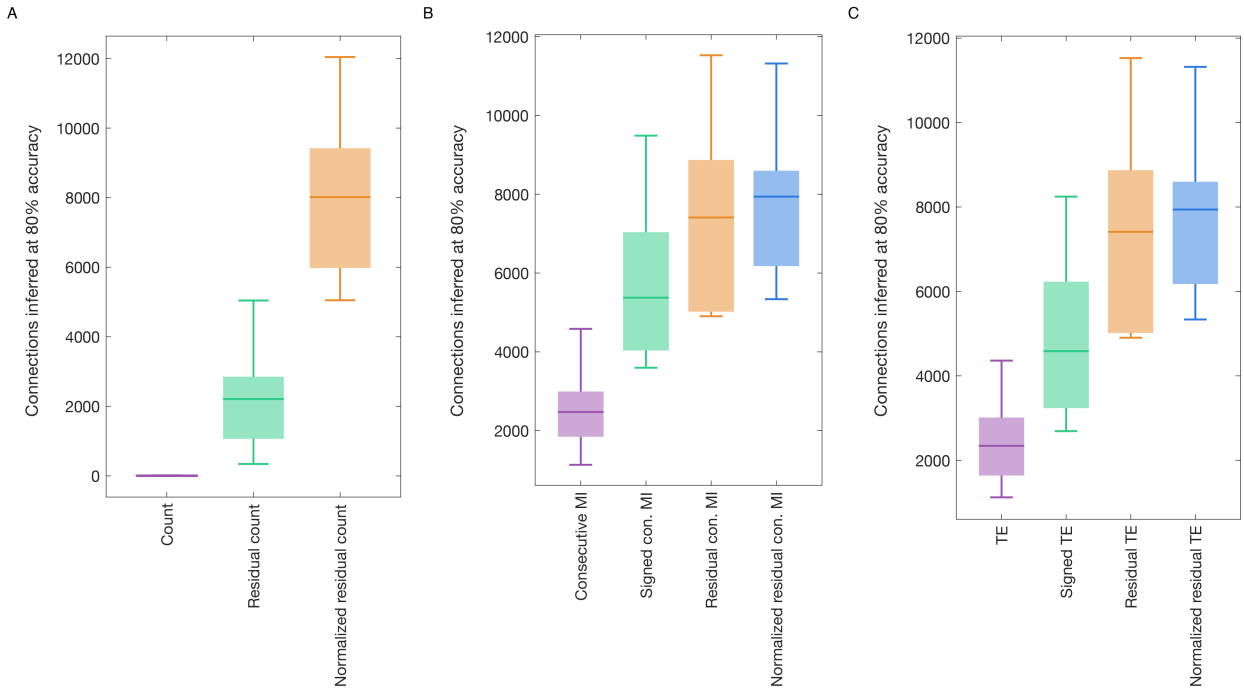

Figure 1.6 - Comparing gains in the regularization pipeline

A) Inference based on raw *Count* initially achieved zero coverage at the 80%-accuracy threshold. Removal of background signal improved coverage substantially, and the greatest gains resulted from z-normalization to compensate for heteroskedasticity. After regularization, *Count* performed as well as the best other individual inference algorithms. **B)** Regularization was also beneficial for the *consecutive Mutual Information* measure, with the greatest gains achieved by signing raw scores to distinguish positive timing relationships from negative ones. **C)** *Transfer Entropy (k=1)* exhibited similar improvements during regularization, benefiting from signing and removal of background signal.

Comparing similarity and temporal preferences across measures

We compared the collection of strongest pairwise relationships for each regularized inference measure. Thresholding was performed independently for each measure to yield sparseness-matched binary subgraphs. We first compared the similarity of detected synaptic connections between each pair of measures. Different measures of pairwise timing statistics highlighted non-identical, overlapping collections of putative synaptic pairs (Fig. 7). Qualitatively, L2 distances between measures were stable across simulated datasets (Fig. 7A and 7B). Interestingly, *count* and *simultaneous mutual information (sMI)* were most dissimilar, reflecting sensitivity to different

temporal structure. Consistent with this interpretation, the most similar measures were *correlation*, *consecutive MI*, and *TE1*.

To investigate further, we compared the performance of each inference metric across timescales. The mean and standard error of coverage and relative coverage are shown in Fig. 7C and 7D. An interesting trend was revealed: At high temporal resolution, the measures focusing on consecutive time-bins, namely *count* and *consecutive MI* performed best. But as time-resolution decreased, optimal performance shifted increasingly towards measures that combine consecutive and simultaneous time-bins, such as *confluent MI* and *TE2*. For small bins, synaptic integration and recruitment often straddle time-bin borders; whereas for large bins, a given synaptic interaction is more likely to occur entirely within single time-bins. In addition to choice of statistical measure, performance depends on the correspondence between relevant timescales of synaptic integration versus the timescale of spike binning[40].

Of all the measures, *simultaneous MI* was unique in that it contains only information on the spikes that occur in the same time-bin (with no consecutive time-bin information). One might hypothesize that, with large time-bins, *sMI* would therefore best detect synaptic interactions. However, its inherent symmetry gave rise to frequent errors in directionality. These errors in assuming bidirectional connectivity prevented it from significantly crossing the 80% accuracy threshold, and it thus achieved zero coverage at this cutoff. While *sMI* is fully symmetric in its raw form, the normalization process removed this symmetry, which explains how the refined score achieved non-zero coverage at the largest time scale.

Comparing topological preferences across measures

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

We hypothesized that the inferred subgraphs of synaptic connections might differ in systematic ways depending on the algorithm employed and that these differences would be detectable by comparing the topological organization of the sub-graphs. If inferred networks differed in the lagged relationships that they were sensitive to it was possible they contained complementary information about the location of true connections. In previous work we had found that specific topological motifs found in simulated model activity were also found in experimental data collected from somatosensory cortex, demonstrating that higher order dynamical organization has the potential to generalize across network construction (Chambers and MacLean 2016). We compared estimates of reciprocal connection probability and local clustering to understand if different measures made similar estimates of these quantities (Methods). Inferred topologies were characterized by a diversity of reciprocity and local clustering (Fig. 7E-H). Unsurprisingly, the *simultaneous MI* measure exceeded all other measures in the level of reciprocity represented among its strong entries, since through symmetry it tends to predict bidirectional connectivity. Note that ground-truth reciprocity is 0.2. In contrast, measures sensitive to time lagged statistical relationships tended to be characterized by lower levels of reciprocity (Fig. 7E and 7G). Since pairwise reciprocity sets a lower bound expectation for local clustering, it is not surprising that measures followed a similar rank ordering for both estimates, although fractional differences were smaller for estimates of local clustering. Emphasizing that the two metrics are related but distinct, we note that *count* estimates of local clustering were relatively high in relation to its estimate of reciprocity. Overall, inferred topologies exhibited non-identical statistical features depending on the algorithm employed.

We next tested whether these differences were dominated by detection errors, while encompassing statistically similar subsets of true synaptic connections. To investigate, we repeated

the analysis above for true positives only. After this step, inferred topologies were matched in sparseness by thresholding and binarized to prevent any uncontrolled differences in edge density. These two steps, thresholding and binarization were conducted solely for these comparisons (Methods). After pruning false positives from inferred topologies, estimates of reciprocity obeyed a qualitatively similar rank ordering (Fig. 7G). *sMI* continued to be dominated by reciprocal edges even after pruning false positives, exhibiting a strong selective preference for bidirectional synaptic pairs. Although it does not faithfully represent the groundtruth statistics of synaptic recruitment (investigated below), this feature could be useful in applied experimental settings: e.g. for targeting multi-cellular patch clamp recordings to find reciprocal pairs, potentially generating higher experimental yields. In contrast, after pruning false positives *sMI* exhibited far lower local clustering (Fig. 7H), revealing that false inference of reciprocal connectivity made a misleading impact in quantifying local clustering. The *count* subnetwork identified connections related by high local clustering. *Correlation* and *TE1* measures identified synaptic pairs less likely to be bidirectionally connected or clustered tightly together. Since different inference measures appeared to prefer distinct subsets of the synaptic network, we next tested whether their heterogeneous strengths could potentially be pooled to yield higher sensitivity than was achievable with any measure alone.

Ensemble approach for combining measures

We found that each inference algorithm revealed overlapping but distinct sets of causal connections with different biases. This diversity suggested that an ensemble approach, leveraging complementary sources of information across measures, could potentially improve upon inferences based on any single measure. We employed a stochastic search strategy (Methods) to find a weighting scheme for combining measures. The weights we obtained largely paralleled the

independent performance of the component measures (Table 4), with consecutive binning being favored at high temporal resolution and confluent or simultaneous binning being favored at low temporal resolutions. Surprisingly, significant weight was assigned to *sMI* at larger exposures despite *sMI* not being able to achieve any real coverage on its own (see Fig. 7C and 7D). This result emphasized the utility of Ensemble approaches in cases in which low performing algorithms can still improve the ensemble performance.

Table 1.4 – Ensemble weights.
Data reported here as mean ± SD across six simulated datasets, each ran 5 times through simulated annealing.

	5 ms	10 ms	20 ms	40 ms	80 ms
Count	0.1411 ± 0.084	0.1990 ± 0.125	-0.0029 ± 0.116	0.0025 ± 0.051	0.0199 ± 0.037
consecutive MI	0.2782 ± 0.094	0.2396 ± 0.116	0.2098 ± 0.071	0.1335 ± 0.129	0.1260 ± 0.121
simultaneous MI	0.0035 ± 0.016	0.0080 ± 0.015	0.0591 ± 0.048	0.2496 ± 0.125	0.2607 ± 0.112
confluent MI	0.2058 ± 0.064	0.2260 ± 0.080	0.2487 ± 0.111	0.1841 ± 0.120	0.2696 ± 0.117
TE (k=1)	0.1084 ± 0.118	0.1116 ± 0.124	0.1484 ± 0.063	0.0884 ± 0.062	0.0323 ± 0.057
TE (k=2)	-0.1086 ± 0.036	0.0359 ± 0.066	0.0984 ± 0.082	0.1446 ± 0.081	0.1191 ± 0.092
Correlation	-0.1302 ± 0.046	-0.1052 ± 0.092	-0.1329 ± 0.083	-0.0555 ± 0.128	-0.0527 ± 0.108

Weights were pooled over all models for principal components analysis, to identify model-independent features of the score landscape. Weight covariance for 10 ms simulated temporal resolution is shown in Fig. 8B, with entries in the main diagonal masked out for visualization purposes. The measures *Count*, *confluent MI*, and *TE1* exhibited particularly high covariance, suggesting that they provide complementary information about synaptic connections. This observation is consistent with their different statistical preferences. *TE2* tended to covary negatively with these measures, suggesting it was being leveraged to disambiguate pairs without a true connection. The complementary information yielded by *TE1* versus *TE2* reiterated the power

of studying functional coupling at multiple time-lags[41]. The search over weight space is illustrated for the first two PCA dimensions with performance indicated by color (Fig. 8C). Although score was not factored explicitly into the dimensionality reduction, effective weightings clustered together as a function of the first principle dimension, suggesting a large basin of feasible weights. Diverse combinations of weights had the potential to pool measures productively.

[41] Ito et al., "Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model."

Figure 1.7 - Inference algorithms differ in their statistical preferences

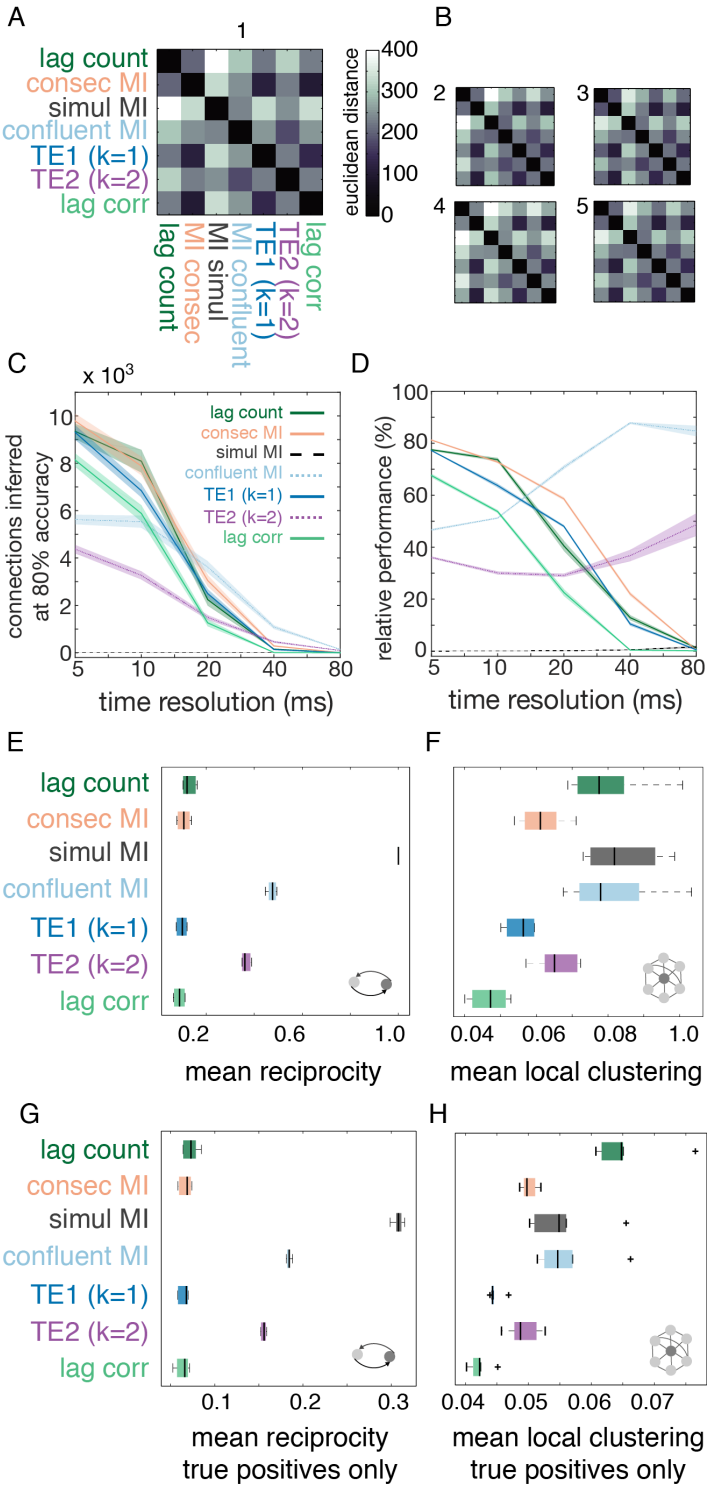

A) Comparing the strongest putative connections inferred by each regularized inference algorithm, after sparseness-matching and binarization, on the basis of their pairwise Euclidean distances. **B)** Stable similarity and dissimilarity relationships manifested across simulated datasets. **C)** Inference algorithms exhibited diverse levels of coverage at the 80%-accuracy threshold, with all algorithms except simultaneous Mutual Information performing better at high temporal resolutions. Regularized *Count*, *consecutive Mutual Information*, and *Transfer Entropy (k=1)* were the best performing measures given high temporal resolution, while *confluent Mutual Information* was the best performing measure in conditions of coarser temporal resolution. Lines and shading represent means and standard errors across six simulated datasets. **D)** The relative contribution of each inference algorithm to the coverage pooled across all algorithms revealed that measures considering more than one time-bin such as *confluent Mutual information* and *Transfer entropy (k=2)* were able to infer more connections at coarser time resolutions. Lines and shading represent means and standard errors across six models. **E)** Different levels of reciprocity were observed across collections of putative connections inferred from different regularized algorithms. Implausibly high reciprocity characterized *simultaneous Mutual Information*, due to its emphasis on bidirectional relationships. **F)** Levels of mean local clustering differed across regularized algorithms to a lesser degree, with *simultaneous Mutual Information* exhibiting highest levels of clustering. **G)** Comparison of reciprocity was repeated after excluding false positives and re-matching for sparseness. High levels of reciprocity continued to characterize *simultaneous Mutual Information*, revealing its strong bias for true reciprocally connected pairs. **H)** Comparison of local clustering was repeated after excluding false positives and re-matching for sparseness. Elevated local clustering

Figure 1.7 - continued

among simultaneous Mutual Information was revealed to be partially an artifact of its insensitivity to directionality. Collections of true inferred connections were most clustered for the lagged *Count* measure, and least clustered for the lagged *Correlation* measure.

Improved sensitivity with, and generalization of Ensemble inference

As we have previously reported, the recruitment network is characterized by elevated clustering in the local synaptic neighborhood[42]. The ensemble method recapitulates these features better than the best individual measures (Fig. 8D-F). The ensemble also exhibited improved sensitivity at the 80%-precision cutoff, with larger relative-improvements as sampling rates decreased. To examine the improvements offered by the Ensemble method, we plotted the coverage for the Ensemble score compared to its best performing component score (Fig. 9D). Across all simulated networks, all trials, and all exposures, the Ensemble method increased coverage, with absolute gains being around 1000 neuronal pairs over the best measure. Given the low component performance at large exposures, this represented a larger relative gain at these time-scales which is relevant for inferring connections using common experimental imaging techniques.

[42] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."

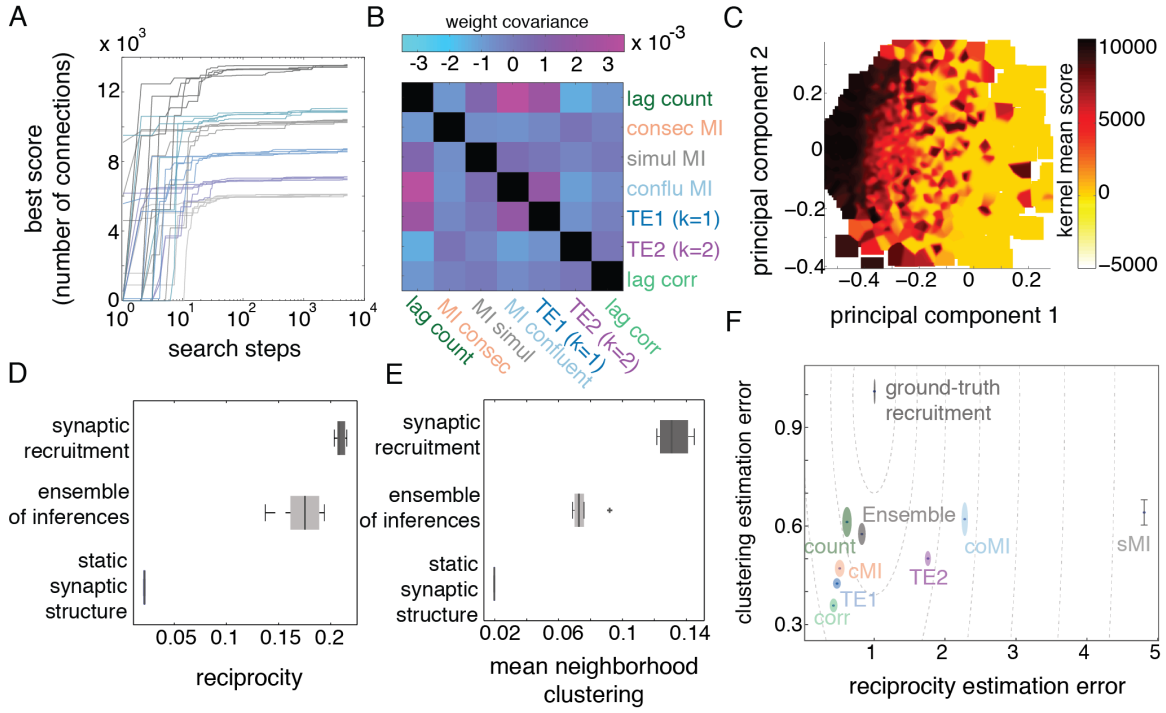

Figure 1.8 - Pooling individual measures to generate Ensemble predictions

A) Optimal weightings were obtained using random-walk search with simulated annealing on step sizes. Independent searches converged on similar best-scores for each simulated dataset, grouped by color. Appropriate weights were discovered early, but modest gains persisted throughout the search procedure. **B)** Covariance matrix quantifying common structure in weights throughout the search procedure, after concatenating all simulated networks and trials. Entries along the main diagonal were masked out for visualization purposes. **C)** Among pooled weights, good scores tended to cluster along one extreme of the first principle dimension, while robust to differences along the second principle dimension. Good solutions also appeared in isolated peaks throughout the solution space. Diverse weighted combinations yielded good performance in inferring synaptic connectivity. **D)** Ensemble-generated topologies closely reflected the ground-truth reciprocity statistics of synaptic recruitment in the simulated network. Synaptic recruitment preferentially occurred within a non-representative subset of the underlying random structural network. **E)** Like the ground-truth network of synaptic recruitment, Ensemble-generated topologies were characterized by elevated local clustering. However, like the best individual measures, Ensemble scores somewhat underestimated its true extent. **F)** With respect to reciprocity and local clustering, the two best characterizations of synaptic recruitment statistics were achieved by *Count* and *Ensemble* inference.

To test whether weights learned for one simulated dataset could transfer to other simulated networks, we computed the ensemble scores for each using weights learned for different datasets (transferred weights) and compared the performance to the datasets' performance with their own

weights (original weights). Simulated temporal resolution was matched for the transferred and original sets of weights. We found weights trained on one simulated network approached coverage after transfer to another network, with 97.32 ±0.54 % (mean ± s.d) of pairs inferred with original weights also inferred with transferred weights, at 5 ms time resolution (Fig. 9A and 9B). Generalization of weights across datasets depended on time scales in a similar manner to overall performance. Nonetheless, retained coverage was still impressive at 40 ms, with over 85% of inferred connections preserved (Fig. 8B), suggesting that simulations of realistic networks may be exploited to train ensemble weights for experimental data, even in cases in which temporal resolution is limited.

We also explored whether weights can be transferred across both simulated network and time-scales by examining retained performance after transfer from non-matching time resolution. Once again, the ensemble method was robust to weights generalization, with transferred performance remaining above 80% for time-scales that are similar yet not identical to the time scale the measures were originally computed with (Fig. 9C). For example, models binned at 10 ms with weights transferred from 5 and 20 ms models performed at 87.002 ±3.578 % and 86.744 ± 1.890 % of their original performance, respectively. Ensemble weights generalized across simulated datasets, and matching temporal resolution at least coarsely, was advantageous for transfer performance.

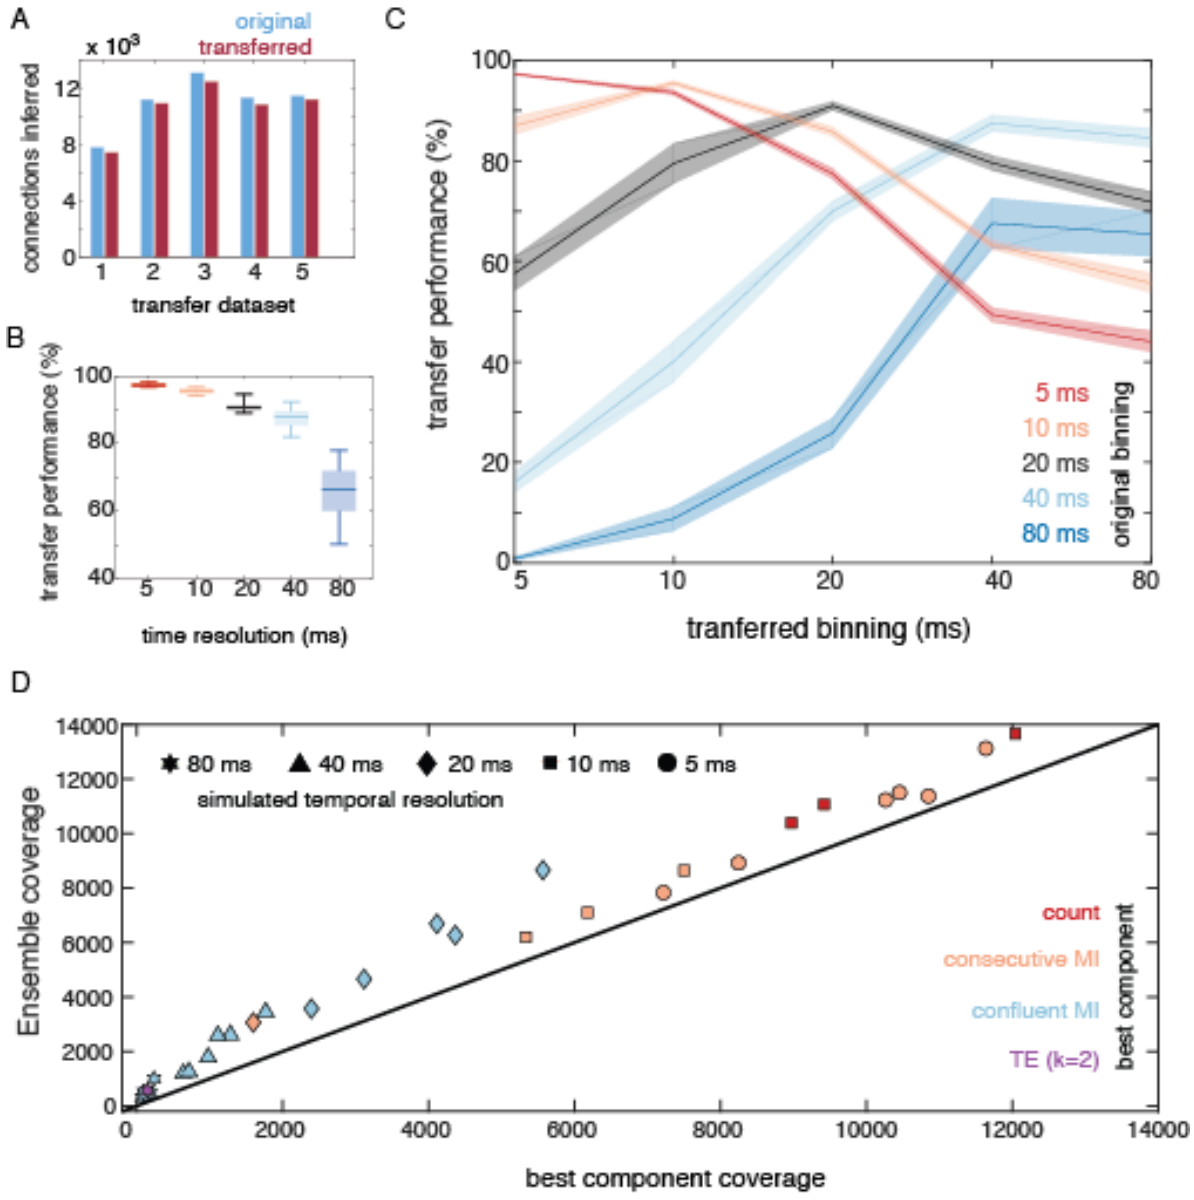

Figure 1.9 - Ensemble weightings generalize across simulated datasets and out-perform individual measures

A) Ensemble weightings were optimized for number of connections inferred using stochastic search (original performance, blue). Weights found for a random dataset with 5 ms time-resolution were used to construct Ensemble-generated adjacencies for each other dataset also binned at 5 ms (transfer performance, red). Transfer performance approached the original performance based on optimized weighting schemes. **B)** Given matched temporal resolution, transfer performance was nearly as good as original performance for fast temporal resolution. At slower time bins, transfer performance was somewhat less effective. **C)** Transfer performance was best when temporal resolution was matched between original and transferred recordings, and mismatch degradation was worst for weights originally learned from recordings with slow temporal resolution. Line and shading represent means and standard errors across six simulated datasets. **D)** Ensemble-generated topologies outperformed their best component measures across all bin sizes and model repetitions. Component measures with the best individual performance varied, but regularized *Count,*

Figure 1.9 - continued

consecutive Mutual Information, and *confluent Mutual Information* were often the best performing individual measures. Algorithms incorporating information at multiple timescales fared well, including *confluent Mutual Information* and *Transfer Entropy (k=2)*.

DISCUSSION

Within local cortical circuits, spiking activity propagates through synaptic networks in order to implement computation and shape behavior. Yet individual connections are weak in isolation, and patterns of coordination are complex and variable. Activity, or functional, mapping approaches, such as those presented here, infer probable synaptic recruitment patterns from statistical regularities in spike-timing. In this framework, statistical relationships are leveraged to predict synaptic connections, typically by thresholding to isolate the most reliably coupled pairs. Importantly, functional graphs, which succinctly summarize circuit dynamics, identify the synapses that are actively involved in the recruitment of post-synaptic neurons - i.e. those synapses which drive the post-synaptic neuron to threshold. Explicitly, only those connections revealed by (1) the timing of presynaptic action potentials, (2) the integrative properties of the post-synaptic neuron, and (3) the membrane potential of the post-synaptic neuron can ever be recovered by an inference approach. These are a limiting set of criteria that dramatically lessen the connections that can be recovered using these methods. Consequently, these approaches are not a realistic means to fully reconstruct a synaptic wiring diagram. Rather this is an approach that identifies synapses involved in implementing computation, and transmitting information during the specific epoch of dynamics that the graph summarizes – i.e. the recruiting network. As a result these connections are particularly interesting from a functional perspective. In this work, we compared methods for inferring excitatory synaptic connections, in order to understand the strengths and weaknesses of each. For mapping activity propagation through networks, we present two approaches to improve

the inference of excitatory connections: a regularization pipeline to improve the performance of individual inference algorithms, and an Ensemble stacking procedure that combines the best features of diverse measures.

Refinements of inference measures

To improve inference of excitatory connections we applied a sign to the information theoretic measures which disambiguates positive interactions from negative interactions. We note that negatively signed interactions could potentially provide insight into inhibition within the network, a long standing challenge in connectivity research, but also point out that disambiguating active inhibition from lack of excitation is a non-trivial endeavor. In this work we focus on the more tractable goal of mapping excitatory connections alone. Assigning valence to inferred relationships enhanced our recovery of excitatory connections.

Inference algorithms were further strengthened by removing a source of noise, reflecting background timing coincidences not reflecting monosynaptic interactions. After re-expressing measures to conform better to normality, we regressed out the mean component of this background signal by averaging over pre- and postsynaptic weighted degree. Since connectivity is sparse even among near neighbors, this averaging procedure was dominated by background influences specific to each neuron's firing rate and response profile. This tendency was not linear in magnitude across degrees, however, resulting in heteroskedasticity of residual scores. We corrected for the variability in standard deviation over the residuals with z-normalization, which further increased the accuracy of inferred interactions. This step was particularly effective for count, transformed by regularization into one of the single most effective indicators of connectivity. This approach is

similar in spirit to the normalized count procedure described in prior work[43]. These steps can be thought of in terms of informing inference algorithms not only by regularities within specific connections, but also by the statistics of the entire network. In this regard the current study differs from previous inference attempts in neuronal networks and, particularly, improves on previous procedures for thresholding. We suggest similarly inspired next steps would be to incorporate priors about higher-order structures such as motifs and clusters as well as accounting for log-normal distribution of weights[44] to further facilitate detection.

Ensemble Method

Different inference algorithms capture and summarize subtly different attributes of collective activity, and it is important to be mindful about these features when interpreting functional connectivity. Temporal resolution appears to be a particularly key design variable, and bin sizes of 25 – 50 ms are preferable for identifying synaptic connections, likely reflecting synaptic integration times (Chambers & MacLean 2015). Of course, even at much slower temporal resolution, the same algorithms can be useful for quantifying average timing relationships among active neurons. For mapping population dynamics, our results suggest that a host of productive statistical measures exist, which can be leveraged to infer likely patterns of synaptic recruitment. These measures are useful in isolation, and they become even more incisive in combination.

We found that each inference method isolated collections of putative underlying synaptic connections which are non-identical. This is an ideal situation for the application of Ensemble methods. Here we used linear combinations of multiple measures in order to improve predictions

[43] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."
[44] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."

of putative synaptic connections beyond the best single inference method. Non-linear combinations of measures have the potential to synergize further, leading to further gains in performance, but will face increasingly severe difficulties in generalization - a problem typical of classification in high-dimensional spaces given limited training data[45]. Ensemble approaches employ a diversity of methods for pooling and one common method is majority vote on individual classifiers[46]. However, Ensemble approaches are also frequently applied to real-valued outputs rather than binary classifier decisions[47]. One of the most common ways to leverage multiple classifiers in combination is known as bagging where the Ensemble score is the mean over all real-valued scores of individual measures. We demonstrate that better performance can be obtained in a weighted combination of the inference approaches. This latter framework for ensemble learning is known as stacking. In general, stacking algorithms are characterized by a pooling step, known as a combiner algorithm. As an example, high performers in the Netflix Prize employed linear regression and, later, backpropagation in a feedforward neural network for their combiner algorithm[48]. In order to optimize stacking weights, we employed an annealing random walk search strategy, a common approach in the field of Ensemble learning, including evolutionary approaches such as genetic algorithms and stochastic hill-climbing[49]. The question of how best to pool diverse connectivity inference algorithms remains an area necessitating further study. In practice, since recording conditions vary from one lab to another, experimenters with different experimental protocols than the ones studied here *in silica*, should re-train a combiner algorithm based on their own internal experimental design and selection of available inference algorithms with the

[45] Vapnik, *The Nature of Statistical Learning Theory*.
[46] Liaw and Wiener, "Classification and Regression by RandomForest."
[47] Mendes-Moreira et al., "Ensemble Approaches for Regression."
[48] Töscher and Jahrer, "The BigChaos Solution to the Netflix Grand Prize."
[49] Ruta and Gabrys, "Classifier Selection for Majority Voting."

understanding that each algorithm has different biases and the appropriate ensemble will depend on the expected underlying connectivity and time constant of synaptic integration.

The measures we studied in this work framed inference of synaptic connectivity as a binary classification algorithm. However, inference has the potential to also be cast as a regression problem using some continuous-valued measure of presynaptic influence (e.g. induced EPSP) which may lead to further gains, particularly in coverage. However, scarcity of ground-truth data and population-level coordination (obstacles facing any connection-inference approach) will make it challenging to move beyond binary classification to infer true synaptic connections. In addition to the algorithms we compared in this work, a host of other effective approaches have been described in the literature including particle methods[50] , GLMs[51,52] and Bayesian methods[53,54,55]. It may be that Ensemble predictions about underlying connectivity could be improved further by including these approaches in the pooling step.

Inference for circuit reconstruction/insights into information processing

Understanding how connection structure gives rise to synaptic recruitment remains a central goal for the study of neocortical circuits. Computation and behavior are enacted by propagating activity, so understanding synaptic recruitment mechanistically within active cortical networks is fundamental to the study of behaving animals. The importance of this issue was recognized

[50] Gerstein and Aertsen, "Representation of Cooperative Firing Activity among Simultaneously Recorded Neurons."
[51] Gerhard et al., "Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone."
[52] Zaytsev, Morrison, and Deger, "Reconstruction of Recurrent Synaptic Connectivity of Thousands of Neurons from Simulated Spiking Activity."
[53] Chambers and MacLean, "Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints."
[54] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."
[55] Mishchencko, Vogelstein, and Paninski, "A Bayesian Approach for Inferring Neuronal Connectivity from Calcium Fluorescent Imaging Data."

early[56,57], but technical obstacles limited its active study. In the last five years, progress is being made in this area through the study of functional relationships in active populations. In other words it is not only whether or not a connection is present, but it is also key to consider which connections are coactive or otherwise functionally related, and causal to spike propagation in a given context.

A confluence of evidence argues that the function of an individual connection depends on its arrangement within the local synaptic neighborhood. For example, a given connection will make a different impact if it is isolated versus arranged within a local cluster[58]. In neural cultures, frequent ignition sites were associated with elevated local clustering, for which a model of convergent amplification was proposed[59]. Culture activity maps were characterized by elevated clustering and short mean path lengths[60]. Intriguingly, clustering motifs, reciprocity, and heavy-tailed weight distributions may emerge through self-organizing plasticity processes[61]. These non-random features have received high-levels of interest, and may be consistent with several global topological organization schemas[62]. In neocortical tissue from mouse sensory cortices, spontaneous lagged firing relationships were found to be characterized by elevated modularity and hierarchical features[63,64]. Similarly, multi-electrode array recordings were marked by rich club

[56] Gerstein and Perkel, "Simultaneously Recorded Trains of Action Potentials."
[57] Palm, "From Neural Dynamics to Cell Assemblies."
[58] Pajevic and Plenz, "The Organization of Strong Links in Complex Networks."
[59] Orlandi et al., "Noise Focusing and the Emergence of Coherent Activity in Neuronal Cultures."
[60] Pajevic and Plenz, "Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches."
[61] Miner and Triesch, "Plasticity-Driven Self-Organization under Topological Constraints Accounts for Non-Random Features of Cortical Synaptic Wiring."
[62] Vegue, Perin, and Roxin, "On the Structure of Cortical Micro-Circuits Inferred from Small Sample Sizes."
[63] Sadovsky and MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry," August 28, 2013.
[64] Gururangan, Sadovsky, and MacLean, "Analysis of Graph Invariants in Functional Neocortical Circuitry Reveals Generalized Features Common to Three Areas of Sensory Cortex."

structure[65] and broad degree distributions[66]. In the latter work, functional hub neurons played a crucial role in supplying inputs to computationally important neurons downstream. Related theoretical work suggests functional hubs may take on distinct roles shaped by their assortativity or disassortativity relationships[67]. Converging sources of evidence have identified generalizable non-random features within connected neural systems consistent with the idea that these are key features to consider when describing the flow of activity through the circuit. Beyond these higher order functional relationships we have previously shown that inference is biased toward stronger connections[68]. As a result inference methods sometimes pose difficulties in interpretation[69], and they are limited in accuracy as well as sensitivity to weak synaptic connections, which are crucial for realistic spiking dynamics[70,71]. Continued development of tools to understand the large-scale organization of synaptic networks is an important area for further investigation.

The function of individual connections also depends on recent dynamics of the local circuit. On short timescales, unexpectedly effective recruitment can arise when inputs to a neuron are temporally coordinated[72], e.g. through the channel-dynamics of action-potential generation[73] or through interactions with the dendritic arbor[74]. Higher-order connectivity, particularly fan-in clustering, may favor coordinated inputs in this way, supporting effective synaptic integration[75].

---

[65] Nigam et al., "Rich-Club Organization in Effective Connectivity among Cortical Neurons."
[66] Timme et al., "High-Degree Neurons Feed Cortical Computations."
[67] Piraveenan, Prokopenko, and Zomaya, "Assortative Mixing in Directed Biological Networks."
[68] Chambers and MacLean, *"*Multineuronal Activity Patterns Identify Selective Synaptic Connections under Realistic Experimental Constraints.*"*
[69] James, Barnett, and Crutchfield, "Information Flows?"
[70] Sadovsky and MacLean, *"*Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry,*"* August 28, 2013.
[71] Teramae, Tsubo, and Fukai, *"*Optimal Spike-Based Communication in Excitable Networks with Strong-Sparse and Weak-Dense Links.*"*
[72] Rossant et al., "Sensitivity of Noisy Neurons to Coincident Inputs."
[73] Fontaine, Peña, and Brette, "Spike-Threshold Adaptation Predicted by Membrane Potential Dynamics In Vivo."
[74] Major, Larkum, and Schiller, "Active Properties of Neocortical Pyramidal Neuron Dendrites."
[75] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."

On moderate timescales, short-term plasticity imposes additional complexity in predicting the function of individual connections, where synaptic efficacy depends in part on a hidden state[76]. In addition, diverse sources of neuromodulation can reorganize synaptic efficacy in real-time, in ways which are difficult to understand a priori. For example, circuit-level reorganization can occur in response to acetylcholine, impacting the salience of extrinsic inputs versus local recurrent drive[77].

The complexity of structure-function relationships in neocortex is daunting. Isolated structural features can take on unexpected functional roles in the context of the local active network. Yet out of the extreme diversity of mechanisms shaping cellular communication, stable statistical relationships in spike-timing emerge. Maps of synaptic recruitment are a promising avenue to summarize the complex effects imparted by these many interlocking subcellular processes. For progress to continue in understanding large active neural circuits, it is important that statistics of spike timing among large populations continue to be investigated actively by the neuroscience community.

[76] Buonomano and Maass, "State-Dependent Computations."
[77] Runfeldt, Sadovsky, and MacLean, "Acetylcholine Functionally Reorganizes Neocortical Microcircuits."

REFERENCES

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

Chambers, Brendan, and Jason N. MacLean. "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks." *PLOS Computational Biology* 12, no. 8 (August 19, 2016): e1005078. https://doi.org/10.1371/journal.pcbi.1005078.

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. https://doi.org/10.1152/jn.00429.2015.

Destexhe, Alain, Michael Rudolph, and Denis Paré. "The High-Conductance State of Neocortical Neurons in Vivo." *Nature Reviews Neuroscience* 4, no. 9 (2003): 739–51.

Endo, Wagner, Fernando P. Santos, David Simpson, Carlos D. Maciel, and Philip L. Newland. "Delayed Mutual Information Infers Patterns of Synaptic Connectivity in a Proprioceptive Neural Network." *Journal of Computational Neuroscience* 38, no. 2 (April 1, 2015): 427–38. https://doi.org/10.1007/s10827-015-0548-6.

Fast, Andrew, and David Jensen. "Why Stacked Models Perform Effective Collective Classification." In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference On*, 785–90. IEEE, 2008. http://ieeexplore.ieee.org/abstract/document/4781179/.

Fontaine, Bertrand, José Luis Peña, and Romain Brette. "Spike-Threshold Adaptation Predicted by Membrane Potential Dynamics In Vivo." *PLOS Computational Biology* 10, no. 4 (April 10, 2014): e1003560. https://doi.org/10.1371/journal.pcbi.1003560.

Freund, Yoav, and Robert E. Schapire. "A Desicion-Theoretic Generalization of on-Line Learning and an Application to Boosting." In *European Conference on Computational Learning Theory*, 23–37. Springer, 1995. http://link.springer.com/chapter/10.1007/3-540-59119-2_166.

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. https://doi.org/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, George L., and Donald H. Perkel. "Simultaneously Recorded Trains of Action

Potentials: Analysis and Functional Interpretation." *Science* 164, no. 3881 (1969): 828–30.

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." Edited by Olaf Sporns. *PLoS Computational Biology* 10, no. 7 (July 10, 2014): e1003710. https://doi.org/10.1371/journal.pcbi.1003710.

Hansen, L. K., and P. Salamon. "Neural Network Ensembles." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, no. 10 (October 1990): 993–1001. https://doi.org/10.1109/34.58871.

Hofer, Sonja B., Ho Ko, Bruno Pichler, Joshua Vogelstein, Hana Ros, Hongkui Zeng, Ed Lein, Nicholas A. Lesica, and Thomas D. Mrsic-Flogel. "Differential Connectivity and Response Dynamics of Excitatory and Inhibitory Neurons in Visual Cortex." *Nature Neuroscience* 14, no. 8 (August 2011): 1045–52. https://doi.org/10.1038/nn.2876.

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. https://doi.org/10.1371/journal.pone.0027431.

James, Ryan G., Nix Barnett, and James P. Crutchfield. "Information Flows? A Critique of Transfer Entropies." *Physical Review Letters* 116, no. 23 (June 9, 2016): 238701. https://doi.org/10.1103/PhysRevLett.116.238701.

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. "Optimization by Simulated Annealing." *Science* 220, no. 4598 (May 13, 1983): 671–80. https://doi.org/10.1126/science.220.4598.671.

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. https://doi.org/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): 109–24.

Kruskal, Peter B., Lucy Li, and Jason N MacLean. "Circuit Reactivation Dynamically Regulates Synaptic Plasticity in Neocortex." *Nature Communications* 4 (October 10, 2013). https://doi.org/10.1038/ncomms3574.

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 (September 2010): 615–27. https://doi.org/10.1038/nrn2886.

Lalanne, Txomin, Therese Abrahamsson, and P. Jesper Sjöström. “Using Multiple Whole-Cell Recordings to Study Spike-Timing-Dependent Plasticity in Acute Neocortical Slices.” *Cold Spring Harbor Protocols* 2016, no. 6 (June 1, 2016): pdb.prot091306. https://doi.org/10.1101/pdb.prot091306.

Liaw, Andy, and Matthew Wiener. “Classification and Regression by RandomForest” 2 (2002): 5.

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. https://doi.org/10.1152/jn.1997.78.3.1714.

Little, Daniel Y., and Lu Chen. “Identification of Coevolving Residues and Coevolution Potentials Emphasizing Structure, Bond Formation and Catalytic Coordination in Protein Evolution.” *PLOS ONE* 4, no. 3 (March 10, 2009): e4762. https://doi.org/10.1371/journal.pone.0004762.

Major, Guy, Matthew E. Larkum, and Jackie Schiller. “Active Properties of Neocortical Pyramidal Neuron Dendrites.” *Annual Review of Neuroscience* 36 (2013): 1–24.

Marbach, Daniel, James C. Costello, Robert Küffner, Nicole M. Vega, Robert J. Prill, Diogo M. Camacho, Kyle R. Allison, et al. “Wisdom of Crowds for Robust Gene Network Inference.” *Nature Methods* 9, no. 8 (2012): 796–804.

Markram, Henry, Maria Toledo-Rodriguez, Yun Wang, Anirudh Gupta, Gilad Silberberg, and Caizhi Wu. “Interneurons of the Neocortical Inhibitory System.” *Nature Reviews Neuroscience* 5, no. 10 (October 2004): 793–807. https://doi.org/10.1038/nrn1519.

Mendes-Moreira, João, Carlos Soares, Alípio Mário Jorge, and Jorge Freire De Sousa. “Ensemble Approaches for Regression: A Survey.” *ACM Computing Surveys* 45, no. 1 (December 7, 2012): 10:1-10:40. https://doi.org/10.1145/2379776.2379786.

Miner, Daniel, and Jochen Triesch. “Plasticity-Driven Self-Organization under Topological Constraints Accounts for Non-Random Features of Cortical Synaptic Wiring.” *PLOS Computational Biology* 12, no. 2 (February 11, 2016): e1004759. https://doi.org/10.1371/journal.pcbi.1004759.

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, 1229–61.

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. https://doi.org/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): 582–90.

Pajevic, Sinisa, and Dietmar Plenz. “Efficient Network Reconstruction from Dynamical Cascades Identifies Small-World Topology of Neuronal Avalanches.” *PLoS Comput Biol* 5, no. 1 (2009): e1000271.

Pajevic, Sinisa, and Dietmar Plenz. “The Organization of Strong Links in Complex Networks.” *Nature Physics* 8, no. 5 (2012): 429–36.

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. https://doi.org/10.1007/978-3-642-81792-2_12.

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): 5419–24.

Piraveenan, Mahendra, Mikhail Prokopenko, and Albert Zomaya. “Assortative Mixing in Directed Biological Networks.” *IEEE/ACM Trans. Comput. Biol. Bioinformatics* 9, no. 1 (January 2012): 66–78. https://doi.org/10.1109/TCBB.2010.80.

Rossant, Cyrille, Sara Leijon, Anna K. Magnusson, and Romain Brette. “Sensitivity of Noisy Neurons to Coincident Inputs.” *Journal of Neuroscience* 31, no. 47 (November 23, 2011): 17193–206. https://doi.org/10.1523/JNEUROSCI.2482-11.2011.

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. https://doi.org/10.1152/jn.00071.2014.

Ruta, Dymitr, and Bogdan Gabrys. “Classifier Selection for Majority Voting.” *Information Fusion*, Diversity in Multiple Classifier Systems, 6, no. 1 (March 1, 2005): 63–81. https://doi.org/10.1016/j.inffus.2004.04.008.

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

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. https://doi.org/10.1152/jn.00334.2011.

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. https://doi.org/10.1523/JNEUROSCI.1977-13.2013.

Schapire, Robert E. "The Strength of Weak Learnability." *Machine Learning* 5, no. 2 (June 1, 1990): 197–227. https://doi.org/10.1007/BF00116037.

Shimono, Masanori, and John M. Beggs. "Functional Clusters, Hubs, and Communities in the Cortical Microconnectome." *Cerebral Cortex* 25, no. 10 (2015): 3743–57.

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.

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. https://doi.org/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.

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. https://doi.org/10.1371/journal.pcbi.1004858.

Töscher, Andreas, and Michael Jahrer. "The BigChaos Solution to the Netflix Grand Prize," January 1, 2009.

Tukey, John W. "Exploratory Data Analysis," 1977. https://pdfs.semanticscholar.org/2f40/fd06bc9fd00a27437e14ed171e96e4fd9326.pdf.

Vapnik, Vladimir. *The Nature of Statistical Learning Theory*. Springer science & business media, 2013. https://books.google.com/books?hl=en&lr=&id=EqgACAAAQBAJ&oi=fnd&pg=PR7&dq=dimensionality+generalization+learning+machine&ots=g2M1ew8U23&sig=m0ir4Jli

LF3_FiWM04Fm3MBO9P8.

Vegue, Marina, Rodrigo Perin, and Alex Roxin. “On the Structure of Cortical Micro-Circuits Inferred from Small Sample Sizes.” *BioRxiv*, 2017, 118471.

Zaytsev, Yury V., Abigail Morrison, and Moritz Deger. “Reconstruction of Recurrent Synaptic Connectivity of Thousands of Neurons from Simulated Spiking Activity.” *Journal of Computational Neuroscience* 39, no. 1 (August 1, 2015): 77–103. https://doi.org/10.1007/s10827-015-0565-5.

Zhang, Xiujun, Juan Zhao, Jin-Kao Hao, Xing-Ming Zhao, and Luonan Chen. “Conditional Mutual Inclusive Information Enables Accurate Quantification of Associations in Gene Regulatory Networks.” *Nucleic Acids Research* 43, no. 5 (March 11, 2015): e31. https://doi.org/10.1093/nar/gku1315.

CHAPTER II

Network analysis of murine cortical dynamics implicates untuned neurons in visual stimulus coding

This work was previously published: Levy, M., Sporns, O., & MacLean, J. N. (2020). Network analysis of murine cortical dynamics implicates untuned neurons in visual stimulus coding. *Cell reports*, *31*(2), 107483.

ABSTRACT

Unbiased and dense sampling of large populations of layer 2/3 pyramidal neurons in mouse primary visual cortex (V1) reveals two functional sub-populations: neurons tuned and untuned to drifting gratings. Whether functional interactions between these two groups contribute to the representation of visual stimuli is unclear. To examine these interactions, we summarize the population partial pairwise correlation structure as a directed and weighted graph. We find that tuned and untuned neurons have distinct topological properties, with untuned neurons occupying central positions in functional networks. Implementation of a decoder that utilizes the topology of these functional networks yields accurate decoding of visual stimuli. We further show that decoding performance degrades comparably following manipulations of either tuned or untuned neurons. Our results demonstrate that untuned neurons are an integral component of V1 functional networks and suggest that network interactions contain information about the stimulus that is accessible to downstream elements.

INTRODUCTION

Neurons in sensory cortices collectively encode information about the external world. In primary visual cortex (V1) neurons are thought of as selective, or tuned, to an orientation or direction of a moving bar or drifting gratings if the neuron exhibits consistently increased firing rate in response to a direction or orientation[1]. Not all neurons are tuned to a particular statistical feature of a visual stimulus shown in an experiment, with untuned neurons comprising roughly 20-50% of pyramidal neurons in mouse V1 varying by lamina[2,3,4,5], and their role in visual coding remains understudied.

Imaging techniques allow us to densely sample large numbers of neurons in an unbiased manner, facilitating simultaneous examination of the full circuit response including both tuned and untuned neurons[6]. This is helpful because one way by which untuned neurons are hypothesized to contribute to coding is through their correlation structure with both tuned and untuned neurons[7]. Indeed, studies that include pairwise correlations have demonstrated superior decoding performance compared with decoders assuming independent units[8,9,10]. Theoretical studies have suggested that the correlation structure of a network of neurons can itself hold information about the stimulus, especially when the spatial decay of correlations is considered Josić et al[11]. Experimental evidence corroborates this postulation: noise correlations between direction

[1] Hubel and Wiesel, "Receptive Fields of Single Neurones in the Cat's Striate Cortex."
[2] Zariwala et al., "Visual Tuning Properties of Genetically Identified Layer 2/3 Neuronal Types in the Primary Visual Cortex of Cre-Transgenic Mice."
[3] Sun et al., "Thalamus Provides Layer 4 of Primary Visual Cortex with Orientation- and Direction-Tuned Inputs."
[4] Ringach et al., "Spatial Clustering of Tuning in Mouse Primary Visual Cortex."
[5] Niell and Stryker, "Highly Selective Receptive Fields in Mouse Visual Cortex."
[6] Olshausen and Field, "How Close Are We to Understanding V1?"
[7] Zylberberg, "Untuned but Not Irrelevant: The Role of Untuned Neurons in Sensory Information Coding."
[8] Shi, Niu, and Wan, "Effect of the Small-World Structure on Encoding Performance in the Primary Visual Cortex."
[9] Graf et al., "Decoding the Activity of Neuronal Populations in Macaque Primary Visual Cortex."
[10] Chen, Geisler, and Seidemann, "Optimal Decoding of Correlated Neural Population Responses in the Primate Visual Cortex."
[11] Josić et al., "Stimulus-Dependent Correlations and Population Codes."

selective neurons in macaque MT were found to depend on the direction shown, indicating that those correlations themselves can be tuned[12]. Finally, correlational couplings between neurons, including coupling between tuned and untuned neurons, can be used to predict neuronal single trial responses regardless of whether couplings represented tuned or untuned inputs[13]. Yet the extent to which pairwise correlations across a large and functionally diverse neuronal population may contribute to decoding remains unknown.

We generated functional networks (FNs) as a summary of network activity since FNs maintain neuron-specific labels while simultaneously capturing all pairwise correlations. Specifically, in this framework neurons are nodes and statistical dependencies in the activity between neurons are edges resulting in a weighted and directed matrix which can be structurally evaluated using graph theoretic tools[14,15] (for a review see [16]). Here we find that functional network topology is specific to a given visual stimulus. Untuned neurons inhabit central positions within the topology acting as functional hubs because of their propensity to form a rich-club of strong connections and their increased ranking in random walks. Using a two-stage model comprised of a generative and a decoding component we demonstrate that information about the stimulus is represented in the pattern of functional connections between tuned and untuned neurons. Hence, the analysis of functional networks, which naturally encompass circuit-wide interactions across multiple neuronal classes, provides an approach that unifies the neuron-centric and population-centric frameworks in visual system neuroscience.

---

[12] Ponce-Alvarez et al., "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons."
[13] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[14] Kotekal and MacLean, "Recurrent Interactions Can Explain the Variance in Single Trial Responses."
[15] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[16] Bassett and Sporns, "Network Neuroscience."

MATERIALS AND METHODS

Animals

Data collection was performed in accordance to the guidelines of, and approved by the Institutional Animal Care and Use Committee at the University of Chicago. Experimental animals were 8 mice (4 males, 4 females) between ages P84-P191, expressing GCaMP6s under the Thy-1 promotor[17]. Animals had ad libitum access to food and water. Information about the animals was previously described in Dechery and MacLean, 2018.

Cranial window surgeries and two-photon calcium imaging

Fully described in Dechery and MacLean, 2018. Briefly, a 3 mm diameter craniotomy over left V1, in which the anatomical location was verified using intrinsic signal imaging[18]. During imaging sessions, animals were head-fixed, awake and allowed to run voluntarily on a linear treadmill. A field of view (FOV) in V1 was found and compared against fiduciary markers from the images obtained by intrinsic signal imaging, and neurons were automatically detected by software written in house. To facilitate cell detection we divided the FOV into a 4X4 grid, zooming on each grid section individually. Neurons were then identified in each grid zoom, and their coordinates in pixel space were transformed back into the FOV coordinate system, enabling us to perform an accurate line-scan on the entire population. Two photon calcium imaging was performed with 910 nm excitation wavelength with a scanning rate of 25-33Hz, depending on the number of neurons imaged (72-347). Each frame is thus ~30 ms long.

[17] Dana et al., "Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging in Vivo."
[18] Kalatsky and Stryker, "New Paradigm for Optical Imaging."

Visual stimuli

We presented the mice with 8-10 repetition of the same movie, comprised of pseudo-random presentations (trials) of drifting gratings in 12 directions (0.04 deg/cycle, 2Hz), spaced 30 degrees apart. Drifting gratings (Stimulus) lasted 5 seconds and were interleaved with 3 seconds of uniform mean luminance screen (Gray). This provided 24-30 trials of each drifting grating direction.

Functional networks

Separate functional networks (FN) were inferred from data of stimulus and gray epochs (Fig. 2D). We trimmed fluorescence traces from stimulus epochs to use the first 2 seconds of grating presentation, and traces from gray epochs to use the last 2 seconds of mean luminance screen. This was done to focus on the initial response to the drifting grating and to lessen the likelihood that fluorescence values had not returned to baseline from the preceding stimulus epoch, respectively. The trimmed individual traces from each epoch type were then concatenated in the order they were presented, separately for stimulus and gray and in each movie, and then smoothed with a running average window of 10 frames to discard discontinuities inserted by the concatenation procedure. This resulted in grand stimulus traces and grand gray traces for each movie. We then computed partial correlation between every pair of neurons, partialling out the mean response of the neurons in the pair in all other movies and the mean response of the population within the same movie (Fig. 2D). Directionality was assigned to the partial correlation score by examining the peak (maximum value) in the cross-correlogram (see [19]).

[19] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."

Separate functional networks for each direction of drifting gratings were constructed by parsing out the traces of activity in response to trials of each of the 12 directions (Fig. 2C, E). Only the first 2 seconds of each trial were kept, and concatenated with other trails of the same direction. In most datasets we had 30 trials for each direction, and we concatenated every 5 consecutive trials together. Those grand traces were then smoothed as described above and partial correlation was computed, regarding every grand trace as a movie. As the majority of our analysis tools such as graph alignment, rich club analysis and random walks are designed for positive edges only, we discarded negative edges from all FNs for all analyses.

Graph alignment

We defined the graph alignment score A of each pair of networks M and $N$with k neurons as:

$$A = \frac{\sum_{i=1}^{k}\sum_{j=1}^{k} min\left(M_{ij}, N_{ij}\right)}{\sum_{i=1}^{k}\sum_{j=1}^{k} M_{ij} + N_{ij}}$$

We then normalized the alignment score using a distribution of alignment scores from 100 degree-sequence-preserving (randmio_dir.mat from [20]) permuted FN as follows:

$$norm(A) = \frac{A - \langle \hat{A} \rangle}{1 - \langle \hat{A} \rangle}$$

Where $\langle \hat{A} \rangle$ is the mean of the distribution from permuted FNs, so that alignment scores quantify similarity beyond what is expected by chance[21].

Rich club analysis

[20] Rubinov and Sporns, "Complex Network Measures of Brain Connectivity."
[21] Gemmetto et al., "Multiplexity and Multireciprocity in Directed Multiplexes."

We sorted neurons in descending order by their combined in and out degree, and then we sequentially considered degrees starting from the largest. For each degree threshold, d, we calculated the rich club coefficient as:

$$R_d = \frac{\sum_{k(i)\geq d}^{N} \sum_{k(j)\geq d}^{N} w_{ij}}{D^2 - D}$$

Where k is the total degree of the neuron, w_{ij} is the FN, and D is the count of neurons with $k \geq d$. We then performed 1000 degree-sequence-preserving (randomio_dir.mat from [22]) randomizations of the network, and calculated the rich club coefficient in each of them. Significance was computed as the probability of the rich-club coefficient of real data being larger than the coefficient obtained from the population of randomized FNs.

Random walks

The ranking P of each neuron $i \in N$ was initialized randomly in (0,1] and then calculated iteratively as:

$$P_i = (1-q)\sum_{j=1}^{N} \frac{P_j}{s_j^{out}} w_{ji} + qz_i + (1-q)z_i \sum_{j=1}^{N} P_j \delta\left(s_j^{out}\right)$$

where $s_j^{out} = \sum_{k=1}^{N} w_{jk}$, and w_{ji} are the edges from neuron j to neuron i and q is the damping factor[23]. We picked $q=0.1$ for all analysis except when we tested the robustness to this factor. This algorithm differs from PageRank because it allows assignment of fixed authority to each node in the network, denoted by z_i, which we set to 1 for all neurons and all analysis unless stated otherwise. The last term on the r.h.s corrects for neurons that do not have outgoing edges with

[22] Rubinov and Sporns, "Complex Network Measures of Brain Connectivity."
[23] Radicchi et al., "Diffusion of Scientific Credits and the Ranking of Scientists."

$\delta(x) = \begin{cases} 1 & x = 0 \\ 0 & otherwise \end{cases}$ and thus disappears if there are no such neurons in the network, which was the case for all but one FNs. The procedure repeats iteratively until the L2 norm between the ranking vector at iteration *k* and *k-1* is smaller than a threshold set to $1x10^{-6}$.

Two-phase decoder model

The decoder model was built in two phases: generation of artificial spike trains using a simple recurrent neural network model, and decoding direction using a feed-forward pattern recognition network. In the first phase, each of the 12 direction-specific FNs was used as a recurrent neural network. The probability of firing for each neuron was governed by:

$$p_i(t) = \left(1 + exp\left(-\sum_{j \neq i}^{N} \left(s_j(t-1) \cdot w_{ji} \right) - \sum_{k=1}^{M} \left(s_k(t-1) \cdot a_{ki} \right) \right) \right)^{-1}$$

once the probability of spiking $p_i(t)$ was calculated, $x \in [0,1]$ was drawn at random from a uniform distribution. Then spiking at time t, $s_i(t)$ was set as:

$$s_i(t) = \begin{cases} 1 x \leq p_i(t) \\ 0 x > p_i(t) \end{cases}$$

Each neuron i received excitatory and inhibitory inputs, described in the first and second terms of the exponent, respectively. The connectivity structure of N excitatory neurons w_{ij} was governed by the FNs inferred from data and frozen. In contrast, M inhibitory neurons were added artificially, since we did not image inhibitory neurons, to balance the activity in the network. Inputs from excitatory to inhibitory neurons were set according to fixed connection probability P_{EI} and the weight values were drawn from a lognormal distribution with parameters estimated from the FN. Weights from inhibitory to excitatory and inhibitory neurons (a_{ki}) followed a random wiring procedure with probabilities P_{IE} and P_{II}, respectively. The weights were again drawn from a

lognormal distribution estimated from the edge weights in the functional network, and then multiplied by a scaling factor, g.

The parameters $M, P_{EI}, P_{IE}, P_{II}$ and g were chosen by a grid search approach in which we set biologically plausible ranges for all the parameters[24], and then ran the model with all possible combinations of values within these ranges. We then examined the resulting dynamics in excitatory neurons for three properties present in real cortical dynamics: persistent activity[25], realistic firing rates[26,27,28] and asynchrony[29,30,31]. Persistent, untruncated simulated activity was evaluated by only examining parameters that produced spiking up to 140 time-frames, which are equivalent to 5 seconds of stimulus presentation considering our average scanning rate. Realistic firing rates were achieved by only considering parameters that produced excitatory firing rate within two standard deviations of the data firing rate as estimated from the fluorescence traces by the OASIS inference algorithm[32] . Finally, asynchrony was guaranteed by examining the simulated rasters by eye and discarding parameters that resulted in simultaneous, locked spikes in more than 20% of the population. This procedure to pick that parameters that satisfied those criteria was done separately for each FN, and on average across networks and datasets $M = 0.28N$, $P_{EI} = 0.39$, $P_{IE} = 0.49$, $P_{II} = 0.11$, and $g = -1.81$. We failed to find parameters values that produced

[24] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks"; Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."
[25] Gutkin et al., "Turning On and Off with Excitation."
[26] Roxin et al., "On the Distribution of Firing Rates in Networks of Cortical Neurons."
[27] Koch and Fuster, "Unit Activity in Monkey Parietal Cortex Related to Haptic Perception and Temporary Memory."
[28] Griffith and Horn, "An Analysis of Spontaneous Impulse Activity of Units in the Striate Cortex of Unrestrained Cats."
[29] Zerlaut et al., "The Spectrum of Asynchronous Dynamics in Spiking Networks as a Model for the Diversity of Non-Rhythmic Waking States in the Neocortex."
[30] Renart et al., "The Asynchronous State in Cortical Circuits."
[31] Ecker et al., "Decorrelated Neuronal Firing in Cortical Microcircuits."
[32] Friedrich, Zhou, and Paninski, "Fast Online Deconvolution of Calcium Imaging Data."

realistic activity for 1/20 dataset, and it was excluded from this analysis. Even though the parameters were fixed for each network, the matrix a was build anew, and thus different, in each trial.

We initiated trials by setting $s_i(1) = 1$ for $\forall i \in N$ that spiked in the first 5 frames in imaged data. We thus had 30 trials for each direction, with differing starting conditions that were directly informed by data. Typically, 12.41±8.91% of neurons received an initial spike. We let activity propagate for 140 frames, at which point the simulation was terminated. For the performance per frame analysis (Fig. 7B), we used the spikes from a single time frame, i.e. a binary vector as inputs. We did not consider the first five frames to allow the dynamics to develop. To examine the effect of bin size (Fig. 7C), we binned frames into a rate vector to be used as input, again discarding the first 5 frames for all bin sizes. The decoder degradation analysis (Fig. 7D-F) was performed with rates binned over 100 time frames, from time frame 20 to 120. This analysis is described in more detail after the second phase of the decoder.

In the second phase of the decoder we used the binary vector or binned activity as input to a feed-forward neural network, in which N excitatory neurons in the input layer were connected in an all-to-all manner to 12 output units. The weights from the input to the output units were initialized randomly between (0,1), and then trained in a supervised learning paradigm with 90% of the inputs for each direction used for training, and 10% held out as a test-set. Training was performed using the MATLAB Network toolbox, with an objective function of minimizing the cross-entropy of the output and the correct targets, and computing that by adjusting the weights in the direction of the conjugate gradient. Once the network was trained, we took the decoded direction to be the identity of the output unit with the largest probability.

We tested the contribution of tuned and untuned neurons to the performance of the model in two ways: 1) we permuted the trained weights from tuned or untuned neurons to the output layer, and 2) we trained the feed-forward network with a subset of tuned or untuned neurons. In both cases, we took steps to manipulate the same number of weights or units. Most datasets contained more tuned neurons, so we performed 100 manipulations, each time picking U tuned neurons to manipulate, with U being the count of untuned neurons.

Control decoders

The cell identity decoder (Fig. 1C) was built by constructing a binary vector for each trial in each direction to be used as input to the second phase of our decoder, that is, the feed-forward pattern recognition network. For each trial we averaged the fluorescence traces of each neuron across the first 1.5 seconds of grating presentation. We then identified the most active neurons and set the binary input vector to 1 for neurons in this group, and 0 otherwise. We performed this procedure with the most active neurons defined at the top 5% to 30%, in steps of 5. Partitioning the inputs to training and test sets and decoding performance were implemented as described for the two-phase decoder.

The maximum-likelihood (ML) decoder was constructed as described in Avitan et al., 2016; Ponce-Alvarez et al., 2018. 90% of trials, picked at random were used as a training set and the remaining 10% of trials were used for testing. We trimmed the trials for each direction to the initial 2 seconds, and estimated conditional probability distributions for the fluorescence of each neuron given the grating direction using MATLAB's ksdensity function. To decode, we averaged the activity of each neuron over the 2 seconds in each test trial. We then collected the probabilities

to obtain this activity from the estimated distributions for the 12 directions, and multiplied these probabilities for the neurons, assuming they are independent. The decoded direction was taken as:

$$direction = \underset{dir \in [1,12]}{argmax} \left(\prod_{i=1}^{N} p(f_i|dir) \right)$$

where N is the number of neurons, and f is the time-average df/f.

The Poisson control decoder was constructed by inferring spikes from fluorescent activity (OASIS[33]), and then permuting the spikes for each neuron, thus preserving firing rates of individual units. For the frame analysis (Fig. 7B), we used the spiking in a single frame as a binary vector input to the decoder as described for phase II. For the bin size analysis (Fig. 7C), we binned the spikes and plugged in the rates as the input to the decoder. Partitioning the inputs to training and test sets and decoding performance were implemented as described for the two-phase decoder.

Tuning properties

We sorted neurons into tuned and untuned according to the following process: for each neuron we trimmed the fluorescence traces to the first 2 seconds of each stimulus epoch (out of 5 seconds) and the last 2 seconds of gray epochs (out of 3 seconds). We than averaged each trace and tested each stimulus condition against the preceding gray epoch in a paired-sample t-test. p values were Bonferroni corrected for 12 directions ($p<0.05/12$). Across datasets, 17.85+/-15.64% (M±SD) of neurons did not pass this test. We manually examined the cell bodies of those neurons in their respective fields of view, and their fluorescence traces, and subsequently 4 neurons across all 20 datasets were deemed as artifacts and excluded. We included all other neurons as untuned neurons in all our analyses. Cells which passed this procedure were then tested by repeated-measures

[33] Friedrich, Zhou, and Paninski.

analysis of variance for direction or orientation tuning significance ($p<0.01$) to further split this population into tuned and untuned neurons. Responses of significantly tuned cells were then iteratively fit with a circular Gaussian[34], with $R^2 = 0.72 \pm 0.19$ (M±SD). We took $\theta \in [0,330]$ in steps of 30 closest to the peak of the Gaussian to be the categorical tuning of the cell.

RESULTS

Visual cortical responses to a range of visual stimuli overlap in neuron identities

To evaluate the circuit encoding of visual stimuli we imaged populations of 73-347 layer 2/3 excitatory pyramidal neurons expressing GCaMP6s in murine primary visual cortex (see Methods; data previously described in [35]). Mice were awake and free to ambulate while they viewed drifting gratings in twelve directions (Stimulus) interleaved with a gray mean luminance screen (Gray). To evaluate the viability of a coding scheme based on the identities of neurons, we examined the overlap of the most active population between directions. The most active neurons were defined as cells with time-averaged fluorescence values in the top *n*th percentile. We varied *n* between 5 and 30 percent. We found that overlap depends on stimulus similarity (Fig. 1A and 1B), with the largest counts of shared neurons found between directions 180 degrees apart (which have the same orientation; 70.51±8.44%) and neighboring directions (63.96±10.01%). In contrast, 53.24±12.28% of the most active neurons were found to be shared between orthogonal directions. The extent of overlap would suggest that the changing sets of neurons active from trial to trial across multiple stimuli would represent an inefficient coding scheme. Indeed, a feed-forward neural net decoder (methods) that used the identities of the neurons active in response to each grating direction as

[34] Mazurek, Kager, and Van Hooser, "Robust Quantification of Orientation Selectivity and Direction Selectivity."
[35] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."

inputs failed to reach realistic levels of performance. This continued to be the case when we varied the threshold of what constitutes the active population (Fig. 1C). We confirmed this result with a more stringent maximum-likelihood (ML) decoder (methods, and [36,37]), decoding from the mean activity of all neurons rather than binarizing and setting the most active neurons to 1 and all the other neurons to 0 as we did with the decoder in Fig. 1C. The ML decoder had a mean performance of 28.12±11.61% across datasets (Fig. 1D), indicating, together with the feed-forward decoder that neither the identities of the most active neurons nor their level of activity was sufficient to decode from a population.

[36] Ponce-Alvarez et al., "Whole-Brain Neuronal Activity Displays Crackling Noise Dynamics."
[37] Avitan et al., "Limitations of Neural Map Topography for Decoding Spatial Information."

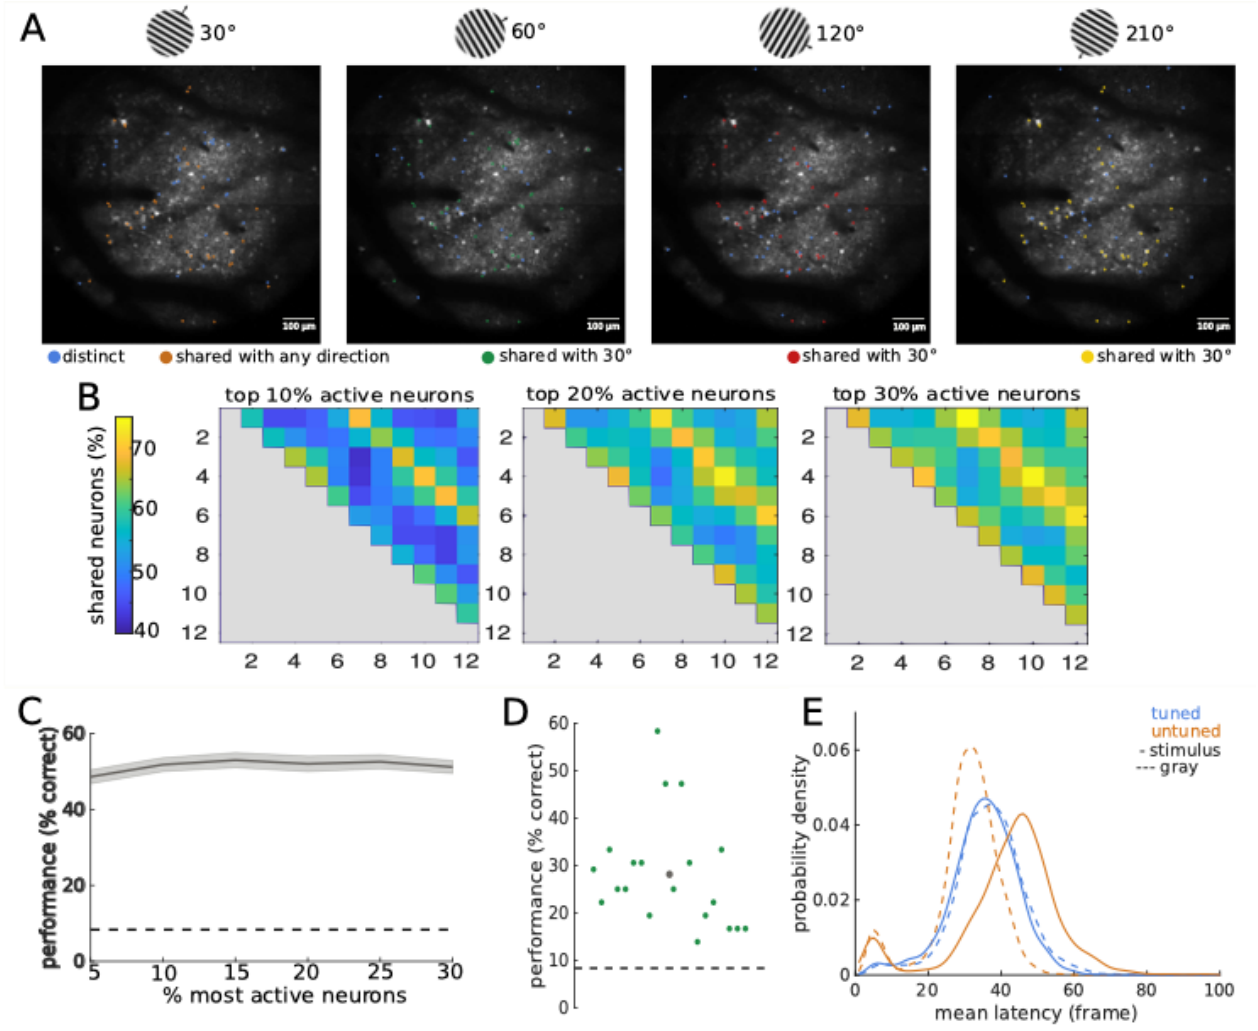

Figure 2.1 – Overlap between the most active cells in different stimuli

A. An example field of view (FOV) from dataset #3, with 347 neurons. Scale bars stand for 100µm. FOV visualization was done by stitching together images we acquired for cell detection in a 4X4 grid, as described in the methods. The top 20% of the most active neurons on average across trials are marked for drifting gratings at 30,60,120 and 210 degrees. Blue neurons are unique to the direction represented, whereas orange, green, red and yellow neurons are shared between stimuli, i.e. they are in the most active group in at least two directions. Note that 60 degrees is a neighboring direction to 30 degrees, 120 degrees is the orthogonal direction, and 210 degrees has the same orientation as 30 degrees.

B. Heatmaps quantifying the overlap in identities of the most active cells between each pair of directions of drifting grating, where the most active population was defined at the top 10% (left), 20% (middle) and 30% (right). Directions are on both axes of the heatmaps, with 1 being 30 degrees.

C. Performance of a decoder trained with the identities of the most active cells. The input to the decoder was a binary vector where the most active cells in each trial received 1 and all other cells received 0. The most active cell pool was

Figure 2.1 – continued

defined as the top 5%-30% cells with the largest fluorescence values averaged over the first 1.5 seconds of stimulus epochs. Line represents mean across datasets (N=20) and shading stands for the standard error of the mean.

D. Performance of a maximum-likelihood decoder. In this decoding approach we built conditional probability distributions for the activity of each neuron in response to each direction from a training set. Decoding was performed by taking the direction that produced the maximum-likelihood of a mean response in test set trials. Green dots represents datasets (N=20) and the larger gray dot stands for the mean across datatsets. Dashed line marks chance level.

E. Tuned neurons (blue) did not differ in their latency to activate in stimulus (solid, 33.95±34.66 frames) and gray (dashed, 33.03±27.05 frames) trials, whereas untuned neurons (orange) activated significantly earlier in gray epochs (27.33±24.92 frames) compared to drifting gratings (38.69±38.64 frames, $F=9245.39$, $p<0.001$) .

Functional networks summarize circuit pairwise correlations

We next explored a second coding scheme based on the co-activity between neurons. We sorted the population into significantly tuned (59.79±19.95%) and untuned (40.21±19.95%) neurons (Fig. 2A-C, see Methods). When comparing the two sub-populations we found that untuned neurons were consistently activated with a shorter latency during presentation of gray epochs (27.33±24.92 frames) but with a longer latency during presentation of drifting gratings (38.69±38.64 frames) as compared to tuned neurons (gray: 33.03±27.05, stimulus: 33.95±34.66 frame; $F=9245.39$, $p<0.001$, Fig. 1E). This longer latency of response to drifting gratings suggested that untuned neurons play a distinct functional role from tuned neurons in visual processing. To evaluate the potential roles of these two functional classes while also considering their interrelationships, we constructed FNs for each direction of drifting grating. We also separately considered FNs constructed from all stimulus epochs together (Stimulus) and FNs constructed from gray epochs (Gray; Fig. 2E). Briefly, we parsed the relevant epochs from fluorescence traces, and calculated the partial correlation coefficient between each pair of neurons, factoring out the average responses of the neurons in the pair (analogous to signal correlation) as

well as the overall level of population activity (corresponding to global state changes such as those that accompany locomotion, Fig. 2D). Thus, the FN summarized the partial pairwise correlation that is independent of stimulus and internal state and is analogous to noise correlations. We assigned directionality to connections based on the lag that resulted in a peak in the neuronal pairwise cross-correlation (Methods), with lag 0 taken to indicate a bi-directional connection.

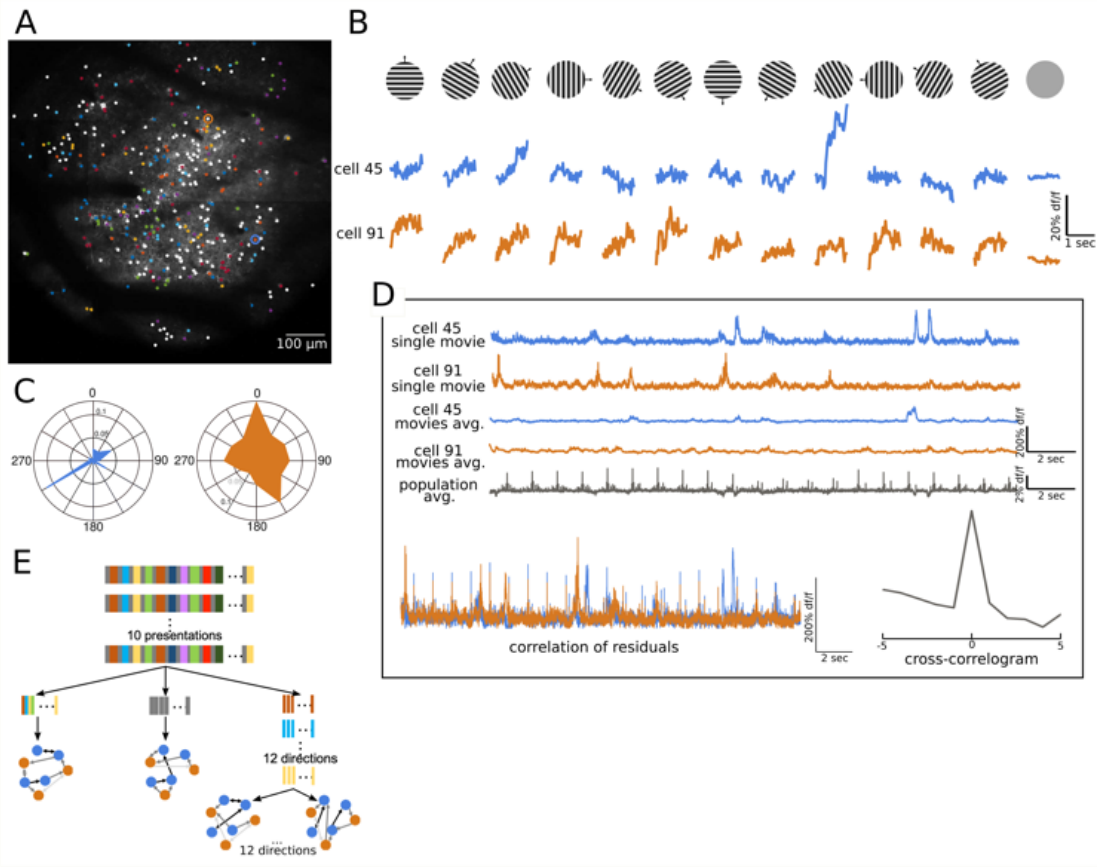

Figure 2.2 – Constructing a network comprising two functional classes

A. An example field of view (FOV) from dataset #3, with 347 neurons. Scale bar denotes 100µm. FOV visualization was done by stitching together images we acquired for cell detection in a 4X4 grid, as described in the methods. Neurons tuned to different directions are marked in different colors, and untuned neurons are marked in white. A tuned and untuned neuron are circled in blue and orange, respectively, for further illustration.

B. Trial-averaged fluorescence traces for 12 directions of drifting gratings and a luminance-matched gray screen for a tuned neuron (top, blue) and untuned neuron (bottom, orange). Those neurons are marked in panel A.

C. The resulting tuning curves for the neurons shown in A and B.

Figure 2.2 – continued

D. An illustration of edge weight and direction computation for the pair of neurons shown in A,B and C. Top: the activity of the tuned neuron (blue) and untuned neuron (orange) in a single movie. Middle: the averaged activity of those two neurons in all other movies (N=9 for most datasets). Bottom, gray: average activity of all other neurons in the population in the same single movie plotted on the top. For each neuron, we regress out the average activity of the neuron in other movies and the population average in the same movie. We then correlate the residual activity of the neurons in the pair, and this value is used as the edge weight. To assign directionality, we examine the peak in the cross-correlogram (bottom right). In this example, the peak is at zero and results in a bi-directional edge.

E. An illustration of the different functional networks we construct from fluorescent activity. Stimulus FN (left) is inferred from all visually evoked activity regardless of stimulus identity, gray FN (middle) is inferred from all gray epochs, and stimulus-specific or direction FNs are inferred only from trials of specific direction of drifting gratings.

First order topological features reflect single neuron response properties

Examining Stimulus FNs, we observed that edge weights between tuned neurons were quantitatively related to the similarity between preferred directions (Fig. 3A) in agreement with previous studies[38,39]. This remained the case when FNs were constructed from spikes inferred from the calcium fluorescence traces (see methods and Fig. S1). The strongest edge weights were present between neurons tuned to the same direction while neurons tuned to orthogonal directions were connected by weaker edge weights (see also [40]). Notably, gray FNs exhibited less structure in the arrangement of weights related to tuning properties likely reflecting the fact that similarly tuned neurons are more likely to be synaptically connected[41].

We next evaluated whether tuned and untuned neurons possessed unique topological signatures in and across FNs. We found that both tuned and untuned neurons had similar in-degrees in both stimulus FNs and gray FNs. However, tuned and untuned neurons differed in their out-

[38] Nauhaus et al., "Stimulus Contrast Modulates Functional Connectivity in Visual Cortex."
[39] Cossell et al., "Functional Organization of Excitatory Synaptic Strength in Primary Visual Cortex."
[40] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[41] Ko et al., "The Emergence of Functional Microcircuits in Visual Cortex."

degrees, with untuned neurons displaying significantly larger out-degrees in gray FNs as compared to stimulus FNs (Fig. S2). To clarify the source of this difference we examined the composition of incoming and outgoing edges in both functional classes of neurons. Specifically, we looked at two neuron-centric metrics: 1) realized edges, defined as edge count out of the potential pool (e.g. the count of incoming edges from untuned neurons over the number of untuned neurons) and 2) relative degree, which is the proportion of edges from a certain class out of the total edges of the neuron (e.g. the count of outgoing edges to tuned neurons over the number of outgoing edges of the neuron in question. For more details see methods).

We found that the average tuned neuron realizes more of the available potential outgoing edges to untuned neurons in Stimulus (0.32±0.14) compared to Gray (0.27±0.15, Fig. 3B) FNs, and untuned neurons constitute a larger portion of a tuned neurons outgoing pool in Stimulus (0.35±0.19) than in Gray (0.29±0.20, Fig. 3E) FNs. Conversely, untuned neurons showed increased realized out edges in Gray FNs regardless of whether the target was tuned: to untuned neurons (0.34±0.17 vs. 0.32±0.15 in Stimulus); to tuned neurons (0.33±0.17 vs. 0.26±0.17 in Stimulus, Fig. 3B). This latter trend also manifested as bigger relative out-degree of untuned to tuned neurons in Gray (0.47±0.20) as compared to stimulus (0.40±0.21, Fig. 3E) FNs. Edge weights reflected the same trends (Fig. 3C). Hence, neurons of the two functional classes show distinct correlation structure depending on the condition considered.

We next evaluated whether tuned neurons have a topological signature beyond partial pairwise correlation values. Tuned neurons displayed elevated realized in degrees in Stimulus FN dependent on the tuning of other tuned neurons. Elevated realized in degree were greatest from neurons preferring the same direction (0.50±0.25), followed by incoming edges from neurons preferring the direction 180 degrees away e.g. 30° and 210° (0.44±0.29), and then neurons tuned

to neighboring directions (30 degrees apart, e.g. 30° and 60°. 0.39±0.22). Tuned neurons showed low realization rate of incoming edges from neurons selective to orthogonal directions (0.23±0.23, Fig. 3D), mirroring the dependency of edge weight on the difference in tuning (Fig. 3A). In Gray FNs, however, realized in degree for tuned neurons depended less on the difference in tuning preference (same: 0.42±0.25, neighboring: 0.37±0.23, same orientation: 0.36±0.28), and notably higher realization was found for incoming edges from neurons preferring orthogonal directions (0.30±0.24, Fig. 3D). The composition of incoming connections for tuned neurons also depended on tuning similarity with large relative degrees for neurons tuned to the same (0.18±0.15) and neighboring (0.17±0.12) directions. Surprisingly, only a small portion of the average tuned neuron in-degree in Stimulus FNs was due to edges arising from neurons selective to directions 180 degrees away, i.e. sharing the same orientation (0.07±0.08, Fig. 3F). Again, in Gray FNs this effect was attenuated (same: 0.15±0.16, neighboring: 0.15±0.12, same orientation: 0.06±0.08), and relative in-degrees from neurons tuned to orthogonal direction was larger (0.07±0.08 vs. 0.05±0.07 in Stimulus, Fig. 3F). Realized and relative out-degrees were similar to in-degrees and are depicted in Fig. S3.

These results demonstrate that functional topology is non-random and specific to epochs of visual stimulation. Furthermore, neurons of different functional roles, tuned and untuned, differ in their correlations profile, in a manner that may enhance coding. Specifically, untuned neurons may contribute to visual stimulus encoding by modulating their interactions with tuned neurons, since these interactions are stimulus-specific. It follows that tuning, or a lack thereof, is therefore a manifestation of network interactions as much as it is a single neuron property.

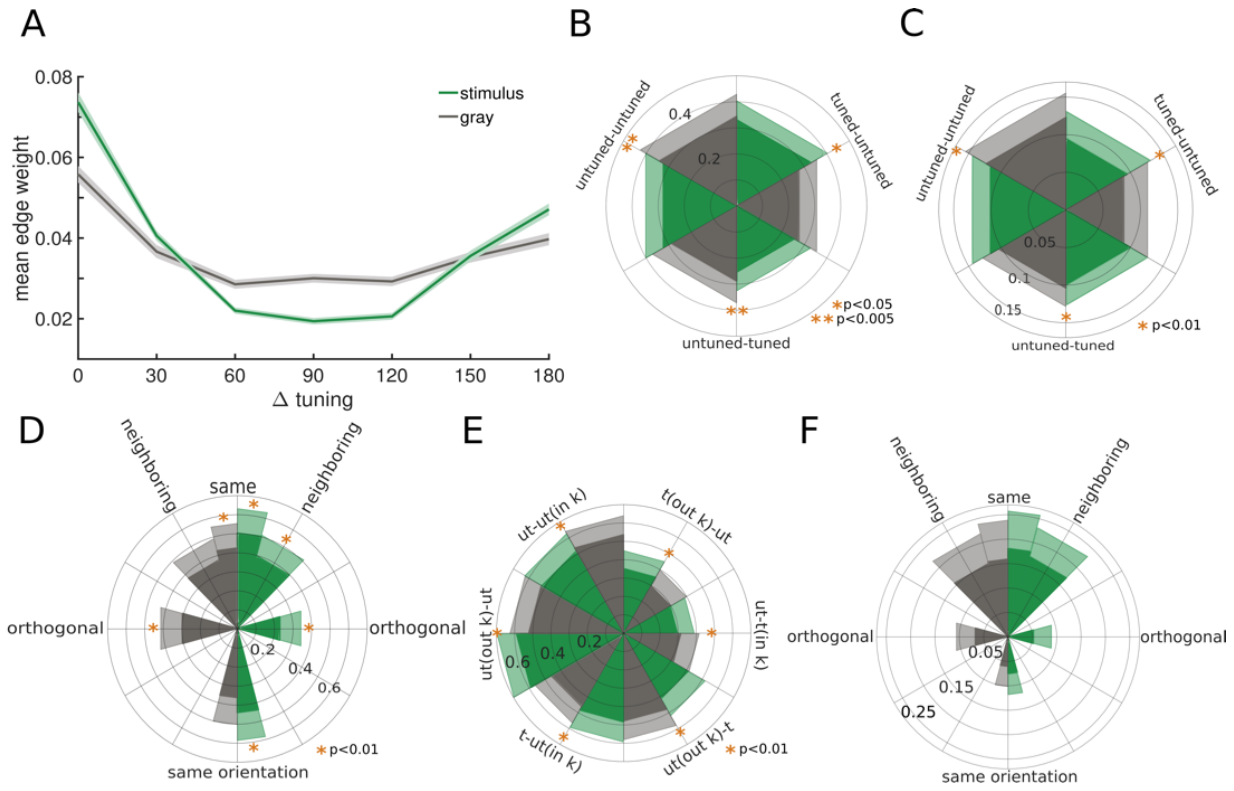

Figure 2.3 – Functional networks inferred from fluorescence contain stimulus features

A. Mean edge weight between pairs of tuned neurons is a function of similarity in their preferred direction in functional networks inferred from stimulus epochs, but not in gray epochs. Lines represent means across datasets (N=20), and shading represents the standard error of the mean. Note that pairs of neurons 180 degrees apart are composed of neurons that prefer the same orientation.

B. Untuned neurons have larger realized out-degrees in gray FNs (gray), both to tuned and untuned targets. The opposite is true for tuned neurons, with bigger realized out-degrees to untuned neurons in Stimulus (green) FNs. Slices in the polar plot present averages across neurons, shading stands for standard deviation. Asterisks denote significance in Kolmogorov-Smirnov test.

C. Same as in B for edge weights between pairs of neurons.

D. Realized in-degrees for tuned neurons depend on the neuronal source and the input (Stimulus vs. Gray). Slices show the means across neurons, and shading represents the standard deviation. Orange asterisks mark the mean as significantly different from all other conditions within the same FN (same color) according to Tukey-Kramer test. Two-way analysis of variance (FN type and incoming edge type) was also significant ($F=89.98$, $p<0.001$). In gray FNs, realized in-degrees from neurons tuned to neighboring directions and neurons tuned to the same orientation did not significantly differ from each other, but were different from the other two incoming edge types.

E. Relative degrees in Stimulus (green) and Gray (gray) FNs for tuned and untuned neurons. Slices and shading represent the means and standard deviations across neurons, respectively. t and ut denote tuned and untuned neurons, respectively. Asterisks stand for significant difference between Stimulus and Gray FNs according to Kolmogrov-Smirnov test, Bonferroni corrected.

Figure 2.3 - continued

F. Same as in D. for relative in-degrees. F=24.19, p<0.01. In both Stimulus and Gray FNs relative in-degree was similar from neurons tuned to the same and neighboring directions, and those two categories were significantly different (p<0.01 Tukey-Kramer corrected) than the relative in-degree from orthogonal and same orientation selective neurons. In turn, orthogonal and same orientation inputs also did not differ in their relative portion.
See also Fig. S1-S3.

Functional Networks are stimulus-specific

To further investigate the stimulus specificity of FNs we compared the functional networks that had been generated from each separate direction of drifting grating. Comparing the edge weights between pairs of neurons tuned to the same direction, we found the largest edge weights in FNs constructed from the direction the pair was tuned to (0.067±0.11). Edge weights in FNs constructed from neighboring directions were found to be increased (0.058±0.10) compared with FNs inferred from orthogonal directions (0.051±0.09, Fig. 4A). Noise correlations between pairs of neurons which are similarly tuned are thus stimulus dependent.

We next asked whether the topology of FNs as a whole is stimulus dependent as well, and sought to quantify this separately for tuned and untuned neurons. Graph alignment allows for a principled comparison of FNs identifying common edges between graphs[42] (Methods). This metric preserves node identities, ranges between 0 and 1 and is normalized to control graphs to evaluate whether alignment is larger than expected by chance. We measured alignment between each pair of FN constructed from trials of different drifting gratings directions. Alignment scores were calculated for FNs containing all neurons, and in sub-networks with only the tuned or untuned neurons, where edges between untuned or tuned neurons were set to zero, respectively. To explore the relevance of FN topology to stimulus coding, we were especially interested in alignment

[42] Gemmetto et al., "Multiplexity and Multireciprocity in Directed Multiplexes."

between FN for neighboring, opposite (180 degrees apart) and orthogonal directions. We found that edges were highly preserved across networks inferred from neighboring directions (0.225±0.053 alignment score), were less similar (0.200±0.053) for orthogonal directions and finally were highly similar for opposite directions (0.237±0.0056), which have the same orientation ($F=22.87$, $p<0.01$, Fig. 4B). Alignment scores were driven by edges amongst tuned neurons, as the tuned sub-FNs showed the same stimulus specificity ($F=47.57$, $p<0.01$, Fig. 4C). In contrast, the functional connectivity structure of untuned neurons was highly preserved regardless of stimulus similarity ($F=3.16$, NS. Fig. 4D) suggesting unique roles of the two functional classes in stimulus coding.

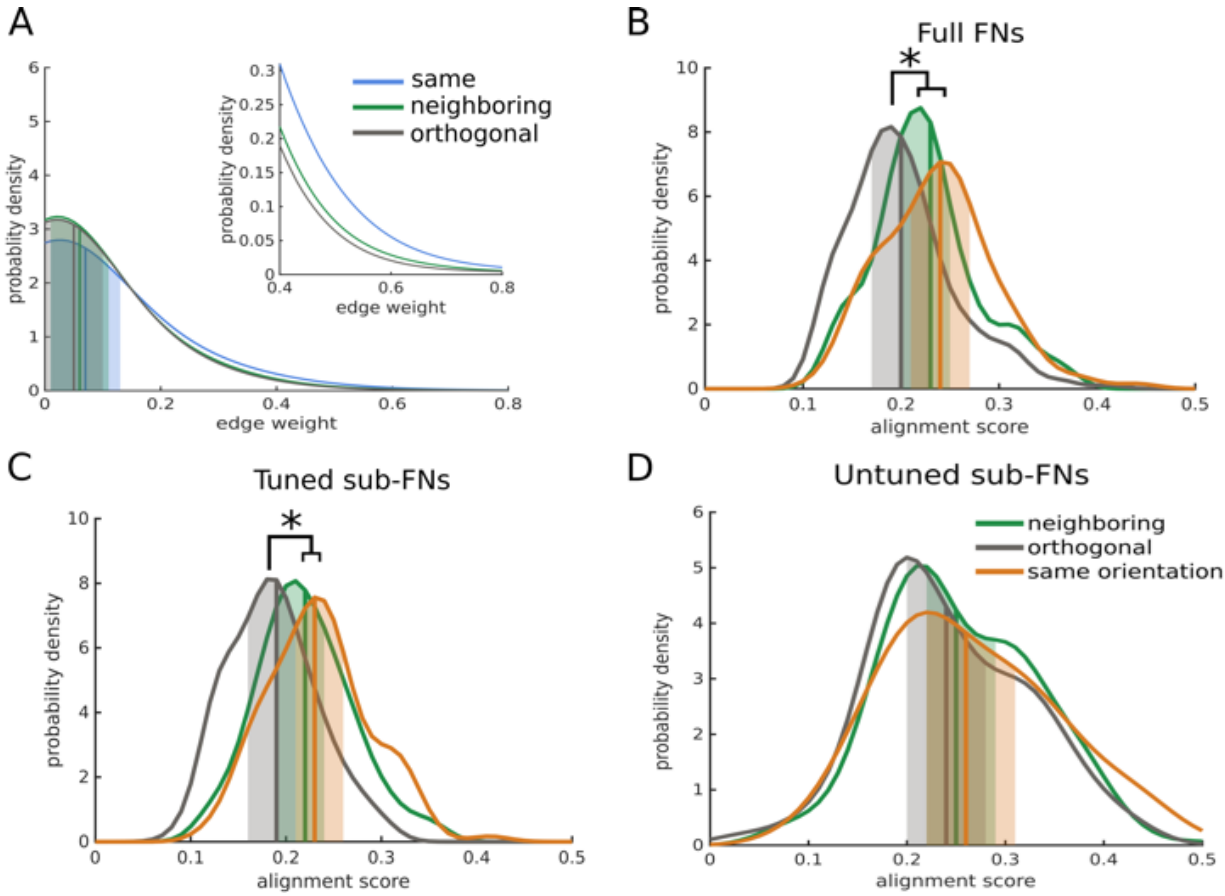

Figure 2.4 – Functional networks are stimulus-specific

A. Pairs of neurons preferring the same direction have larger edge weights in FNs constructed from the same (blue) direction they are tuned for, followed by FNs inferred from the neighboring directions (green). Pairs have the smallest

Figure 2.4 - continued

edges weights in FNs of orthogonal directions ($p<0.01$, all three groups different according to Kolmogrov-Smirnov test, Bonferroni corrected). Inset: zoom of large edge weights. Vertical lines and shading are means and standard deviations, respectively.

B-D. Probability density distributions of normalized alignment scores for pairs of FNs that are built from trials of neighboring directions (green), trials of directions 180 degrees apart (same orientation, orange) and trials of orthogonal directions (gray), in full FNs (D, $F=22.87$, $p<0.01$), sub-FNs including only tuned neurons (E, $F=47.57$, $p<0.01$) and sub-FNs with only untuned neurons (F, $F=3.16$, NS). Vertical lines and shading are means and standard deviations, respectively.

Untuned neurons form a rich club of large edge weights

Recent studies of neuronal networks constructed from whole human brain imaging data[43,44], cortical slices[45], cortical cultures[46] and fronto-parietal cortex in monkeys[47] have found a rich club topology, in which the nodes with the largest degrees are also densely connected amongst themselves. Here we examined whether FNs in L2/3 in V1 *in vivo* also exhibit a rich club organization (Methods). We found that all datasets displayed significant rich club topology spanning the majority of neurons in the sample, in both the stimulus and gray FNs (Fig. 5A-C). To probe the position of untuned neurons within the rich club, we iteratively thresholded the networks according to edge weight. Consequently, at each iterative stage we included smaller and smaller weights and networks became increasingly dense. We then sorted the neurons by their degrees in the resultant networks and examined the composition of the neurons with the *k*th percentile of top degrees (Fig. 5D). Untuned neurons were found to be more prevalent in the group of neurons with

[43] Heuvel et al., "High-Cost, High-Capacity Backbone for Global Brain Communication."
[44] Heuvel and Sporns, "Rich-Club Organization of the Human Connectome."
[45] Nigam et al., "Rich-Club Organization in Effective Connectivity among Cortical Neurons."
[46] Faber et al., "Computation Is Concentrated in Rich Clubs of Local Cortical Networks."
[47] Dann et al., "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates."

the strongest weights and largest degrees, as evident in the weight-thresholded networks. As more small weights were included in the networks, untuned neurons no longer possessed the largest degrees in the network (Fig. 5E-F). When pruning the networks in reverse order, keeping the smallest weights at each iteration[48], untuned neurons were less likely to be amongst the most connected neurons in both Stimulus and Gray FNs when only small edge weights were included (Fig. S4). This indicates that untuned neurons take part in a rich club of strong weights, putting them in a prime position for integration and computation[49,50]

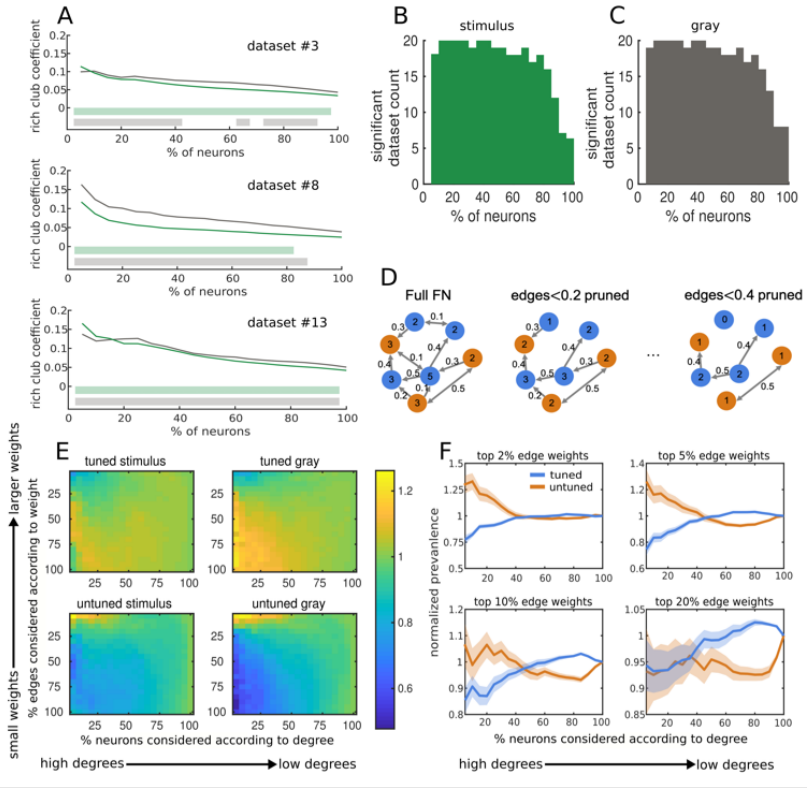

Figure 2.5 – Rich club structure in functional networks

A. Examples from three datasets of rich club coefficients in stimulus (green) and gray (gray) FNs across percentiles of neurons, sorted in descending order by their degrees. Lines represent rich club coefficients and shading represent significance with $p<0.01$, tested against a distribution of rich club coefficients from 1000 networks permuted to preserve degree-sequence distributions. **B.** Count of datasets (out of 20) which display significant rich club coefficient in stimulus networks. Significance was determined as described in A.

C. Same as B., for gray functional networks.

D. A cartoon of the rich club sparsification procedure: left, we start with all the edges in the network. Edge weight is specified next to the arrow and the degree of the neurons is written on each node. Tuned neurons are in blue and

[48] Pajevic and Plenz, "The Organization of Strong Links in Complex Networks."
[49] Faber et al., "Computation Is Concentrated in Rich Clubs of Local Cortical Networks."
[50] Dann et al., "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates."

Figure 2.5 - continued

untuned neurons are in orange. At each step we discard the smallest edges and re-count the degrees of the neurons. We then ask what is the prevalence of tuned and untuned neurons within the group of neurons with the largest degrees, and normalize this prevalence by their frequency in the population.

E. Results of the analysis described in D. Untuned neurons form a rich club of strong weights. Warmer colors represent prevalence that is over the expected prevalence (value greater than 1). As functional networks move from bottom to top on the heatmaps they become more sparse, with only the strongest edges remaining on the top. As FNs move from left to right more neurons are examined, with the fewest, largest degree neurons on the left side. See also Fig. S4.

F. Four cross-sections of E, showing that as networks are sparsified to hold only the strongest weights (top), untuned neurons (orange) are more prevalent amongst high-degree neurons than expected by their prevalence in the population. All FNs in this analysis are stimulus FNs, lines represent means across datasets (N=20) and shading stands for the standard error.

Untuned neurons are hubs

In order to better quantify the network contribution of untuned neurons we measured the centrality of tuned and untuned neurons within the FN topology. Centrality can be evaluated by a number of different metrics that focus on local or global network patterns. Of particular relevance in the present context is the family of PageRank (eigenvector-based) algorithms, which take into account the network embedding of any given node. PageRank has been shown to capture the importance of nodes in a variety of biological systems. For example, in protein networks, PageRank identifies proteins underlying traits and predictive of prognosis[51] and in ecological networks it identifies species that are crucial for biodiversity[52]. In neuroscience, NeuronRank, a measure inspired by the PageRank algorithm, was found to be correlated with firing rates of single neurons[53], and more importantly, of the population[54] in networks of integrate-and-fire neurons. Here we used a

[51] Wang and Marcotte, "It's the Machine That Matters."
[52] Domínguez-García and Muñoz, "Ranking Species in Mutualistic Networks."
[53] Fletcher and Wennekers, "From Structure to Activity."
[54] Gürel, De Raedt, and Rotter, "Ranking Neurons for Mining Structure-Activity Relations in Biological Neural Networks."

variation of PageRank[55], that allowed us to assign an individual authority value to nodes *a priori*. This proved useful for testing a model in which tuned neurons are hypothesized to be a more reliable source of information. However, for the majority of the analyses we kept the authority scores (z; see Methods) equal across all neurons unless stated otherwise. We converted the raw ranking scores into relative rankings, which quantifies the percentage of neurons that possess a smaller ranking than the neuron in question.

We found that untuned neurons had larger rankings in FNs inferred from stimulus epochs (0.542±0.292) as compared to tuned neurons (0.465±0.281, $t=8.599$, $p<0.001$; Fig. 6A). In contrast, tuned neurons displayed comparable rankings in gray FNs (0.518±0.287) to untuned neurons (0.468±0.288; Fig. S5). This effect was robust across model parameters such as the damping value, which reflects time spent in the system (q; See Methods. Fig. 6B). The effect suggests that activity converges mostly onto untuned neurons across multiple timescales during stimulus presentation. This result was also robust to permuting incoming edges of untuned neurons (Fig. 6C), consistent with the rich club structure of this sub-population. Interestingly, permuting incoming edges of tuned neurons resulted in a mild increase in the ranking of tuned neurons and a substantial decrease in the ranking of untuned neurons (Fig. S6). Taken together with the robustness of ranking to configurations of incoming edges in untuned neurons, these results suggested that network topology is organized to form a backbone of untuned hubs. Finally, to asses a model of V1 in which one functional sub-population is a more trusted source of information, we varied the ratio between the authority values (z; see Methods) of tuned and untuned neurons. The two sub-populations showed opposite trends: when the ratio favored tuned neurons, they showed increased rankings in gray functional networks, whereas untuned neurons showed smaller

[55] Radicchi et al., "Diffusion of Scientific Credits and the Ranking of Scientists."

rankings. In contrast when untuned neurons were assigned larger authority values, they had even larger rankings in networks built from stimulus epochs, while tuned neurons had decreased ranking scores in random walks on these networks (Fig. 6D). Importantly, even when tuned neurons were set to be 4 times more influential than untuned neurons, untuned neurons remained highly ranked in networks built from stimulus epochs.

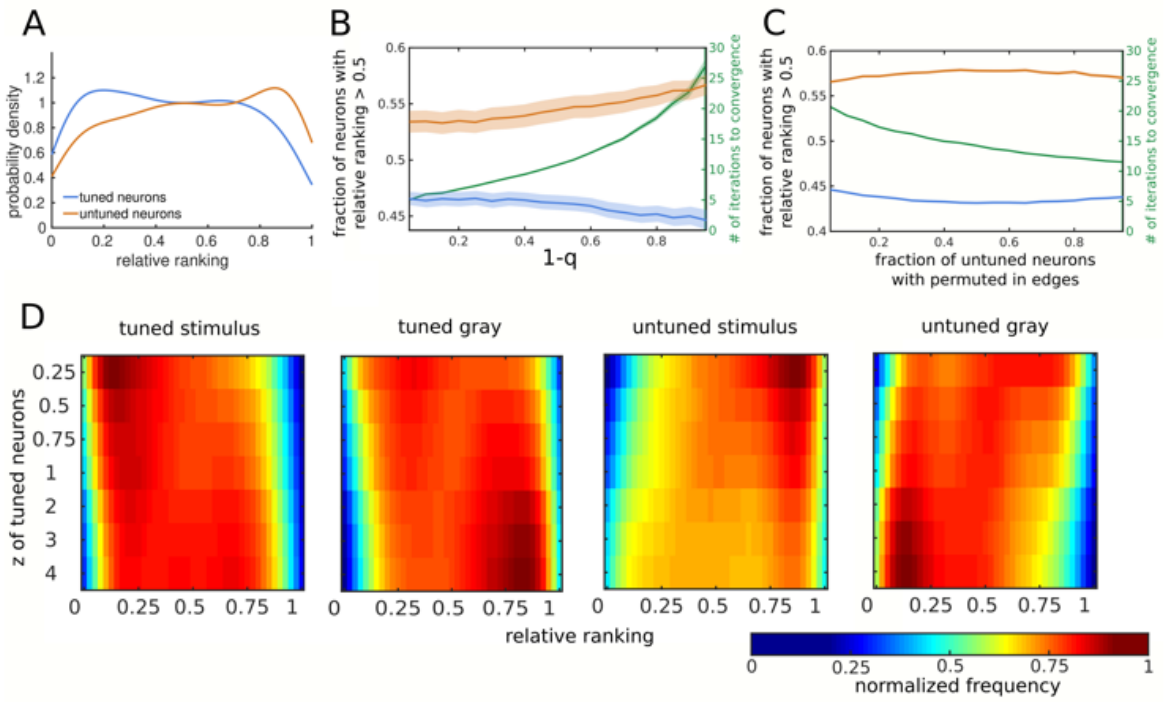

Figure 2.6 – Random walks on stimulus networks tend to converge onto untuned neurons

A. Probability density functions of the relative ranking for tuned (blue) and untuned (orange) neurons in networks constructed stimulus epochs. Untuned neurons show increased relative ranking, which is a metric for pooling activity, $t=8.599$, $p<0.001$. No such effect is present in gray FNs (Fig. S5).

B. Larger ranking among untuned neurons in stimulus epochs is robust to the damping factor, and can thus be thought of as occurring at multiple timescales. Lines represent the means of tuned (blue) and untuned cells (orange), and shading stands for the standard error. The right-side axis (green) shows the number of steps it took for the random walk to converge.

C. Untuned neurons in stimulus networks are ranked highly regardless of their incoming edges. We gradually increased the portion of untuned neurons whose incoming edges we randomly permuted, in steps of 5%. For each step we ran 30 trials, in which we chose neurons at random and permuted their incoming edges at random. Lines represent means over datasets, with shading for the standard error.

Figure 2.6 - continued

D. Differences in the ranking of tuned and untuned neurons persist under a model in which one of the sub-population is more influential. For each value of $x \in [0.25, 4]$ we set z(tuned)=x and z(untuned)=1. We then ran the algorithm and computed the relative ranking of each of the sub-populations. Untuned neurons tended to have larger rankings in stimulus FNs, especially when they had larger z values compared with tuned neurons, but also when tuned neurons were assigned authority values as 4 times larger.
See also Fig. S5-S6.

A decoder model of V1 functional networks relies on both tuned and untuned neurons

In order to directly test the idea that the specific topology of a FN comprised of tuned and untuned neurons contains information about the stimulus we constructed a two-stage model: the first stage was a simple generative model, in which we simulated spiking activity within each of the 12 direction-specific FNs. Specifically, we used the edge weights from the networks as synaptic weights in a sparse recurrent neural network (RNN) and initiated activity by activating the small subset of neurons that had the shortest latency response in the experimental data in each stimulus condition (12.41±8.91% of neurons). In addition to using the edges from the FNs, inferred from data, as the weights in the RNN we ensured that all other parameters in the RNN were biologically realistic (see Methods). To do so we performed a grid search and matched the activity produced by the RNN to the first-order statistics of spiking activity recorded from mouse V1. In the second phase we used the spikes produced by each FN as inputs in a decoding framework, which was designed as a feed-forward (FF) neural network with 12 output units, for the 12 directions. The connectivity between N input units and the output layer was all-to-all, and was trained on simulated spiking data produced using conjugate gradient (Methods, Fig. 7A). Since we employed FNs inferred from data and froze the connections within them, the topological properties as well as neuron identities were preserved from the *in vivo* experimental data in the model.

This decoding approach proved to be efficacious. Decoding from a single time-step in our model, which was a binary vector, resulted in performance that was 310.53±72.68% over chance level, and also significantly exceeded the performance of a Poisson decoder in which the neurons were considered independent (Fig. 7B). Binning time-steps and decoding from firing rates of units in each network saturated the performance in most datasets when bin sizes exceed 11.94±3.53 frames (Methods, Fig. 7C), and greatly exceeded the performance of the decoder based on the most active neurons (Fig. 1C).

We next examined the level at which the sub-populations of tuned and untuned neurons contributed to decoder performance. To assess their importance, we trained the weights to the output layer with all neurons, but then permuted the weights from tuned or untuned neurons and tested the decoder on held-out test set data. Performance degradation was equivalent regardless of the sub-population of neurons we permuted (Fig. 7D). In fact, degradation in performance was a linear function of the fraction of cells whose weights we permuted, with no difference between tuned or untuned neurons (Fig. 7E-F). Similar results were found when we trained the network with a subset of neurons, with tuned-only and untuned-only networks performing comparably (Fig. S7).

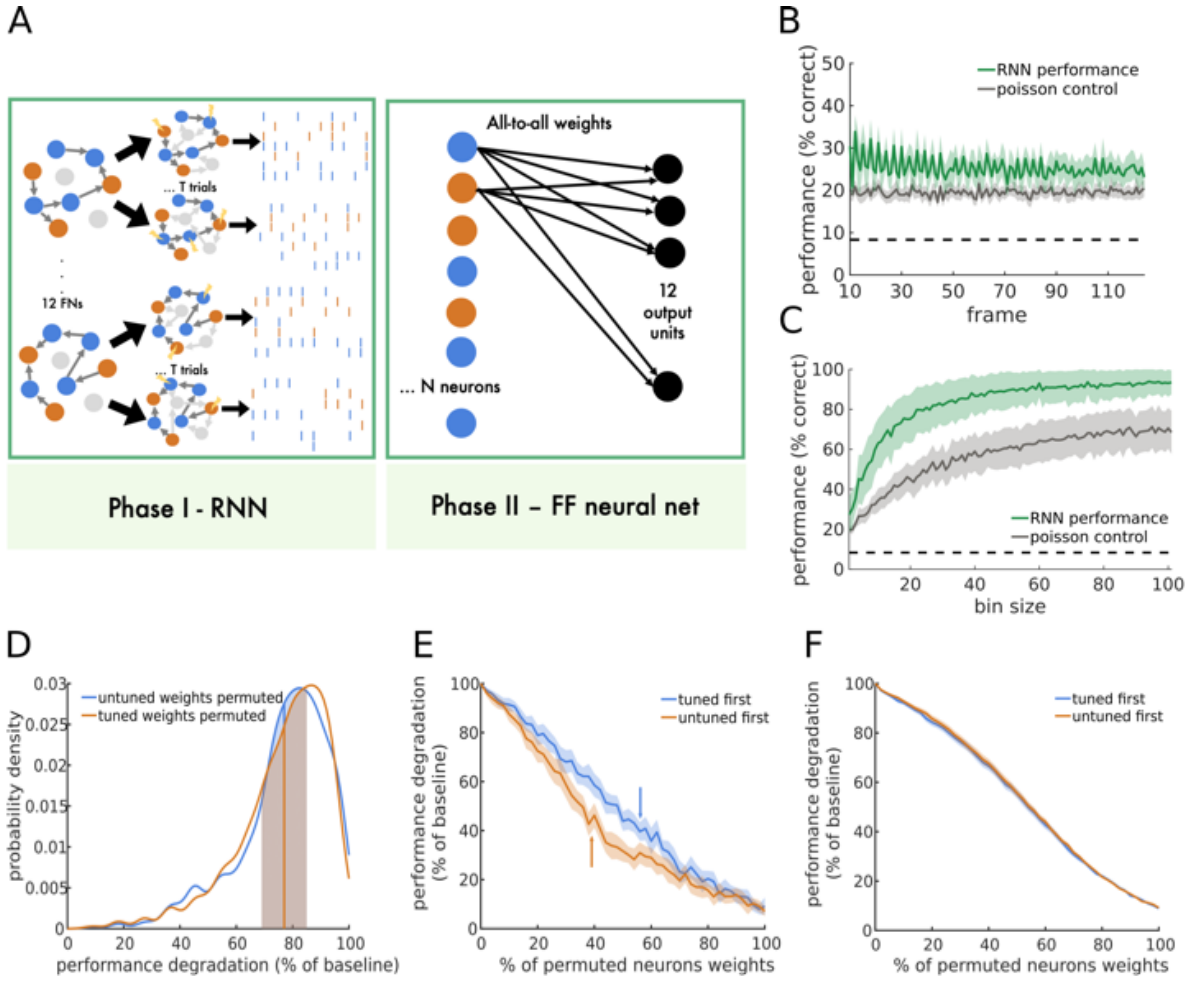

Figure 2.7 – Two-phase model of V1 decodes accurately from both tuned and untuned neurons

A. Illustration of the two-phase decoder model. In Phase I we instantiate each of the direction specific FNs as a recurrent neural network (Methods). Edges between tuned (blue) and untuned neurons (orange) as well as edges within the same functional class are inferred from data, whereas inhibitory neurons and inhibitory edges (light gray) are added in a pseudo-random fashion according to parameters to balance the activity. Inputs (yellow) consist of inserting a spike in neurons that spiked in real data in the first 5 imaging frames. Each one of the 12 direction FNs produced spiking activity that was used as inputs to Phase II. In the second phase we either used spikes from one frame (i.e. binary vector, B), or binned over frames to produce a rate vector (C). Each one of the neurons was connected in an all-to-all manner to 12 output units, and those connection weights were trained with conjugate gradient.

B. Performance of 19 datasets when decoding from a single frame across the population of N neurons (green), compared with a poisson decoder in which we permuted the spikes from the real data (gray). Practically, in this analysis we trained the decoder on binary vectors of size 1XN, taken from time step t. We then decoded the direction

Figure 2.7 - continued

of drifting gratings from a held-out set of binary vectors from the same t. Lines represent the mean across datasets, and shading is for the standard deviation. Note that chance level is 1/12 = 8.3% (dashed).

C. Performance of the decoder is initially a function of bin size, but quickly saturates (green). Performance of the poisson decoder is plotted in gray. We trained the decoder on summed activity from n timed steps, and decoded from a held-out set of summed activity from the same n, starting and ending at the same frames. Lines represent the mean across datasets, and shading is for the standard deviation. Note that chance level is 1/12 = 8.3% (dashed).

D. Decrease in performance is equivalent when the learned weights of tuned (blue) or untuned (orange) neurons to the output layer are permuted on a single neuron base. For each neuron, we randomly shuffled the twelve weights to the twelve output units. Probability density distributions are across 100 permutations in each dataset. At each permutation we also picked neurons at random from the more prevalent sub-population such that the count of permuted neurons in comparable to the smaller sub-population. This means that in 16/19 cases we did not permute all tuned neurons. Vertical lines and shading are for the means and standard deviations, respectively.

E. Example dataset showing performance degradation is linear with the percent of permuted neurons (and therefore permuted weights). For each percentage of neurons from 5% to 100%, we did 30 permutations: we first pseudo-randomly picked the neurons whose weights will be permuted by randomly sampling only tuned (blue) or untuned (orange) neurons, until there were no more neurons left in this sub-population, and then we randomly sampled from the other sub-population until the desired percentage was reached. Finally, we randomly shuffled the weights of picked neurons to the output layer. Lines represent means across 30 permutations for each percentage, and shading stands for the standard deviation. Arrows point at the percentage at which all tuned (blue) neurons are being permuted and we start permuting untuned neurons, and orange is the same for untuned neurons.

F. The same analysis as in E., pooled across datasets. Lines and shading are the means and standard errors, respectively.

See also Fig. S7.

DISCUSSION

Here we tested the hypothesis that functional networks, comprised of both untuned and tuned neurons, are stimulus-specific, and that functional network topology itself contains visual stimulus information that is accessible to and thus decodable by downstream neurons. Our findings indicate that the two sub-populations occupy specific topological positions in FNs, suggesting a collective role for all neurons, regardless of tuning, in the network representation of visual stimuli. This finding emphasizes the need for sufficiently broad and unbiased population sampling when

studying cortical population coding. We note that although we sorted cells into two discrete categories in this study, in reality tuning is a continuous, graded property. Multiple quantification methods and parameter choices will engender varying proportions of tuned neurons, as well as tuning strength, rendering the dichotomy of tuned versus untuned somewhat arbitrary. However, the unique FN topological positions occupied by untuned neurons suggest that this functional designation is as much a manifestation of network interactions[56,57,58] as it is a manifestation of single neuron properties regardless of the specific parameter choices. In that regard, layer 2/3 may be different from layer 4, in which tuning properties are thought to be inherited from the dorsal lateral geniculate nucleus[59]. Finally, it is likely that distinct classes of stimuli such as gratings, dots, or natural movies, will result in distinct functional networks, consistent with the results presented here, and we hypothesize that the assignment of individual neurons to tuned and untuned categories will also change in a stimulus-dependent manner again indicative of the fact that tuning is at least in part a consequence of network interactions. We suggest that our network-based approach can generalize across stimulus class.

This study highlights the crucial role of functional network topology in stimulus coding. Previous studies have included statistical dependencies between neurons and found gains in decoding accuracy[60,61]. We build on these results by generating and instantiating a complete functional topology from data, including tuned and untuned neurons, which preserves higher-order

[56] Cossell et al., "Functional Organization of Excitatory Synaptic Strength in Primary Visual Cortex."
[57] Arakaki, Barello, and Ahmadian, "Capturing the Diversity of Biological Tuning Curves Using Generative Adversarial Networks."
[58] Amsalem et al., "From Neuron Biophysics to Orientation Selectivity in Electrically Coupled Networks of Neocortical L2/3 Large Basket Cells."
[59] Ringach, Shapley, and Hawken, "Orientation Selectivity in Macaque V1."
[60] Shi, Niu, and Wan, "Effect of the Small-World Structure on Encoding Performance in the Primary Visual Cortex."
[61] Graf et al., "Decoding the Activity of Neuronal Populations in Macaque Primary Visual Cortex."

structure and also renders the read-out layer of our decoder naïve to all but the realistic spiking activity the network produces. Consequently, we show that every sensory neuron matters in the context of the active network and in the circuit-level representation of visual stimuli.

We find that untuned cells are strongly connected amongst themselves. This places untuned neurons at the core of the circuit, as evidenced in their propensity to form a rich-club of strong weights. This demonstration of a rich-club property in V1 FNs *in vivo* implicates rich-clubs in processing on multiple spatial scales. Rich-club topologies have been linked to increased integration of information[62], synchronization[63], and fast decision making[64]. Member nodes, or neurons, of the rich-club were shown to perform disproportionate amounts of computation in cortical cultures[65]. Furthermore, a spin glass model has demonstrated that a rich-club organization supports a network's capability to converge to a larger set of attractor states[66] (see also [67]) and hints at the potential role of rich-club neurons in sensory discrimination. Our data agree with this hypothesis. The central position occupied by untuned neurons in visually evoked FNs is reinforced by their larger ranking in dynamic network models based on random walks, which suggests an important role as integrators, or poolers, of information. Taken together with their rich-club structure of strong weights, untuned neurons may be especially crucial when the stimulus is ambiguous or low contrast[68]. Untuned neurons may hence aid the visual coding of more complex, naturalistic stimuli.

---

[62] Heuvel and Sporns, "Rich-Club Organization of the Human Connectome."
[63] Watanabe, "Rich-Club Network Topology to Minimize Synchronization Cost Due to Phase Difference among Frequency-Synchronized Oscillators."
[64] Daniels and Romanczuk, "Quantifying the Impact of Network Structure on Speed and Accuracy in Collective Decision-Making."
[65] Faber et al., "Computation Is Concentrated in Rich Clubs of Local Cortical Networks."
[66] Senden et al., "Rich Club Organization Supports a Diverse Set of Functional Network Configurations."
[67] Ponce-Alvarez et al., "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons."
[68] Nauhaus et al., "Stimulus Contrast Modulates Functional Connectivity in Visual Cortex."

We propose that a model of V1 circuits based on functional networks supports important functional roles in stimulus coding for both tuned and untuned neuronal populations. Moreover, our work suggests that the functional designation of a neuron as being tuned or untuned is a consequence of network topological interactions. Our work represents an example of how the application of network models and graph theory can provide insights and test hypotheses for future investigation of how neuronal populations encode and compute sensory signals.

REFERENCES

Amsalem, Oren, Werner Van Geit, Eilif Muller, Henry Markram, and Idan Segev. "From Neuron Biophysics to Orientation Selectivity in Electrically Coupled Networks of Neocortical L2/3 Large Basket Cells." *Cerebral Cortex* 26, no. 8 (August 1, 2016): 3655–68. https://doi.org/10.1093/cercor/bhw166.

Arakaki, Takafumi, G. Barello, and Yashar Ahmadian. "Capturing the Diversity of Biological Tuning Curves Using Generative Adversarial Networks." *ArXiv:1707.04582 [Cs, q-Bio]*, July 14, 2017. http://arxiv.org/abs/1707.04582.

Avitan, Lilach, Zac Pujic, Nicholas J. Hughes, Ethan K. Scott, and Geoffrey J. Goodhill. "Limitations of Neural Map Topography for Decoding Spatial Information." *Journal of Neuroscience* 36, no. 19 (May 11, 2016): 5385–96. https://doi.org/10.1523/JNEUROSCI.0385-16.2016.

Bassett, Danielle S., and Olaf Sporns. "Network Neuroscience." *Nature Neuroscience* 20, no. 3 (March 2017): 353–64. https://doi.org/10.1038/nn.4502.

Chambers, Brendan, and Jason N. MacLean. "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks." Edited by Jochen Triesch. *PLOS Computational Biology* 12, no. 8 (August 19, 2016): e1005078. https://doi.org/10.1371/journal.pcbi.1005078.

Chen, Yuzhi, Wilson S Geisler, and Eyal Seidemann. "Optimal Decoding of Correlated Neural Population Responses in the Primate Visual Cortex." *Nature Neuroscience* 9, no. 11 (November 2006): 1412–20. https://doi.org/10.1038/nn1792.

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 (February 2015): 399–403. https://doi.org/10.1038/nature14182.

Dana, Hod, Tsai-Wen Chen, Amy Hu, Brenda C. Shields, Caiying Guo, Loren L. Looger, Douglas S. Kim, and Karel Svoboda. "Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging in Vivo." *PloS One* 9, no. 9 (2014): e108697. https://doi.org/10.1371/journal.pone.0108697.

Daniels, Bryan C., and Pawel Romanczuk. "Quantifying the Impact of Network Structure on Speed and Accuracy in Collective Decision-Making." *ArXiv:1903.09710 [Cs, q-Bio]*, March 22, 2019. http://arxiv.org/abs/1903.09710.

Dann, Benjamin, Jonathan A. Michaels, Stefan Schaffelhofer, and Hansjörg Scherberger. "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates." *ELife* 5 (August 15, 2016). https://doi.org/10.7554/eLife.15719.

Dechery, Joseph B., and Jason N. MacLean. “Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons.” Edited by Jeff Beck. *PLOS Computational Biology* 14, no. 5 (May 4, 2018): e1006153. https://doi.org/10.1371/journal.pcbi.1006153.

Domínguez-García, Virginia, and Miguel A. Muñoz. “Ranking Species in Mutualistic Networks.” *Scientific Reports* 5 (February 2, 2015): 8182. https://doi.org/10.1038/srep08182.

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. https://doi.org/10.1126/science.1179867.

Faber, Samantha P., Nicholas M. Timme, John M. Beggs, and Ehren L. Newman. “Computation Is Concentrated in Rich Clubs of Local Cortical Networks.” *Network Neuroscience* 3, no. 2 (September 14, 2018): 384–404. https://doi.org/10.1162/netn_a_00069.

Fletcher, Jack McKay, and Thomas Wennekers. “From Structure to Activity: Using Centrality Measures to Predict Neuronal Activity.” *International Journal of Neural Systems* 28, no. 02 (November 16, 2016): 1750013. https://doi.org/10.1142/S0129065717500137.

Friedrich, Johannes, Pengcheng Zhou, and Liam Paninski. “Fast Online Deconvolution of Calcium Imaging Data.” *PLoS Computational Biology* 13, no. 3 (2017): e1005423.

Gemmetto, Valerio, Tiziano Squartini, Francesco Picciolo, Franco Ruzzenenti, and Diego Garlaschelli. “Multiplexity and Multireciprocity in Directed Multiplexes.” *Physical Review E* 94, no. 4 (October 27, 2016). https://doi.org/10.1103/PhysRevE.94.042316.

Graf, Arnulf B A, Adam Kohn, Mehrdad Jazayeri, and J Anthony Movshon. “Decoding the Activity of Neuronal Populations in Macaque Primary Visual Cortex.” *Nature Neuroscience* 14, no. 2 (February 2011): 239–45. https://doi.org/10.1038/nn.2733.

Griffith, J. S., and G. Horn. “An Analysis of Spontaneous Impulse Activity of Units in the Striate Cortex of Unrestrained Cats.” *The Journal of Physiology* 186, no. 3 (October 1966): 516–34.

Gürel, Tayfun, Luc De Raedt, and Stefan Rotter. “Ranking Neurons for Mining Structure-Activity Relations in Biological Neural Networks: NeuronRank.” *Neurocomputing*, Computational Neuroscience: Trends in Research 2007, 70, no. 10 (June 1, 2007): 1897–1901. https://doi.org/10.1016/j.neucom.2006.10.064.

Gutkin, Boris S., Carlo R. Laing, Carol. L. Colby, Carson C. Chow, and G. Bard Ermentrout. “Turning On and Off with Excitation: The Role of Spike-Timing Asynchrony and Synchrony in Sustained Neural Activity.” *Journal of Computational Neuroscience* 11, no. 2 (September 1, 2001): 121–34. https://doi.org/10.1023/A:1012837415096.

Heuvel, Martijn P. van den, René S. Kahn, Joaquín Goñi, and Olaf Sporns. "High-Cost, High-Capacity Backbone for Global Brain Communication." *Proceedings of the National Academy of Sciences* 109, no. 28 (July 10, 2012): 11372–77. https://doi.org/10.1073/pnas.1203593109.

Heuvel, Martijn P. van den, and Olaf Sporns. "Rich-Club Organization of the Human Connectome." *Journal of Neuroscience* 31, no. 44 (November 2, 2011): 15775–86. https://doi.org/10.1523/JNEUROSCI.3539-11.2011.

Hubel, D. H., and T. N. Wiesel. "Receptive Fields of Single Neurones in the Cat's Striate Cortex." *The Journal of Physiology* 148, no. 3 (October 1959): 574–91.

Josić, Krešimir, Eric Shea-Brown, Brent Doiron, and Jaime de la Rocha. "Stimulus-Dependent Correlations and Population Codes." *Neural Computation* 21, no. 10 (2009): 2774–2804.

Kalatsky, Valery A., and Michael P. Stryker. "New Paradigm for Optical Imaging: Temporally Encoded Maps of Intrinsic Signal." *Neuron* 38, no. 4 (May 22, 2003): 529–45. https://doi.org/10.1016/S0896-6273(03)00286-1.

Ko, Ho, Lee Cossell, Chiara Baragli, Jan Antolik, Claudia Clopath, Sonja B. Hofer, and Thomas D. Mrsic-Flogel. "The Emergence of Functional Microcircuits in Visual Cortex." *Nature* 496, no. 7443 (April 3, 2013): 96–100. https://doi.org/10.1038/nature12015.

Koch, K. W., and J. M. Fuster. "Unit Activity in Monkey Parietal Cortex Related to Haptic Perception and Temporary Memory." *Experimental Brain Research* 76, no. 2 (July 1, 1989): 292–306. https://doi.org/10.1007/BF00247889.

Kotekal, Subhodh, and Jason N. MacLean. "Recurrent Interactions Can Explain the Variance in Single Trial Responses." *PLOS Computational Biology* 16, no. 1 (January 30, 2020): e1007591. https://doi.org/10.1371/journal.pcbi.1007591.

Mazurek, Mark, Marisa Kager, and Stephen D. Van Hooser. "Robust Quantification of Orientation Selectivity and Direction Selectivity." *Frontiers in Neural Circuits* 8 (August 6, 2014). https://doi.org/10.3389/fncir.2014.00092.

Nauhaus, Ian, Laura Busse, Matteo Carandini, and Dario L. Ringach. "Stimulus Contrast Modulates Functional Connectivity in Visual Cortex." *Nature Neuroscience* 12, no. 1 (January 2009): 70–76. https://doi.org/10.1038/nn.2232.

Niell, Cristopher M., and Michael P. Stryker. "Highly Selective Receptive Fields in Mouse Visual Cortex." *The Journal of Neuroscience* 28, no. 30 (July 23, 2008): 7520–36. https://doi.org/10.1523/JNEUROSCI.0623-08.2008.

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. https://doi.org/10.1523/JNEUROSCI.2177-15.2016.

Olshausen, Bruno A., and David J. Field. "How Close Are We to Understanding V1?" *Neural Computation* 17, no. 8 (2005): 1665–99.

Pajevic, Sinisa, and Dietmar Plenz. "The Organization of Strong Links in Complex Networks." *Nature Physics* 8 (2012): 429–36. https://doi.org/10.1038/nphys2257.

Ponce-Alvarez, A., A. Thiele, T. D. Albright, G. R. Stoner, and G. Deco. "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons." *Proceedings of the National Academy of Sciences* 110, no. 32 (August 6, 2013): 13162–67. https://doi.org/10.1073/pnas.1300098110.

Ponce-Alvarez, Adrián, Adrien Jouary, Martin Privat, Gustavo Deco, and Germán Sumbre. "Whole-Brain Neuronal Activity Displays Crackling Noise Dynamics." *Neuron* 100, no. 6 (December 19, 2018): 1446-1459.e6. https://doi.org/10.1016/j.neuron.2018.10.045.

Radicchi, Filippo, Santo Fortunato, Benjamin Markines, and Alessandro Vespignani. "Diffusion of Scientific Credits and the Ranking of Scientists." *Physical Review E* 80, no. 5 (November 11, 2009): 056103. https://doi.org/10.1103/PhysRevE.80.056103.

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. https://doi.org/10.1126/science.1179850.

Ringach, Dario L., Patrick J. Mineault, Elaine Tring, Nicholas D. Olivas, Pablo Garcia-Junco-Clemente, and Joshua T. Trachtenberg. "Spatial Clustering of Tuning in Mouse Primary Visual Cortex." *Nature Communications* 7 (August 2, 2016): 12270. https://doi.org/10.1038/ncomms12270.

Ringach, Dario L., Robert M. Shapley, and Michael J. Hawken. "Orientation Selectivity in Macaque V1: Diversity and Laminar Dependence." *Journal of Neuroscience* 22, no. 13 (July 1, 2002): 5639–51. https://doi.org/10.1523/JNEUROSCI.22-13-05639.2002.

Roxin, Alex, Nicolas Brunel, David Hansel, Gianluigi Mongillo, and Carl van Vreeswijk. "On the Distribution of Firing Rates in Networks of Cortical Neurons." *Journal of Neuroscience* 31, no. 45 (November 9, 2011): 16217–26. https://doi.org/10.1523/JNEUROSCI.1677-11.2011.

Rubinov, Mikail, and Olaf Sporns. "Complex Network Measures of Brain Connectivity: Uses and Interpretations." *NeuroImage*, Computational Models of the Brain, 52, no. 3 (September 1, 2010): 1059–69. https://doi.org/10.1016/j.neuroimage.2009.10.003.

Senden, Mario, Gustavo Deco, Marcel A. de Reus, Rainer Goebel, and Martijn P. van den Heuvel. "Rich Club Organization Supports a Diverse Set of Functional Network

Configurations." *NeuroImage* 96 (August 1, 2014): 174–82. https://doi.org/10.1016/j.neuroimage.2014.03.066.

Shi, Li, Xiaoke Niu, and Hong Wan. "Effect of the Small-World Structure on Encoding Performance in the Primary Visual Cortex: An Electrophysiological and Modeling Analysis." *Journal of Comparative Physiology A* 201, no. 5 (May 2015): 471–83. https://doi.org/10.1007/s00359-015-0996-5.

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii. "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits." Edited by Karl J. Friston. *PLoS Biology* 3, no. 3 (March 1, 2005): e68. https://doi.org/10.1371/journal.pbio.0030068.

Sun, Wenzhi, Zhongchao Tan, Brett D. Mensh, and Na Ji. "Thalamus Provides Layer 4 of Primary Visual Cortex with Orientation- and Direction-Tuned Inputs." *Nature Neuroscience* 19, no. 2 (February 2016): 308–15. https://doi.org/10.1038/nn.4196.

Wang, Peggy I., and Edward M. Marcotte. "It's the Machine That Matters: Predicting Gene Function and Phenotype from Protein Networks." *Journal of Proteomics*, Model organism proteomics, 73, no. 11 (October 10, 2010): 2277–89. https://doi.org/10.1016/j.jprot.2010.07.005.

Watanabe, Takamitsu. "Rich-Club Network Topology to Minimize Synchronization Cost Due to Phase Difference among Frequency-Synchronized Oscillators." *Physica A: Statistical Mechanics and Its Applications* 392, no. 5 (March 1, 2013): 1246–55. https://doi.org/10.1016/j.physa.2012.11.041.

Zariwala, Hatim A., Linda Madisen, Kurt F. Ahrens, Amy Bernard, Edward S. Lein, Allan R. Jones, and Hongkui Zeng. "Visual Tuning Properties of Genetically Identified Layer 2/3 Neuronal Types in the Primary Visual Cortex of Cre-Transgenic Mice." *Frontiers in Systems Neuroscience* 4 (2011). https://doi.org/10.3389/fnsys.2010.00162.

Zerlaut, Yann, Stefano Zucca, Stefano Panzeri, and Tommaso Fellin. "The Spectrum of Asynchronous Dynamics in Spiking Networks as a Model for the Diversity of Non-Rhythmic Waking States in the Neocortex." *Cell Reports* 27, no. 4 (April 23, 2019): 1119-1132.e7. https://doi.org/10.1016/j.celrep.2019.03.102.

Zylberberg, Joel. "Untuned but Not Irrelevant: The Role of Untuned Neurons in Sensory Information Coding." *BioRxiv*, January 1, 2017, 134379. https://doi.org/10.1101/134379.

SUPPLEMENTARY FIGURES

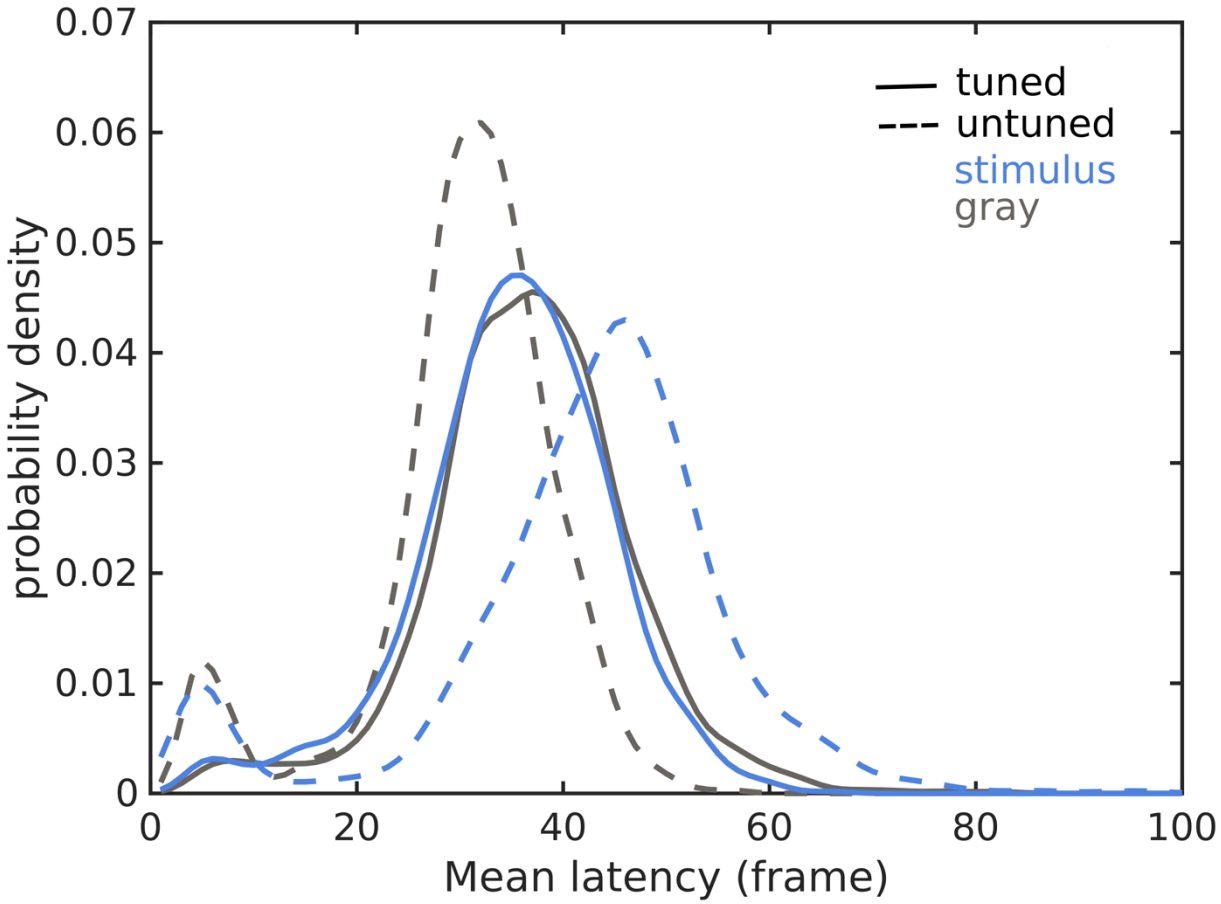

Figure 2.S1 – Tuned and untuned neurons activate with different latencies

For each neuron we inferred spikes (Friedrich et al., 2017) from fluorescence traces. We then identified the first imaging frame in which the neuron spiked in each stimulus or gray trial, and averaged those frames indices in stimulus and gray trials separately. Tuned neurons (solid lines) did not differ in their latency to activate in stimulus (blue, 33.95±34.66 frames) and gray (gray, 33.03±27.05 frames) trials, whereas untuned neurons activated significantly earlier in gray epochs (27.33±24.92 frames) compared to drifting gratings (38.69±38.64 frames, $F=9245.39$, $p<0.001$) .

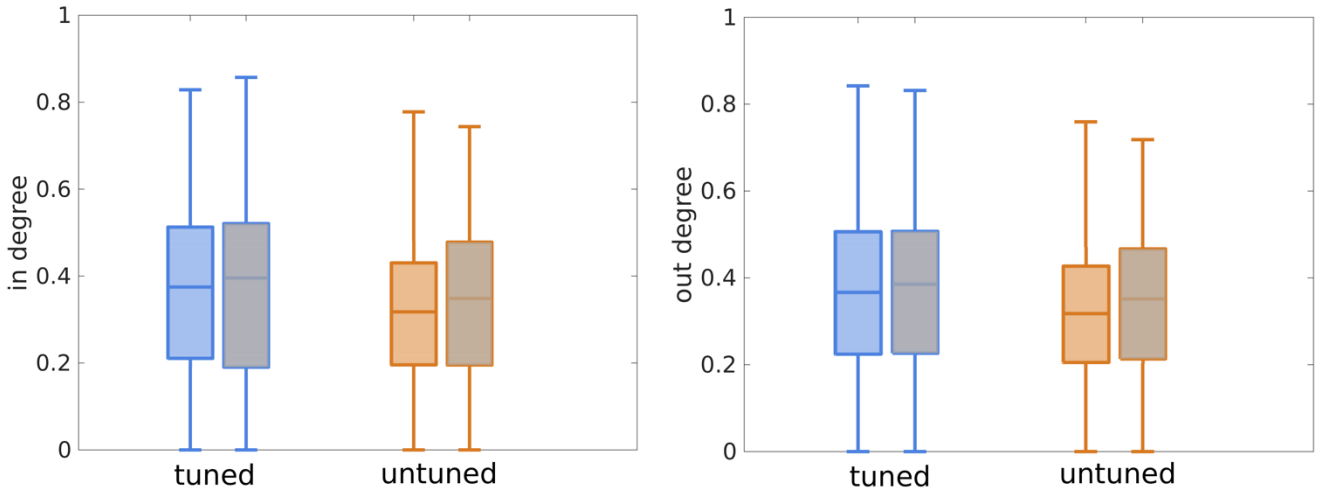

Figure 2.S2 – In- and out- degrees of neurons in sub-FNs are similar for stimulus and gray epochs

We kept only tuned-tuned edges in the tuned sub-FNs (blue) and untuned-untuned edges in the untuned sub-FNs (orange). Stimulus epochs are represented by blue/orange and gray epochs in gray. Boxplots show interquartile range and medians.

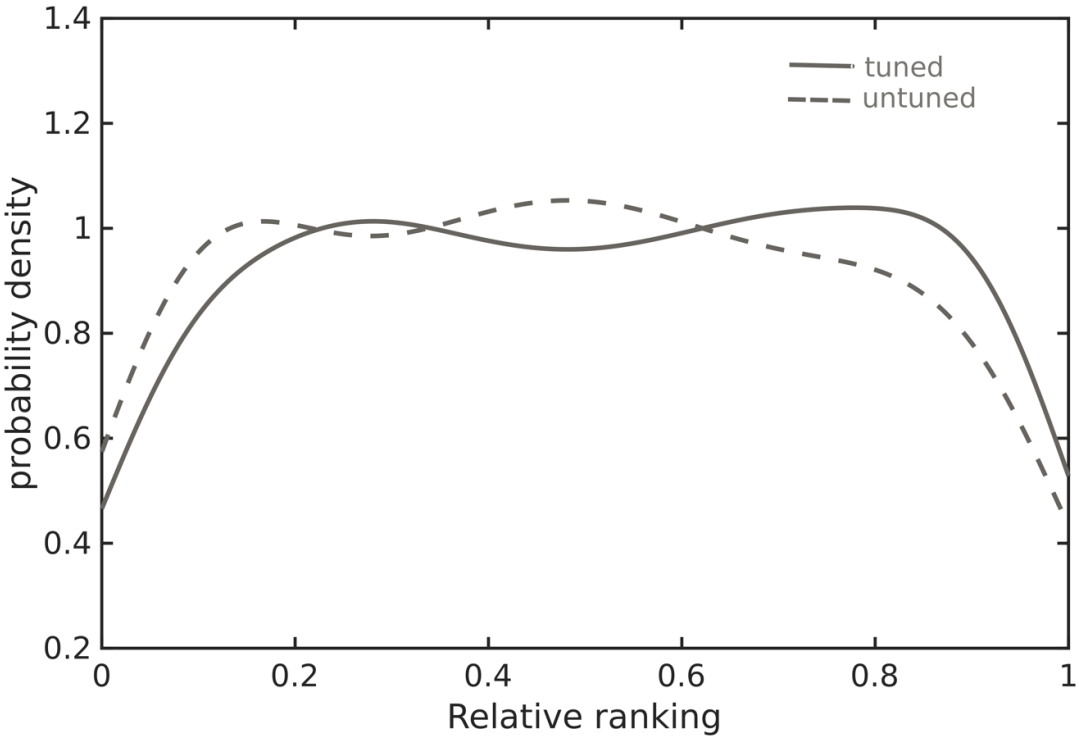

Figure 2.S3 – Tuned neurons show slightly larger relative rankings in gray FNs

We computed the relative ranking of tuned and untuned neurons in FN created from gray epochs (see Methods).

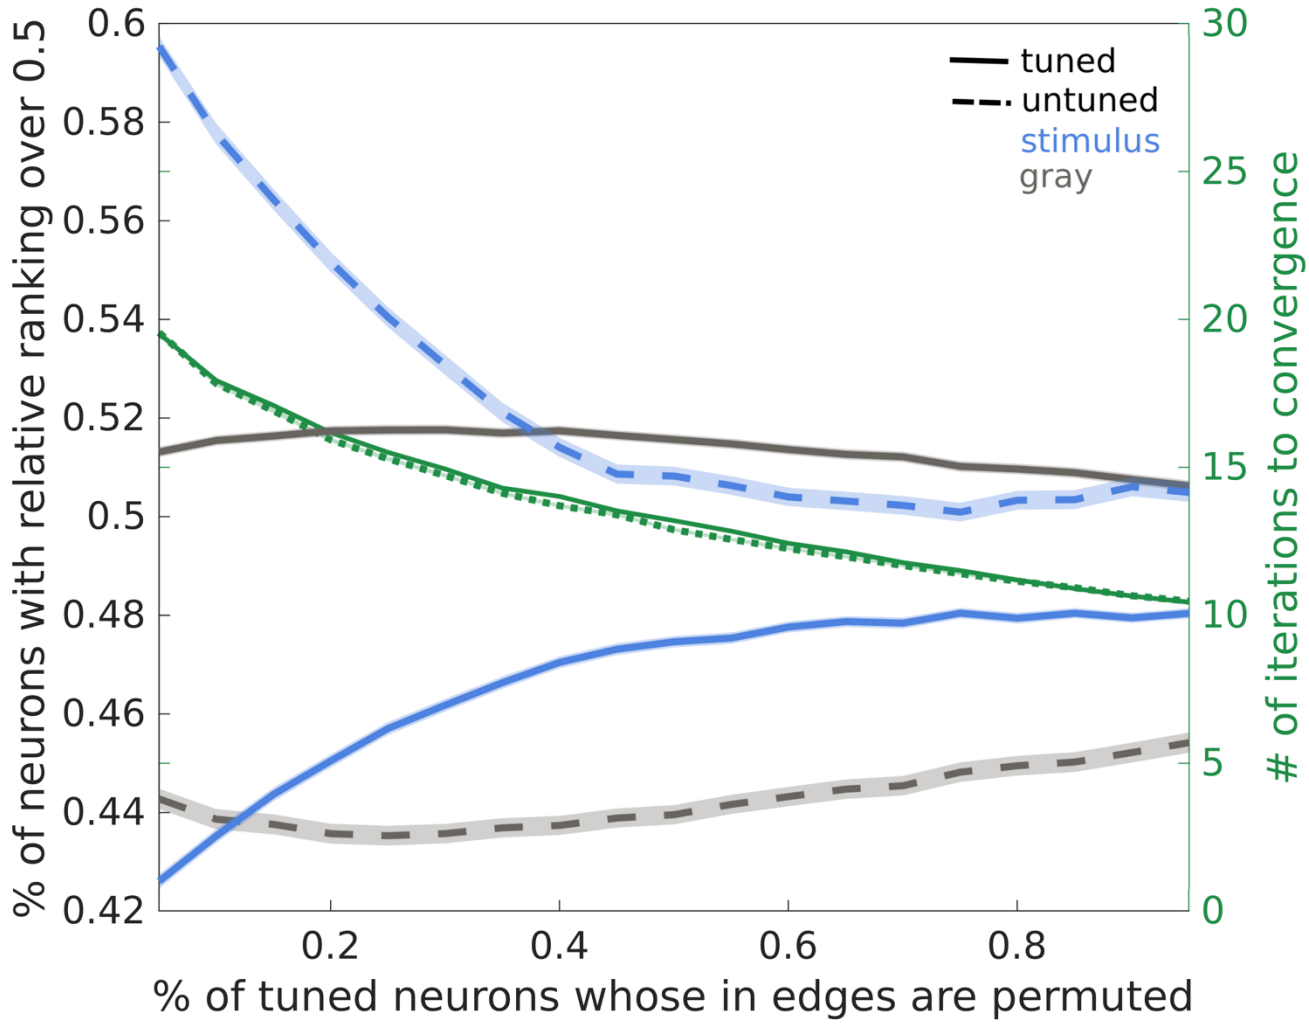

Figure 2.S4 – Permuting in-edges of tuned neurons causes a decrease in untuned neurons rankings

We gradually increased the portion of tuned neurons whose incoming edges we randomly permuted, in steps of 5%. For each step we ran 30 trials, in which we chose neurons at random and permuted their incoming edges at random. Lines represent means over datasets, with shading for the standard error. Green lines represent the time for convergence, with solid for stimulus and dotted for gray.

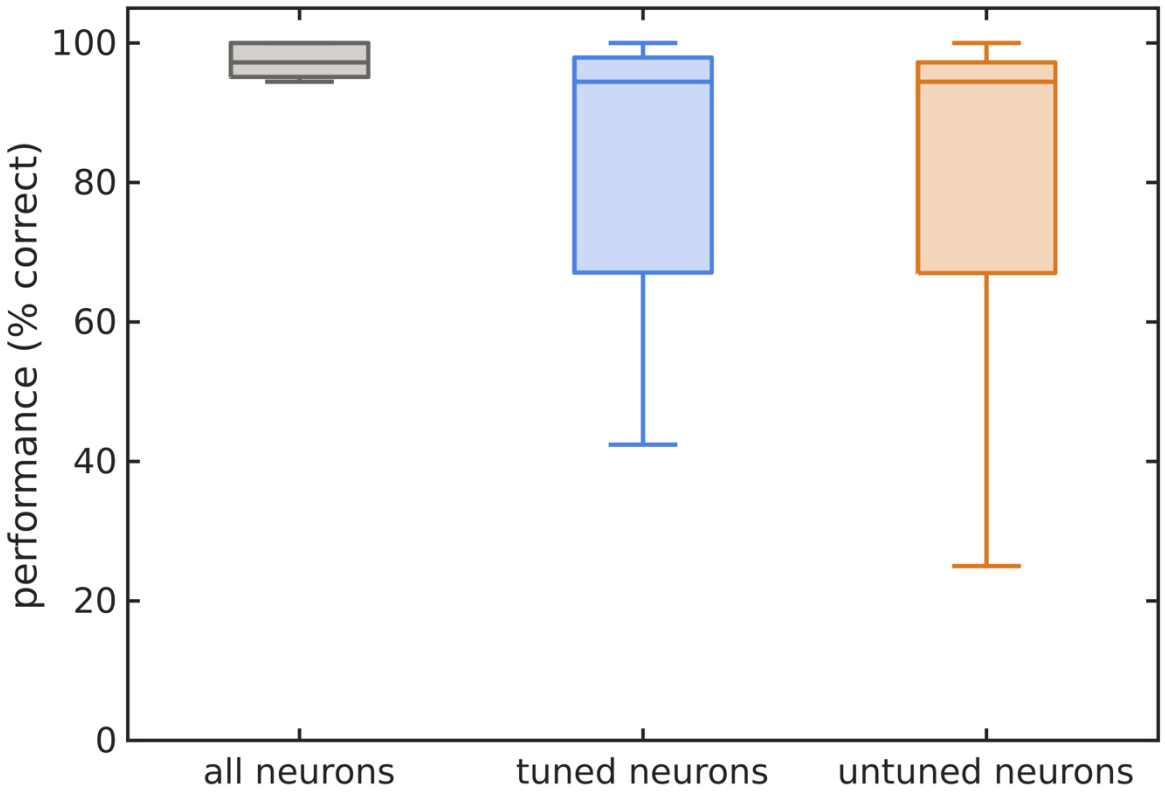

Figure 2.S5 – Training the decoder only with one functional sub-class results in similar degradation in performance

We trained the feed-forward network (second phase of the decoder) only with spikes from tuned (blue) or untuned (orange) neurons. The number of neurons we trained with was kept the same by randomly choosing neurons from the more frequent sub-class. We did 100 such manipulations. Boxplots show interquartile range and medians.

CHAPTER III

Temporal pairwise correlations reveal sparse subsets of highly variable spikes that code for the visual stimulus

This work will be published as a preprint on Summer 2021: Levy, M.*, Guo, J.K.* & MacLean, J.N. (2021). Temporal pairwise correlations reveal sparse subsets of highly variable spikes that code for the stimulus. *BioRxiv* (*co-first authors / equal contribution)

ABSTRACT

The spiking activity of large groups of neurons underlies the representation of both internal and external variables. However, it remains unclear whether specific subsets of spikes can be ascribed to one class of variable or the other. To complicate matters, the activity of individual neurons is highly variable belying a straightforward assignment of spikes to variables. Despite this variance statistical regularities can be found in multineuronal activity and contain stimulus related information. What remains unclear is whether these regularities are meaningful over the time course of a single trial. Here we intersect pairwise statistical regularities and pairs of action potentials recorded during single trials in murine primary visual cortex (V1) and identify a sparse set of spikes. These spikes correspond to particularly strong pairwise statistical regularities between neurons, but still exhibit trial-to-trial variability relative to stimulus. We find that these subsets carry more information about drifting grating direction as compared to the entirety of spikes or any other subset, and can be decoded from with high accuracy. Our findings suggest that consistent spatio-temporal patterns, delineated by pairs of neurons, might be an organizational principle of the coding scheme of neocortex. The proposed scheme elaborates on the Hebbian

assembly hypothesis and proposes a biologically plausible and highly efficient way for the same spike to represent multiple variables.

INTRODUCTION

The neocortical response to ostensibly identical sensory input is highly variable, including spike rate and timing at both the levels of individual and populations of neurons[1,2,3]. Under the assumption that action potentials evoke changes in the activity of downstream neurons, this trial-to-trial variability poses a challenge to assigning meaning to spikes, i.e. both the representation and subsequent downstream decoding of a sensory stimulus. However, neural activity within subsets of neurons, or circuits, is not statistically independent (for reviews see [4,5,6]). Intuitively, knowledge about the activity of functionally connected neighbors (often referred to as peer prediction[7]) should provide better prediction of a target neuron activity. Indeed, successful prediction of population statistics and circuit dynamics from pairwise correlations has been demonstrated in retina[8,9], motor cortex[10] and visual cortex[11,12]. Moreover, covariability is relevant from the perspective of a downstream neuron since it is more likely to spike given multiple

[1] Ventura, "Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons."
[2] Tolhurst, Movshon, and Dean, "The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex."
[3] Deweese and Zador, "Shared and Private Variability in the Auditory Cortex."
[4] Averbeck, Latham, and Pouget, "Neural Correlations, Population Coding and Computation."
[5] Cohen and Kohn, "Measuring and Interpreting Neuronal Correlations."
[6] Salinas and Sejnowski, "Correlated Neuronal Activity and the Flow of Neural Information."
[7] Harris, "Neural Signatures of Cell Assembly Organization."
[8] Ganmor, Segev, and Schneidman, "Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code."
[9] Pillow et al., "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population."
[10] Stevenson et al., "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons."
[11] Kotekal and MacLean, "Recurrent Interactions Can Explain the Variance in Single Trial Responses."
[12] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."

temporally coordinated inputs[13,14,15]. Hence, correlated activity can lead to more stable representation downstream despite single neuron variability.

However, there is an ongoing debate whether correlations between neurons enhance or impair the encoding of internal and external variables[16,17,18]. The controversy might be in part due to the nature of correlations considered; some studies have reported global correlations, arising from an internal state and population-wide fluctuations[19] or modulation of correlations across all pairs[20]. Others have found that correlations can be stimulus-specific[21,22,23]. To clarify, stimulus-specific correlations are not akin to signal correlations, which measure the similarity of neurons across stimuli[24]. Rather, stimulus-specific correlation are conceptually similar to noise correlations, that quantify covariability across trials of the same stimulus. A straightforward way to understand this is to imagine an experiment in which an animal is presented with left or right moving dots. If two neurons covary in left trials but not (or less reliably so) in right trials, the pair displays a stimulus-specific correlation. It is likely that global and stimulus-specific correlations coexist and methods have been proposed to tease them apart[25,26]. It was also shown that stimulus-

[13] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."
[14] Salinas and Sejnowski, "Impact of Correlated Synaptic Input on Output Firing Rate and Variability in Simple Neuronal Models."
[15] Usrey, Reppas, and Reid, "Paired-Spike Interactions and Synaptic Efficacy of Retinal Inputs to the Thalamus."
[16] Ohiorhenuan et al., "Sparse Coding and High-Order Correlations in Fine-Scale Cortical Networks."
[17] Averbeck and Lee, "Effects of Noise Correlations on Information Encoding and Decoding."
[18] Kanitscheider, Coen-Cagli, and Pouget, "Origin of Information-Limiting Noise Correlations."
[19] Ecker et al., "State Dependence of Noise Correlations in Macaque Primary Visual Cortex."
[20] Cohen and Maunsell, "Attention Improves Performance Primarily by Reducing Interneuronal Correlations."
[21] Franke et al., "Inhibition Decorrelates Visual Feature Representations in the Inner Retina."
[22] Levy, Sporns, and MacLean, "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding."
[23] Ponce-Alvarez et al., "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons."
[24] Cohen and Kohn, "Measuring and Interpreting Neuronal Correlations."
[25] Granot-Atedgi et al., "Stimulus-Dependent Maximum Entropy Models of Neural Population Codes."
[26] Goris, Simoncelli, and Movshon, "Origin and Function of Tuning Diversity in Macaque Visual Cortex."

specific correlations have information about the stimulus[27,28] Yet the interplay of these two types of correlations at beyond pairwise level, their structure as well as their relevance to coding remains unclear. To be relevant from a coding perspective, a partition of correlations is required to identify spikes that participate in stimulus or non-stimulus specific correlations with other spikes in a biotically plausible way that can implemented downstream.

Here we utilize a functional network (FN) framework in which spiking dependencies between neurons are summarized as graph where neurons are nodes and correlations are directed and weighted edges[29,30,31,32]. We extend this approach to isolate and select subsets of edges that are stimulus-specific and non-stimulus-specific. In previous work we examined which synapses are activated in a way that contributes to dynamics over entire simulation[33,34,35]. Here we use a similar technique, with several modifications: we work with functional networks on temporally fine-scale, to identify spikes that are an instantaneous manifestation of lagged pairwise correlations. We identify pairs of spikes that correspond to reliable statistical dependencies that are stimulus-specific and these spikes are very sparse and highly variable across trials. Nonetheless, they are more informative of the visual stimulus regardless of the decoding algorithm used. Our findings that trial-to-trial variability in single neurons and population, relative to stimulus, does not constitute an obstacle for the sensory code.

---

[27] Franke et al., "Structures of Neural Correlation and How They Favor Coding."
[28] Montani and Schultz, "Information-Theoretic Analysis of the Role of Correlations in Neural Spike Trains."
[29] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[30] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."
[31] Dann et al., "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates."
[32] Downes et al., "Emergence of a Small-World Functional Network in Cultured Neurons."
[33] Bojanek, Zhu, and MacLean, "Cyclic Transitions between Higher Order Motifs Underlie Sustained Asynchronous Spiking in Sparse Recurrent Networks."
[34] Chambers et al., "Ensemble Stacking Mitigates Biases in Inference of Synaptic Connectivity."
[35] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."

METHODS

Data collection and curation

Animals and protocols are described in full in [36]. Briefly, we performed a craniotomy over the left primary visual cortex (V1) in 8 Tg(Thy1-GCaMP6s)GP4.12Dkim (Jackson Laboratory) mice (4 male, 4 female). These mice continuously express GCaMP6s in excitatory pyramidal neurons of layer 2/3. Upon recovery and verification of V1 location, mice were head-fixed but free to run on a linear treadmill while passively viewing drifting gratings in 12 evenly-spaced directions (80% contrast, 0.04 cyc/deg spatial frequency and 2Hz temporal frequency). Stimulus presentation (ON epochs) were 5s in duration and interleaved with 3s of gray screen (OFF epochs). Imaging of the neuronal population was performed with a line-scan[37] at a wavelength of 910nm (Coherent Chameleon). We inferred spikes from df/f by employing a deconvolution algorithm[38]. All data presented are from 19 datasets unless stated otherwise.

Functional networks

We summarized the pairwise relationship between neurons in a static functional network (FN), computed as the confluent mutual information of the spikes of every pair of neurons i,j. Confluent mutual information (conMI[39], Fig. 1A) is defined as:

$$(1) \qquad conMI_{ij} = \sum_{i(t)\in\{0,1\}} \sum_{j(\hat{t})\in\{0,1\}} p(i(t), j(\hat{t})) \cdot \log_2 \left[\frac{p(i(t), j(\hat{t}))}{p(i(t)) \cdot p(j(\hat{t}))} \right]$$

[36] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[37] Sadovsky et al., "Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging."
[38] Friedrich, Zhou, and Paninski, "Fast Online Deconvolution of Calcium Imaging Data."
[39] Chambers et al., "Ensemble Stacking Mitigates Biases in Inference of Synaptic Connectivity."

where: $j(\hat{t}) = \begin{cases} 1 & j(t-1) = 1 \ OR \ j(t) = 1 \\ 0 & otherwise \end{cases}$

And is therefore none-symmetric. This results in an NxN adjacency matrix, with neurons being nodes and conMI values as directed and weighted edges between them. Edges between neuronal pairs with negative correlation between their spike trains were then set to zero, and the adjacency was further pruned to contain only the top 50% of edge weights, resulting in a final density of 0.38±0.04. Functional network construction was performed separately for trials of each direction of drifting grating, for a total of 12 FNs per dataset.

Edge classification into four sub-FNs

For each direction θ of drifting gratings, 3 fellow directions were defined: the two neighboring directions that are 30 degrees apart from θ, and the opposite direction of θ, 180 degrees apart, which corresponds to the same orientation as θ. For example, for θ=60 the fellow directions are 30, 90 and 240 degrees. We then compared each edge to the values in the three according to the following rules:

(2) $\begin{cases} FN_{ij}^{\theta} \neq 0 \\ \forall f \in fellow, \ FN_{ij}^{f} = 0 \end{cases}$ ⇒ hard unique (**hard u**)

(3) $\begin{cases} FN_{ij}^{\theta} \neq 0 \\ \exists f \in fellow, FN_{ij}^{f} \neq 0 \\ \forall f \in fellow, FN_{ij}^{f} < \frac{1}{2} FN_{ij}^{\theta} \end{cases}$ ⇒ soft unique (**soft u**)

(4) $\begin{cases} FN_{ij}^{\theta} \neq 0 \\ \exists f \in fellow, FN_{ij}^{f} = 0 \\ \exists f \in fellow, FN_{ij}^{f} \geq \frac{1}{2} FN_{ij}^{\theta} \end{cases}$ ⇒ partially shared (**p shared**)

(5) $$\begin{cases} FN_{ij}^{\theta} \neq 0 \\ \forall f \in fellow, FN_{ij}^{f} \neq 0 \\ \forall f \in fellow, FN_{ij}^{f} \geq \frac{1}{2} FN_{ij}^{\theta} \end{cases} \Rightarrow \text{fully shared (}\mathbf{f\ shared}\text{)}$$

In words, hard u contains edges that are present only in θ. Soft u is composed of edges that may be present in one or more of the fellow directions FNs, but are at most half the weight of the edge in θ. Edges that are sorted into p shared are present in some of the fellow directions, and are roughly the same magnitude as in θ. Finally, f shared edges are present in the FN of all fellow directions, and are of similar weight to the weight in the FN for θ. This procedure was carried out for every direction of drifting gratings, resulting in 4 exclusive sub-functional networks (sub-FNs) for each of the 12 directions.

Edge expression and spike sparsification

A temporal graph (TG) is a single-trial moment-to-moment representation of which pairwise relationships are instantiated. To construct a temporal graph, we first built a binary tensor of potential edges (POT) of size NxNxT-1 with T being the time points in the rasters (R). Each time slice $\hat{t}$, summarized the spiking activity at t-1 and t. For each neuron pair i,j at each time point sliding along the duration of the spiking activity:

(6) $$POT_{i,j,\hat{t}} = 1 \text{ if } R_{i,t-1} = 1 \; and \; R_{j,t} = 1$$

else

$$POT_{i,j,\hat{t}} = 0$$

POT is called a tensor of potential edges since it may be the case that i spiking at t-1 has contributed to the spiking activity of j at t. POT is thus a combinatorial representation of the rasters and its density depends on the firing rate in the population.

Each time slice $\hat{t}$ was then intersected with an FN, or sub-FN, as indicated in the results:

(7) $$TG_{\hat{t}} = POT_{\hat{t}} \odot FN^{\theta}$$

TG is thus the same size as the potential edges tensor, but is way sparser and contains weighted edges. Expressed edges are the edges that remain in the TG after intersection out of the total number of edges in the FN.

We extracted spikes that are consistent with a temporal graph (TG) by using a process that is the reverse of the potential edges construction. First, we binarized the TG, setting all non-zero elements to 1. Then, for every neuron pair i,j and every time slice , we defined the sparsified rasters (SR) to be:

(8) $$SR_{i,t-1} = 1 \; and \; SR_{j,t} = 1 \text{ if } TG_{i,j,\hat{t}} = 1$$

In practice, instead of iterating over every pair of neurons, the same can be achieved summing the over rows and columns of each time slice $TG_{\hat{t}}$. A spike at t-1 is then assigned to every neuron with a non-zero column sum, and spike at t is assigned to every neuron with a non-zero row sum.

Note that no new spikes are inserted, rather, spikes in SR are a subset of the spikes in the original raster R. SR contains only the spikes in R that can be explained as a manifestation of the FN or the sub-FN that was used in the calculation of the TG.

Trial-to-trial variability measures

For single cell rate variability (Fig. 3D), for each neuron we counted the spikes over the course of the trial. We computed the coefficient of variation, CV as σ/μ with σ and μ being the standard deviation and mean, respectively, of the neuron spike-count across trials with the same drifting gratings direction.

For temporal precision of single cells, for each neuron we calculated the Victor-Purpura (VP) distance[40] between pairs of trials of the same direction. Since VP is sensitive to the number of spikes (i.e. sparsity) we divided by the mean spike count for the trials in the pair. We chose q=1 for the cost.

Population-level variability was measured by the L2 (Euclidean) norm between each pair of population vectors for trials of the same drifting grating direction. We formed population vectors by counting the spikes for each neuron across the duration of the trial. The L2 metric was normalized by the mean of the total spikes in the trials in the pair to adjust for sparsity differences.

Information quantification

The mutual information (in bits) between the neural response *r* and the stimulus feature *s* was calculated as[41,42]:

$$I(r;s) = \sum_s P(s) \sum_r P(r|s) \log_2 \left[\frac{P(r|s)}{P_r(r)}\right] \quad (9)$$

Neural responses analyzed were binarized vectors of spikes for 5 neurons binned across 10 imaging frames (Fig. 4A), and vectors of spikes for 5 neurons across 3 imaging frames concatenated into 15-by-1 vectors (Fig. 4B). The identities of the 5 neurons were selected by taking those with the highest firing rate across the first 100 imaging frames of all trials in the original spikes and used to compute the mutual information for all four sub-FN consistent spikes. Across all trials, the value of P(s) is uniform and equal to 1/12, and we scaled raw information value by the average number of spikes in a vector to obtain bits per spike.

[40] Victor and Purpura, "Nature and Precision of Temporal Coding in Visual Cortex."
[41] Brenner et al., "Synergy in a Neural Code."
[42] Palmer et al., "Predictive Information in a Sensory Population."

Feedforward pattern recognition neural network decoder

We used the decoder described in [43]. This is a multiclass decoder with N input units for N neurons, and 12 output units, one for each direction of drifting gratings. The input and output layer are connected by all-to-all feedforward connectivity with random initial weights. As inputs, we binned spikes in 10 consecutive imaging frames to a population vector of size N. 90%-10% of the trials were used as training and test sets, respectively. The weights were trained by conjugate gradient with Matlab's machine learning toolbox. We trained the weights for the original spikes and the sub-FN consistent spikes separately.

Support vector machine (SVM) decoder

To confirm the results of the feedforward neural network decoder, we used the same inputs, i.e. population vectors of spikes binned at 10 frames from 90% of the trials, to train a Support Vector Machine (SVM) decoder with a linear kernel. Training was performed with Matlab's fitcecoc.m function. We tested on the remaining 10% of trials.

Naive Bayes decoder

For each imaging frame in the data, the number of spikes for each neuron were binned across 10 frames and concatenated together for all trials in a dataset. The concatenated spike counts and their corresponding stimulus direction were randomly split into training and testing data in a 90-10 ratio.

[43] Levy, Sporns, and MacLean, "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding."

This Naive Bayes Decoder[44,45] used a Gaussian likelihood function. Using the concatenated spike counts in the training data, for each neuron, the mean and standard deviation for the number of spikes binned across 10 frames was calculated for each stimulus direction, giving the mean and standard deviation for the likelihood function for each stimulus direction for each neuron. Values of 0 for the sample variance were replaced with the average variance across all trials for subsequent normal pdf calculations. Then, using the testing data, the posterior probability of each stimulus direction given the observed spike counts were calculated by finding the product of the likelihood functions for all neurons evaluated at the observed spike counts and the stimulus direction giving the maximum posterior probability was the decoded direction. This procedure was repeated 30 times for each starting frame and conducted for the first 100 starting frames for the original spikes and the sub-FN consistent spikes.

In the shuffled control for the Bayesian decoding, prior to binning spike counts, the spikes were shuffled within each trial to preserve trial spike density. For each imaging frame, the shuffled spikes were then concatenated as above and the direction was decoded using the same procedure. The decoder performance of the unshuffled spikes was divided by that of the shuffled spikes on a frame-by-frame basis to obtain the fold change of the decoding accuracy relative to the density-matched control.

Controls

To test for the possibility that any spiking activity that is organized in pairs of consecutive spikes accounts for the decoding performance, we designed a stringent control that preserves the density of the sub-FNs as well as the resulting pairwise structure of spikes. We randomly permuted the

[44] Zhang et al., "Interpreting Neuronal Population Activity by Reconstruction."
[45] Glaser et al., "Machine Learning for Neural Decoding."

edges in each sub-FN before intersecting with each trial. We then binned the spikes in each 10 frames as with the real spikes, and passed them through the SVM decoder as described above.

Statistical analysis and code

Means and standard deviations across datasets or neurons are reported throughout the paper as M±SD unless stated otherwise. p values were bonferroni corrected in cases of multiple comparisons. All analysis was done in Matlab 2018 or later (Mathworks) and Python 3.7.

RESULTS

Functional networks divided according to stimulus specificity have distinct topological hallmarks

We imaged hundreds of excitatory neurons in layer 2/3 of mouse primary visual cortex (V1) in response to 12 directions of drifting gratings[46]. Here we summarize the spiking dynamics of these populations of neurons as functional networks (FNs), with neurons and the statistical dependencies between them as nodes and edges, respectively. We generated a functional network for each of the 12 directions of drifting gratings by computing the confluent mutual information (conMI[47]; Fig 1A) between each pair of neurons in all trails of each grating direction separately. conMI is the mutual information between neuron i at time t and neuron j at time t and t+1, with time dictated by imaging frame.

When comparing each of the twelve FNs we observed that some edges in the FNs are overlapping across stimuli, whereas other edges are present only in a subset or a single FN

[46] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[47] Chambers et al., "Ensemble Stacking Mitigates Biases in Inference of Synaptic Connectivity."

corresponding to a single direction of drifting gratings (Fig. 1B)(11.53±1.62% unique edges, 12.96±2.66% edges shared between two directions, 75.51±3.98% edges that are present in more than 2 FNs). Furthermore, the similarity of FNs correlates with stimulus similarity[48] (Fig. 1C), with increased similarity for directions that are adjacent (e.g. 30 and 90 are adjacent to 60) and directions that have the same bar orientation (e.g. 60 and 240). These data would suggest that these three (fellow) directions are presumably more difficult for downstream circuits to disambiguate. With this in mind we allocated edges in each FN into four sub-FNs depending on the extent to which an edge was shared or unique to one FN corresponding to one direction of drifting grating. Specifically edges were sorted into: hard u(nique) sub-FN which contains edges that only exist in an FN for a single direction, soft u(nique) edges may exist in FNs for fellow directions, but they are stronger, reflecting more reliable statistical dependency, in an FN for one direction. P(artial) shared which consists of edges that are common to a direction and to some, but not all fellow directions, and f(ully) shared edges that are found in all of the FN regardless of direction. This sorting procedure was done for each direction of drifting grating and results in 4 non-overlapping sub-FNs (Fig. 1D). Across datasets and directions, 17.41±5.06% edges were classified at hard u, 6.89±2.49% edges as soft u, and 53.46±4.31%, 22.22±9.02% edges as p shared and f shared, respectively, making soft u the sparsest sub-FN and p shared the densest sub-FN (Fig. 1E). As expected given our segmentation procedure we found that soft u contains a significant subset of the strongest edge weights (AVONA, $p<0.001$, Fig. 1F), meaning more of the edges of this particular sub-FN reflected more reliable statistical dependencies as compared to the other three sub-FNs. Almost all neurons had at least one edge and thus participated in each sub-FN, meaning,

[48] Levy, Sporns, and MacLean, "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding."

sub-FNs included the vast majority of neurons (hard u: 99.39±0.99%, soft u: 95.15±6.06%, p shared: 99.15±1.14% and f shared: 91.80±4.17%).

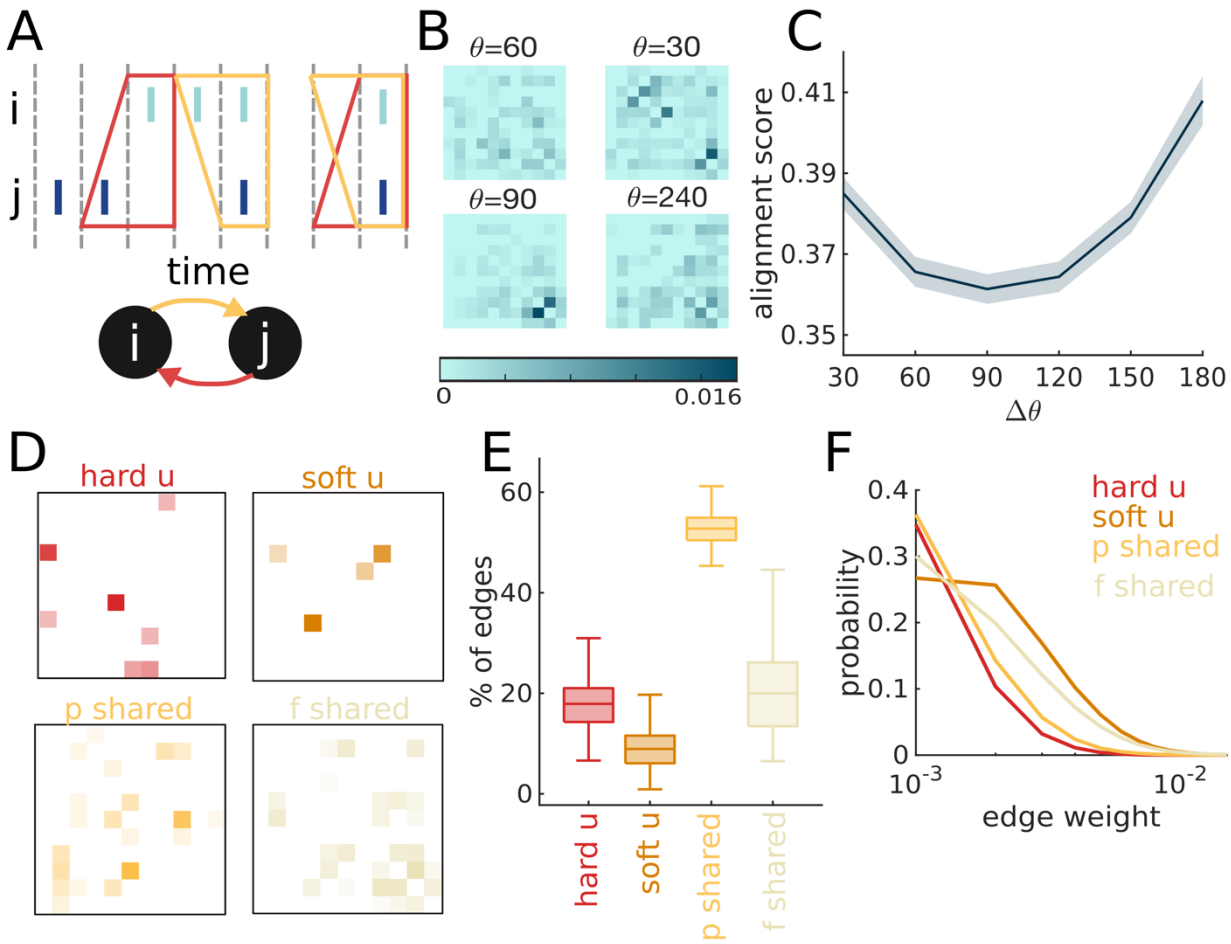

Figure 3.1 - Functional networks (FNs) can be divided to 4 sub-FNs based on edge statistics

A Illustration of edge inference by confluent mutual information (conMI). The edge from neuron i to neuron j is a statistical dependency between i spiking at time t-1 or t, and j spiking at time t. Hence conMI is not necessarily symmetric.

B FNs for 4 directions: the direction of interest (DOI; 60 degrees) and its three fellow directions (30,90 and 240 degrees). 10 neurons are illustrated here for visualization purposes. Colors represent edge weight. Note that some edges are similar across directions whereas others appear only in one FN.

C FN similarity reflects stimulus similarity: alignment score varies in [0,1] with larger numbers indicating increased similarity between two networks. $\Delta\theta$ stands for stimulus similarity in degrees, with 30 degrees being an adjacent direction and 180 degrees being the direction with the same orientation. Line and shading represent mean and the standard error, respectively.

Figure 3.1 - continued

D Illustration of edges allocation into 4 sub-FNs: the FN for 60 degrees from B was split into 4 non-overlapping sub-FNs, according to which edges are unique and strong in the DOI (60 degrees) as compared to the FNs for fellow directions. Colors for sub-FNs are consistent throughout this manuscript.

E Density of the four sub-FNs as the portions of edges from the overall allocated to each of them. Boxplots represent interquartile ranges and midlines mark the median.

F Probability distributions for the edge weights included in each of the four sub-FNs

We found that the number of reciprocal connections varied between the four FN sub-types and depended on the extent to which edges were unique to a direction specific FN. Shared sub-types displayed more bi-directional edges, adjusted for the FN density, and the portion of reciprocal edges was a function of the initial reciprocity before splitting up the FN in sub-FN. In contrast, unique sub-types had exhibited less reciprocal connections that were relatively fixed, regardless of the prevalence of bi-directional connections in the original FN (Fig. S1).

We then evaluated the extent to which each sub-FN may be an indication of the tuning of individual neurons. We counted how many edges in each sub-FN are amongst pairs of tuned (t-t), untuned (ut-ut), and mixed-functionality (t-ut, ut-t) neurons. The prevalence of each edge type was proportional to the density of two of the sub-FNs: soft u and f shared include the same portions of the three edge-types as the overall FN. Hard u exhibits a lower prevalence of tuned-tuned connections, whereas p shared shows an over representation of those pairs (Fig. S2).

Real-time correspondence of pairs of spikes to each of the four sub-FNs is organized in higher-order patterns

We next examined how statistical dependencies are dynamically realized by corresponding pairs of action potentials. To do so we created an instantaneous, potential functional network for every two consecutive frames (methods). In this potential FN, a pair of consecutive spikes where neuron

i spikes at time t and neuron j spikes at time t+1 is represented as 1 in the ij-th location, and 0 otherwise. Intersecting this potential FN with any of the sub-FN retains only the edges that are active at that time, termed expressed edges, and can be thought of as a temporal network (Fig. 2A).

Overall we found elevated expression during the first third of the trial and this correspondence declined thereafter. We then evaluated the extent to which the correspondence between spikes and edges depended on the sub-FN. Since sub-FNs vary in their density, we took the portion of expressed edges out of the total density. We found that soft u (0.20±0.53%) and f shared (0.19±0.53%) had larger levels of expression as compared with the overall FN and the other two sub-FNs (0.16±0.36% and 0.16±0.42% for hard u and p shared respectively, Fig. 2B). Notably, correspondence between the sub-FNs and pairs of action potentials during gray (OFF) visual stimulus epochs was significantly lower (ON epochs=0.18±0.48%, OFF epochs=0.07±0.23%, across subtypes and time, $p<0.001$).

We observed that some edges, that corresponded to pairs of action potentials, are also organized into higher-order patterns. A feedforward pattern (FF) is composed of two edges from i to j at t and t+1, whereas a recurrent pattern (RR) has an edge from i to j at time t, followed by an edge from j to i at time t+1 (Fig. 2A). On average a neuron participated in 0.63±1.37 patterns in a trial (across all subtypes, Fig. 2C). Adjusted for density, more of the edges that corresponded to pairs of spikes in the p shared and f shared FNs also were components of higher-order patterns as compared to the unique subtypes (Fig. 2D). We found that hard u, p shared and f shared exhibited elevated counts of recurrent patterns (7.81±3.34, 11.25±5.02 and 11.95±7.54, respectively), as compared to feedforward patterns (2.87±1.92, 3.22±2.28 and 2.97±2.25, respectively). In contrast, soft u exhibited balanced counts of FF (2.44±1.85) and RR (4.73±3.20) patterns (Fig. 2E and Fig. S3).

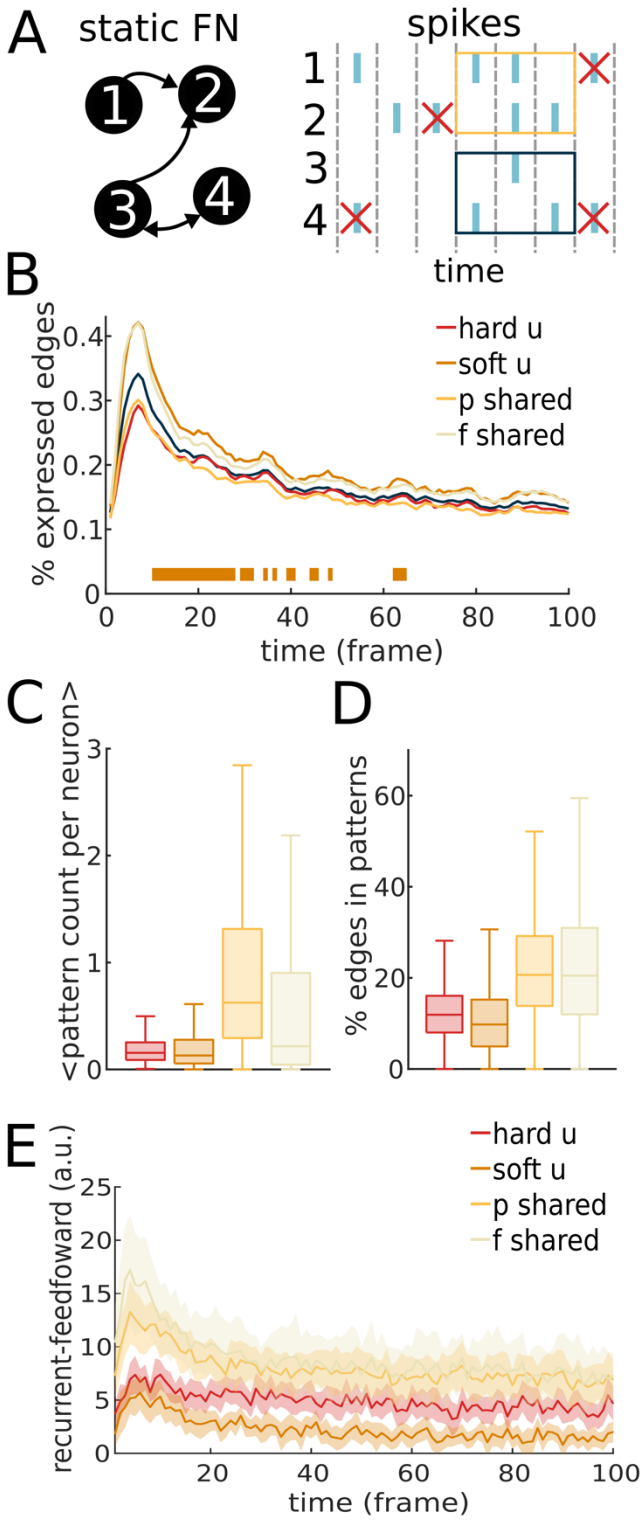

Figure 3.2 - Sub-FNs have different hallmarks of real-time edge expression

A Illustration of a temporal graph (TG) approach: a static FN which is inferred from many trials is intersected with spikes on a single-trial time-point by time-point basis. An edge from neuron 1 to neuron 2 is expressed in the first two time-frames with neuron 1 spiking followed by a spike in neuron j. Any static FN can be used in this procedure, and we have used both the overall FN and the sub-FNs. Spikes that are not an expression of any edges are marked with a red X and later discarded in the spike sparsification process. Example feedforward and recurrent patterns are marked in yellow and dark blue, respectively.

B Percent of expressed edges out of the overall FN (blue) and the four sub-FNs throughout the course of an ON epoch. Lines represent means across trials and datasets. Orange bars on the bottom indicate the time-points at which soft u had significantly larger expression (ANOVA, $p<0.05$).

C Neurons participate in more patterns in shared types. <> denote the mean over trials, and boxes indicate the interquartile range with the median marked as a horizontal line.

D Percent of the expressed edges that are taking part in patterns, out of all expressed edges. Boxes and horizontal lines represent the interquartile range and median, respectively.

E Recurrent pattern counts minus feedforward pattern counts for each time point in the trial. 0 would be a perfect balance between these two patterns. Lines and shading are means and standard errors across datasets.

Spikes that correspond to sub-FNs are sparse and variable

Feedforward connectivity has been implicated in the transmission of information and we postulated that this subset of correlations may be indicative of pairs of action potentials that would be more decodable by downstream elements. In order to test this hypothesis we focused on the action potentials rather than edges and segmented spikes according to correspondence to an edge in a sub-FN. In other words, if neuron i spiked at time t and neuron j spiked at time t+1, we kept those spikes if there was a non-zero edge in the ij-th index in the FN, and discarded spikes that did not correspond to an edge (Fig. 2A). This procedure yielded 4 sets of spikes that corresponded to edges in each sub-FN for each trial. Sub-FN consistent spikes are sparser than the original spikes, with soft u corresponding to the fewest action potentials (Fig. 3A). In this framework a spike can correspond to an edge in more than one sub-FN ; for example, an ij edge may exist in hard u, while an ik edge is nonzero in p shared. In this case, a spike at time t will be retained for both of these subtypes if j and k spiked at time t+1, respectively. Therefore, while edges are exclusive to one of the four sub-FNs, the corresponding spikes are not. We found varying degrees of overlap between sub-FN consistent spikes for different sub-FNs (Fig. 3B-C) and the magnitude of overlap was stationary.

Sparsifying the spikes by keeping only sub-FN consistent spikes did not reduce the trial-to-trial variability of spikes in single neurons as measured by rate or temporal precision (Fig. 3D and 3F, respectively). This was surprising since these spikes were selected according to correspondence to edges in sub-FNs and edges are indicative of pairwise statistical dependencies. Similarly we found that network wide sub-FN consistent spikes were not more reliable than the original spikes as a population vector as measured with the L2 norm (Fig. 3E).

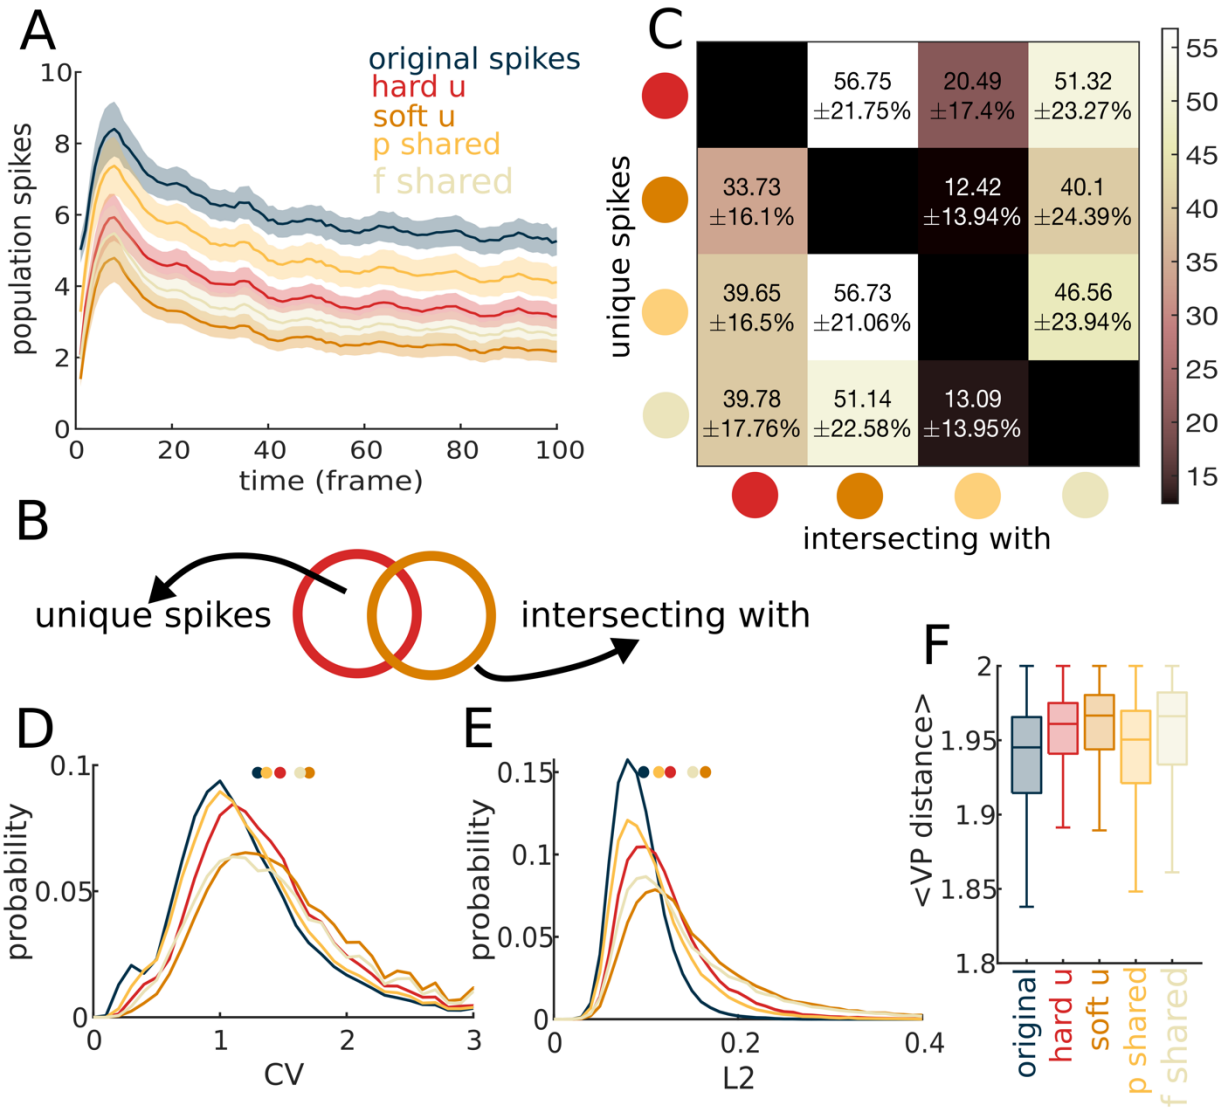

Figure 3.3 - Sub-FN consistent spikes are ultra-sparse and highly variable

A Population spikes over time in the trial for the original spikes and the four sets of sub-FN consistent spikes. Lines and shading are the means and standard deviations across all trials.

B Illustration of how spike overlap is calculated: in each trial, we intersected each of the four sets of sub-FN consistent spikes with the three sets of sub-FN consistent spikes. In this example, we are intersecting hard u (in red) with soft u (in orange). Some of the spikes appear only in hard u (arrow from within the circle) and these are the unique spikes.

C Spikes overlap between sets of four sub-FN consistent spikes. Colors in the grid represent the percent of unique spikes out of all the spikes for the sub-FN in a given row. Numbers are means±standard deviations across trials and datasets.

D Coefficient of variation probability distributions for the firing rate of single neurons in the original spikes and in each of the sets of sub-FN consistent spikes. Dots on the top indicate the means.

Figure 3.3 - continued

E L2 (Euclidean) norm (see Methods) probability distributions for population vectors in the original spikes and in each of the sets of sub-FN consistent spikes. Dots on the top indicate the means.

F Temporal precision as measured by the VP distance. We normalized the metric by the number of spikes (Methods). <> denotes the mean over pairs of trials of the same grating direction. Boxes span the interquartile range and horizontal lines indicate the medians.

Soft-u consistent spikes are more informative of drifting grading direction

Despite variable responses it is clear that information relevant to visual stimulus is present in the response of the visual cortex. In order to directly quantify the information we calculated the mutual information between spikes and stimulus[49] (Methods). We did so for all spikes as well as for the subset of spikes that corresponded to edges in each of the sub-FN. We normalized the raw information by the average number of spikes in the time frames used for analysis to obtain a metric of bits per spike regardless of sparsity. We examined the information in a small group of neurons: we picked 5 neurons that had the largest average firing rates across stimuli and trials in the original spikes, and we kept the identities and order of these 5 neurons fixed. Examining the activity in bins of 10 frames, we found that soft u consistent spikes were more informative about the direction of drifting gratings, as compared to all of the spikes and the other subsets of spikes (Fig. 4A). To quantify the information on faster timescales, we concatenated every 3 subsequent imaging frames preserving the spatio-temporal pattern of spikes. A similar trend emerged, where the information quantity in these spatio-temporal patterns of soft u consistent spikes was increased as compared to all of the action potentials and to other subset of sub-FNs consistent spikes (Fig. 4B). We replicated these results with a separate set of 5 neurons for the original spikes and each of the four subsets of spikes.

[49] Brenner et al., "Synergy in a Neural Code."

Figure 3.4 - Information quantity in subsets of sub-FN consistent spikes

A Bits per spike in sub-FN consistent spikes (different sub-FNs in their respective colors) and original spikes across time. Here we quantified the information in "binary words" of size 5, using the 5 neurons with the largest firing rate and binning over 10 frames. Line and shading represent the mean and standard error, respectively.

B Bits per spike in sub-FN consistent spikes (different sub-FNs in their respective colors) and original spikes across time. Here we quantified the information in spatio-temporal patterns where every 3 consecutive frames were vectorized, which preserves the structure of spiking in the 5 neurons we analyzed.

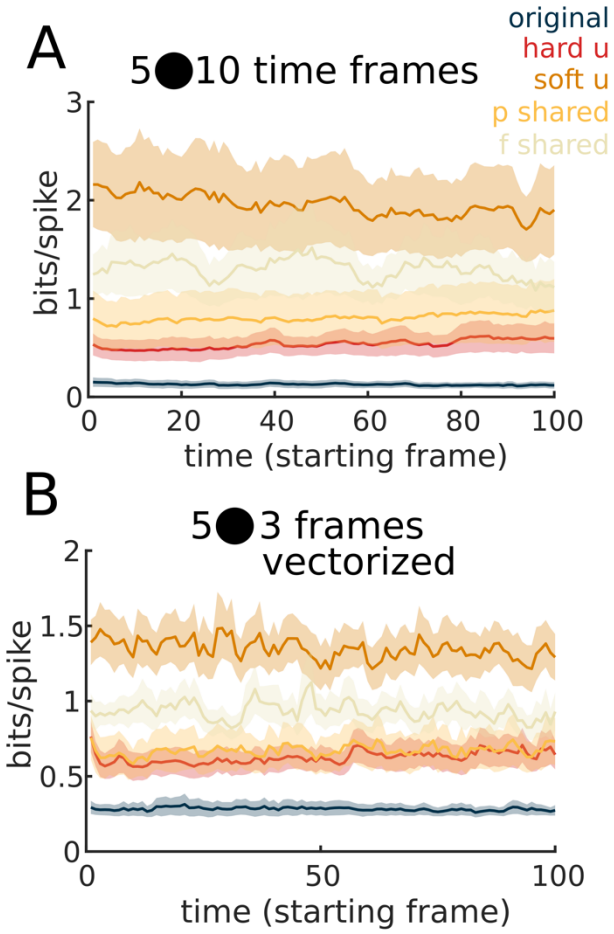

Soft u consistent spikes are decodable at high accuracy

The elevated information quantity in soft u spikes suggests that soft u spikes might be more decodable. Indeed, a feedforward network multiclass decoder trained with conjugate gradient on 90% of the data yielded significantly higher performance for soft u consistent spikes (52.13±14.60%) as compared to all action potentials and all of the other subsets of spikes (34.12±11.8% for the original spikes. 44.37±12.80%, 40.64±11.20% and 42.95±13.06% for hard u, p shared and f shared, respectively, Fig. 5A). We note that decoding performance followed a similar temporal profile to the proportion of edges which corresponded to pairs of action potentials (Fig. 2B).

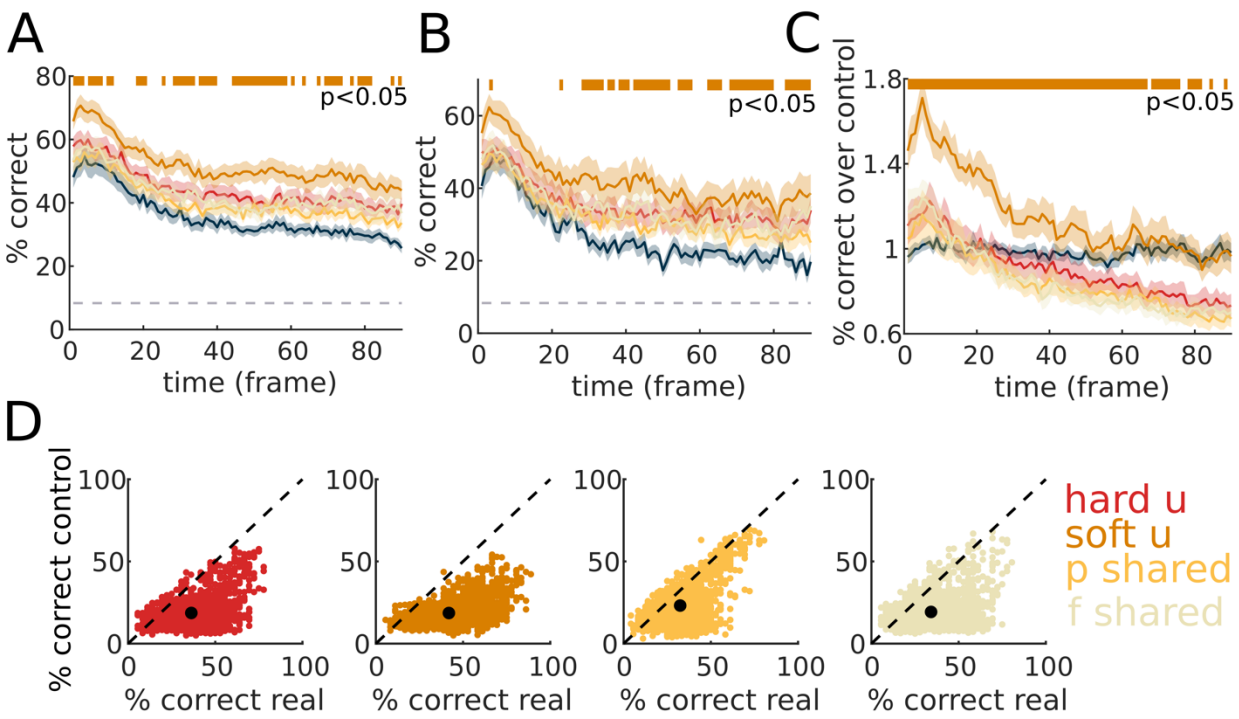

Figure 3.5 - Decoding the direction of drifting gratings from sub-FN consistent spikes

A Performance as percent correct from the test set trials of the original spikes and the sparsified sub-FN consistent spikes in the Feedforward neural network decoder (Methods). Dashed line on the bottom denotes chance level decoding (1/12). The original spikes are in blue. Orange bars on top indicate the time points at which soft u is significantly better (ANOVA, $p<0.05$). Lines and shading are for the means and standard errors across datasets.

B Performance of a second supervised decoder, namely a Support vector machine (SVM). Colors and markings are the same as in A.

C Performance of a Naive Bayes decoder (unsupervised) over sparsity-matched controls (Methods) over time in the trial. Colors and markings are the same as in A.

D Real decoding performance (abscissa) of the four sets of sub-FN consistent spikes against the decoding performance of control spikes (ordinate) in the SVM framework. Control spikes were generated by permuting the sub-FN before intersection with the original spikes and thus preserve the sparsity and the same count of pairs of consecutive spikes but not their specificity (see Methods). Dashed line is unity. Each data point is one decoder, binned at 10 frames. Black dots represent cloud means.

We next confirmed these findings with two additional decoding approaches: We employed a support vector machine (SVM) with a linear kernel which is another supervised decoder (Fig.

5B), and also a naïve Bayes decoder which is an unsupervised method (Fig. 5C). In all three cases the decoders have demonstrated that soft u consistent spikes are more readable downstream. Importantly, in order to verify that this result is due to the existence of specific pairwise statistical dependencies between neurons, we used a control that disentangles the sub-FNs and the spikes. Specifically, we permuted the edges of each sub-FN before selecting the subset of action potentials. This allowed us to preserve density and allowed for other, random pairwise dependencies to be expressed. In all cases decoding performance of this control was less accurate and was most pronounced for the unique types: control performance was 0.57±0.31 and 0.50±0.23 of the real decoding performance of hard u and soft u, respectively (0.76±0.33 and 0.66±0.37 for p shared and f shared, respectively).

DISCUSSION

Here we extended the standard graph-theory approach, in which the activity that occurs over a number of trials is summarized as a functional network, by examining the dynamical correspondence between edges in the FN and action potentials over the time course of single trials. Furthermore, we considered the similarities between functional networks that corresponded to activity evoked by the different directions of drifting gratings and allocated edges into four classes based on their exclusivity to one direction or another. We found that an ultra-sparse set of action potentials that corresponds to edges that are strong FNs of one direction and weak in FNs of adjacent directions (soft u) contained information about the direction of drifting gratings. These subsets of spikes are more readily decoded, than any other set or all of the action potentials evoked by the visual stimulus.

Our findings suggest that some spatio-temporal patterns are only found in one direction. Identification of these patterns directly from the raw spike trains is difficult since they are

interwoven with patterns that are stimulus-nonspecific with a neuron and a spike being in more than one pattern simultaneously. Even a detection method based on the similarity of correlation matrices[50], though conceptually similar to our method, will fall short due to this mingled nature of the spiking activity. Rather, we leverage the stimulus-specificity of some pairwise functional connections[51,52,53] to find these patterns. It follows that, to the extent that neuron pairs with stimulus-specific correlation project to the one downstream target whereas neuron pairs with non-stimulus-specific correlation project to another, these patterns can be disentangled even with overlapping neuron identities and spikes. We have previously found that higher-order relationships are hallmarks of synaptic cooperativity and coordinated integration in a post-synaptic neuron[54], making the proposed readout realistic.

This coding scheme is broadly consistent with the Hebbian assembly hypothesis[55,56,57]. However, in most studies the assembly is assumed to be composed of groups of neurons without considerable overlap in cell identities. We revise the Hebbian assembly hypothesis by demonstrating that pairwise correlations within the group of neurons defines the assembly; soft u sub-FNs encompass the vast majority of neurons recorded indicating that assemblies coding for different directions gratings highly overlap, but differ in their correlation structure. Intuitively, the inclusion of correlations in the implementation of the code increases the coding capacity of a population due to an increased number of possible combinations. Consequently, specific sets of

[50] Grossberger, Battaglia, and Vinck, "Unsupervised Clustering of Temporal Patterns in High-Dimensional Neuronal Ensembles Using a Novel Dissimilarity Measure."
[51] Ponce-Alvarez et al., "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons."
[52] Franke et al., "Structures of Neural Correlation and How They Favor Coding."
[53] Levy, Sporns, and MacLean, "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding."
[54] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."
[55] Miller et al., "Visual Stimuli Recruit Intrinsically Generated Cortical Ensembles."
[56] Harris, "Neural Signatures of Cell Assembly Organization."
[57] Hebb, *The Organization of Behavior*.

spikes code for a variable, not specific neurons. This coding framework is reminiscent of packet coding[58] in which no limitation on neuronal specificity for a variable is set.

Interestingly, soft u consistent spikes (and other sub-FN consistent spikes) did not exhibit high fidelity spatio-temporal sequences, but nonetheless contained more information about the stimulus. This suggests that several spatio-temporal patterns, different in their structure and identity of spiking neurons, can code for the same stimulus. This is in agreement with studies in retina that have found semantic, rather than structural similarity between spikes that code for the same stimulus[59,60]. One implication of this observation is that neuronal and population variability per stimulus does not capture the variance meaningful to the neural code[61]. The reliability that matters to this code is variance of pairwise spikes between neurons[62,63], supporting robustness.

In addition to sparseness, the coding scheme proposed here is particularly efficient since one spike can participate in coding multiple, independent variables. In that sense the code is multiplexed. The high complexity of the natural world and the fact that sensory and motor related brain activity is broadly distributed[64,65] demonstrate the value of a multiplexed code. Multiplexing can be considered for the format of the coding scheme[66,67], the content of the code (i.e. the feature),

[58] Luczak, McNaughton, and Harris, "Packet-Based Communication in the Cortex."
[59] Berry and Tkačik, "Clustering of Neural Activity."
[60] Ganmor, Segev, and Schneidman, "A Thesaurus for a Neural Population Code."
[61] Montijn et al., "Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective."
[62] Sadovsky and MacLean, "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity."
[63] Desbordes et al., "Timing Precision in Population Coding of Natural Scenes in the Early Visual System."
[64] Steinmetz et al., "Distributed Coding of Choice, Action and Engagement across the Mouse Brain."
[65] Musall et al., "Single-Trial Neural Dynamics Are Dominated by Richly Varied Movements."
[66] Kayser et al., "Spike-Phase Coding Boosts and Stabilizes Information Carried by Spatial and Temporal Spike Patterns."
[67] Ainsworth et al., "Rates and Rhythms."

or both[68,69]. While in our work no variables other than the direction of drifting gratings were systematically examined, we postulate that stimulus-nonspecific (shared) correlations might code for other variables, such as internal state. Other researchers have suggested that demultiplexing can be carried out through connectivity to downstream targets[70]. Such specific connectivity patterns still need to be demonstrated, but it is worth noting that cortical pyramidal project to multiple targets[71].

In summary, here we use functional networks to isolate pairs of spikes that occur during a single trial in a stimulus dependent manner. We find that only a small subset of these pairwise statistical dependencies and the corresponding pairs of spikes occur at any given single point in time, resulting in very sparse and variable dynamics. Nonetheless, these sparse sets of spikes carry more information about the stimulus than all of the spikes recorded during the same single trial. Our coding scheme elaborates upon and unifies several theories of neural coding including Hebbian assemblies, semantic coding and packet coding while simultaneously being consistent with the necessity of synaptic integration for a code to be mechanistically viable.

[68] Lankarany et al., "Differentially Synchronized Spiking Enables Multiplexed Neural Coding."
[69] Insanally et al., "Spike-Timing-Dependent Ensemble Encoding by Non-Classically Responsive Cortical Neurons."
[70] Naud and Sprekeler, "Sparse Bursts Optimize Information Transmission in a Multiplexed Neural Code."
[71] Kasthuri et al., "Saturated Reconstruction of a Volume of Neocortex."

REFERENCES

Ainsworth, Matt, Shane Lee, Mark O. Cunningham, Roger D. Traub, Nancy J. Kopell, and Miles A. Whittington. "Rates and Rhythms: A Synergistic View of Frequency and Temporal Coding in Neuronal Networks." *Neuron* 75, no. 4 (August 23, 2012): 572–83. https://doi.org/10.1016/j.neuron.2012.08.004.

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. https://doi.org/10.1038/nrn1888.

Averbeck, Bruno B., and Daeyeol Lee. "Effects of Noise Correlations on Information Encoding and Decoding." *Journal of Neurophysiology* 95, no. 6 (June 1, 2006): 3633–44. https://doi.org/10.1152/jn.00919.2005.

Berry, Michael J. II, and Gašper Tkačik. "Clustering of Neural Activity: A Design Principle for Population Codes." *Frontiers in Computational Neuroscience* 14 (2020). https://doi.org/10.3389/fncom.2020.00020.

Bojanek, Kyle, Yuqing Zhu, and Jason MacLean. "Cyclic Transitions between Higher Order Motifs Underlie Sustained Asynchronous Spiking in Sparse Recurrent Networks." *PLOS Computational Biology* 16, no. 9 (September 30, 2020): e1007409. https://doi.org/10.1371/journal.pcbi.1007409.

Brenner, Naama, Steven P. Strong, Roland Koberle, William Bialek, and Rob R. de Ruyter van Steveninck. "Synergy in a Neural Code." *Neural Computation* 12, no. 7 (July 1, 2000): 1531–52. https://doi.org/10.1162/089976600300015259.

Chambers, Brendan, Maayan Levy, Joseph B. Dechery, and Jason N MacLean. "Ensemble Stacking Mitigates Biases in Inference of Synaptic Connectivity." *Network Neuroscience*, November 2, 2017, 1–49. https://doi.org/10.1162/NETN_a_00032.

Chambers, Brendan, and Jason N. MacLean. "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks." Edited by Jochen Triesch. *PLOS Computational Biology* 12, no. 8 (August 19, 2016): e1005078. https://doi.org/10.1371/journal.pcbi.1005078.

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

Cohen, Marlene R., and John H. R. Maunsell. "Attention Improves Performance Primarily by Reducing Interneuronal Correlations." *Nature Neuroscience* 12, no. 12 (December 2009): 1594–1600. https://doi.org/10.1038/nn.2439.

Dann, Benjamin, Jonathan A. Michaels, Stefan Schaffelhofer, and Hansjörg Scherberger. "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates." *ELife* 5 (August 15, 2016). https://doi.org/10.7554/eLife.15719.

Dechery, Joseph B., and Jason N. MacLean. "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons." Edited by Jeff Beck. *PLOS Computational Biology* 14, no. 5 (May 4, 2018): e1006153. https://doi.org/10.1371/journal.pcbi.1006153.

Desbordes, Gaëlle, Jianzhong Jin, Chong Weng, Nicholas A. Lesica, Garrett B. Stanley, and Jose-Manuel Alonso. "Timing Precision in Population Coding of Natural Scenes in the Early Visual System." *PLOS Biology* 6, no. 12 (December 16, 2008): e324. https://doi.org/10.1371/journal.pbio.0060324.

Deweese, Michael R., and Anthony M. Zador. "Shared and Private Variability in the Auditory Cortex." *Journal of Neurophysiology* 92, no. 3 (September 1, 2004): 1840–55. https://doi.org/10.1152/jn.00197.2004.

Downes, Julia H., Mark W. Hammond, Dimitris Xydas, Matthew C. Spencer, Victor M. Becerra, Kevin Warwick, Ben J. Whalley, and Slawomir J. Nasuto. "Emergence of a Small-World Functional Network in Cultured Neurons." *PLOS Computational Biology* 8, no. 5 (May 17, 2012): e1002522. https://doi.org/10.1371/journal.pcbi.1002522.

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. https://doi.org/10.1016/j.neuron.2014.02.006.

Franke, Felix, Michele Fiscella, Maksim Sevelev, Botond Roska, Andreas Hierlemann, and Rava Azeredo da Silveira. "Structures of Neural Correlation and How They Favor Coding." *Neuron* 89, no. 2 (January 20, 2016): 409–22. https://doi.org/10.1016/j.neuron.2015.12.037.

Franke, Katrin, Philipp Berens, Timm Schubert, Matthias Bethge, Thomas Euler, and Tom Baden. "Inhibition Decorrelates Visual Feature Representations in the Inner Retina." *Nature* 542, no. 7642 (February 8, 2017): 439–44. https://doi.org/10.1038/nature21394.

Friedrich, Johannes, Pengcheng Zhou, and Liam Paninski. "Fast Online Deconvolution of Calcium Imaging Data." *PLoS Computational Biology* 13, no. 3 (2017): e1005423.

Ganmor, E., R. Segev, and E. Schneidman. "Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code." *Proceedings of the National*

Academy of Sciences 108, no. 23 (June 7, 2011): 9679–84. https://doi.org/10.1073/pnas.1019641108.

Ganmor, Elad, Ronen Segev, and Elad Schneidman. "A Thesaurus for a Neural Population Code." *Elife* 4 (2015): e06134.

Glaser, Joshua I., Ari S. Benjamin, Raeed H. Chowdhury, Matthew G. Perich, Lee E. Miller, and Konrad P. Kording. "Machine Learning for Neural Decoding." *Eneuro* 7, no. 4 (July 2020): ENEURO.0506-19.2020. https://doi.org/10.1523/ENEURO.0506-19.2020.

Goris, Robbe L.T., Eero P. Simoncelli, and J. Anthony Movshon. "Origin and Function of Tuning Diversity in Macaque Visual Cortex." *Neuron* 88, no. 4 (November 2015): 819–31. https://doi.org/10.1016/j.neuron.2015.10.009.

Granot-Atedgi, Einat, Gašper Tkačik, Ronen Segev, and Elad Schneidman. "Stimulus-Dependent Maximum Entropy Models of Neural Population Codes." *PLOS Computational Biology* 9, no. 3 (March 14, 2013): e1002922. https://doi.org/10.1371/journal.pcbi.1002922.

Grossberger, Lukas, Francesco P. Battaglia, and Martin Vinck. "Unsupervised Clustering of Temporal Patterns in High-Dimensional Neuronal Ensembles Using a Novel Dissimilarity Measure." *PLOS Computational Biology* 14, no. 7 (July 6, 2018): e1006283. https://doi.org/10.1371/journal.pcbi.1006283.

Harris, Kenneth D. "Neural Signatures of Cell Assembly Organization." *Nature Reviews Neuroscience* 6, no. 5 (May 2005): 399–407. https://doi.org/10.1038/nrn1669.

Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory*. Psychology Press, 1949.

Insanally, Michele N, Ioana Carcea, Rachel E Field, Chris C Rodgers, Brian DePasquale, Kanaka Rajan, Michael R DeWeese, Badr F Albanna, and Robert C Froemke. "Spike-Timing-Dependent Ensemble Encoding by Non-Classically Responsive Cortical Neurons," n.d., 31.

Kanitscheider, Ingmar, Ruben Coen-Cagli, and Alexandre Pouget. "Origin of Information-Limiting Noise Correlations." *Proceedings of the National Academy of Sciences* 112, no. 50 (December 15, 2015): E6973–82.

Kasthuri, Narayanan, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongil Lee, et al. "Saturated Reconstruction of a Volume of Neocortex." *Cell* 162, no. 3 (July 2015): 648–61. https://doi.org/10.1016/j.cell.2015.06.054.

Kayser, Christoph, Marcelo A. Montemurro, Nikos K. Logothetis, and Stefano Panzeri. "Spike-Phase Coding Boosts and Stabilizes Information Carried by Spatial and Temporal Spike Patterns." *Neuron* 61, no. 4 (February 2009): 597–608. https://doi.org/10.1016/j.neuron.2009.01.008.

Kotekal, Subhodh, and Jason N. MacLean. "Recurrent Interactions Can Explain the Variance in Single Trial Responses." *PLOS Computational Biology* 16, no. 1 (January 30, 2020): e1007591. https://doi.org/10.1371/journal.pcbi.1007591.

Lankarany, Milad, Dhekra Al-Basha, Stéphanie Ratté, and Steven A. Prescott. "Differentially Synchronized Spiking Enables Multiplexed Neural Coding." *Proceedings of the National Academy of Sciences* 116, no. 20 (May 14, 2019): 10097–102.

Levy, Maayan, Olaf Sporns, and Jason N. MacLean. "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding." *Cell Reports* 31, no. 2 (April 14, 2020): 107483. https://doi.org/10.1016/j.celrep.2020.03.047.

Luczak, Artur, Bruce L. McNaughton, and Kenneth D. Harris. "Packet-Based Communication in the Cortex." *Nature Reviews Neuroscience* 16 (2015): 745–55.

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 (September 23, 2014): E4053–61.

Montani, Fernando, and Simon R. Schultz. "Information-Theoretic Analysis of the Role of Correlations in Neural Spike Trains." In *Concepts and Recent Advances in Generalized Information Measures and Statistics*, 375–407. Accessed July 5, 2021. https://www.eurekaselect.com/118914/chapter.

Montijn, Jorrit S., Guido T. Meijer, Carien S. Lansink, and Cyriel M.A. Pennartz. "Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective." *Cell Reports* 16, no. 9 (August 2016): 2486–98. https://doi.org/10.1016/j.celrep.2016.07.065.

Musall, Simon, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. "Single-Trial Neural Dynamics Are Dominated by Richly Varied Movements." *Nature Neuroscience* 22, no. 10 (October 2019): 1677–86. https://doi.org/10.1038/s41593-019-0502-4.

Naud, Richard, and Henning Sprekeler. "Sparse Bursts Optimize Information Transmission in a Multiplexed Neural Code." *COMPUTATIONAL BIOLOGY*, *Proceedings of the National Academy of Science*, 10.

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. https://doi.org/10.1038/nature09178.

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. https://doi.org/10.1073/pnas.1506855112.

Pillow, Jonathan W., Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M. Litke, E. J. Chichilnisky, and Eero P. Simoncelli. "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population." *Nature* 454, no. 7207 (August 21, 2008): 995–99. https://doi.org/10.1038/nature07140.

Ponce-Alvarez, A., A. Thiele, T. D. Albright, G. R. Stoner, and G. Deco. "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons." *Proceedings of the National Academy of Sciences* 110, no. 32 (August 6, 2013): 13162–67. https://doi.org/10.1073/pnas.1300098110.

Sadovsky, A. J., and J. N. MacLean. "Mouse Visual Neocortex Supports Multiple Stereotyped Patterns of Microcircuit Activity." *Journal of Neuroscience* 34, no. 23 (June 4, 2014): 7769–77. https://doi.org/10.1523/JNEUROSCI.0169-14.2014.

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. https://doi.org/10.1152/jn.00334.2011.

Salinas, E., and T. J. Sejnowski. "Impact of Correlated Synaptic Input on Output Firing Rate and Variability in Simple Neuronal Models." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 20, no. 16 (August 15, 2000): 6193–6209.

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. https://doi.org/10.1038/35086012.

Steinmetz, Nicholas A., Peter Zatka-Haas, Matteo Carandini, and Kenneth D. Harris. "Distributed Coding of Choice, Action and Engagement across the Mouse Brain." *Nature* 576, no. 7786 (December 2019): 266–73. https://doi.org/10.1038/s41586-019-1787-x.

Stevenson, Ian H., Brian M. London, Emily R. Oby, Nicholas A. Sachs, Jacob Reimer, Bernhard Englitz, Stephen V. David, et al. "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons." *PLOS Computational Biology* 8, no. 11 (November 15, 2012): e1002775. https://doi.org/10.1371/journal.pcbi.1002775.

Tolhurst, D. J., J. A. Movshon, and A. F. Dean. "The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex." *Vision Research* 23, no. 8 (1983): 775–85. https://doi.org/10.1016/0042-6989(83)90200-6.

Usrey, W. Martin, John B. Reppas, and R. Clay Reid. "Paired-Spike Interactions and Synaptic Efficacy of Retinal Inputs to the Thalamus." *Nature* 395, no. 6700 (September 1998): 384–87. https://doi.org/10.1038/26487.

Ventura, V. "Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons." *Journal of Neurophysiology* 94, no. 4 (April 20, 2005): 2928–39. https://doi.org/10.1152/jn.00644.2004.

Victor, J. D., and K. P. Purpura. "Nature and Precision of Temporal Coding in Visual Cortex: A Metric-Space Analysis." *Journal of Neurophysiology* 76, no. 2 (August 1, 1996): 1310–26.

Zhang, Kechen, Iris Ginzburg, Bruce L. McNaughton, and Terrence J. Sejnowski. "Interpreting Neuronal Population Activity by Reconstruction: Unified Framework With Application to Hippocampal Place Cells." *Journal of Neurophysiology* 79, no. 2 (February 1, 1998): 1017–44. https://doi.org/10.1152/jn.1998.79.2.1017.

SUPPLEMENTARY FIGURES

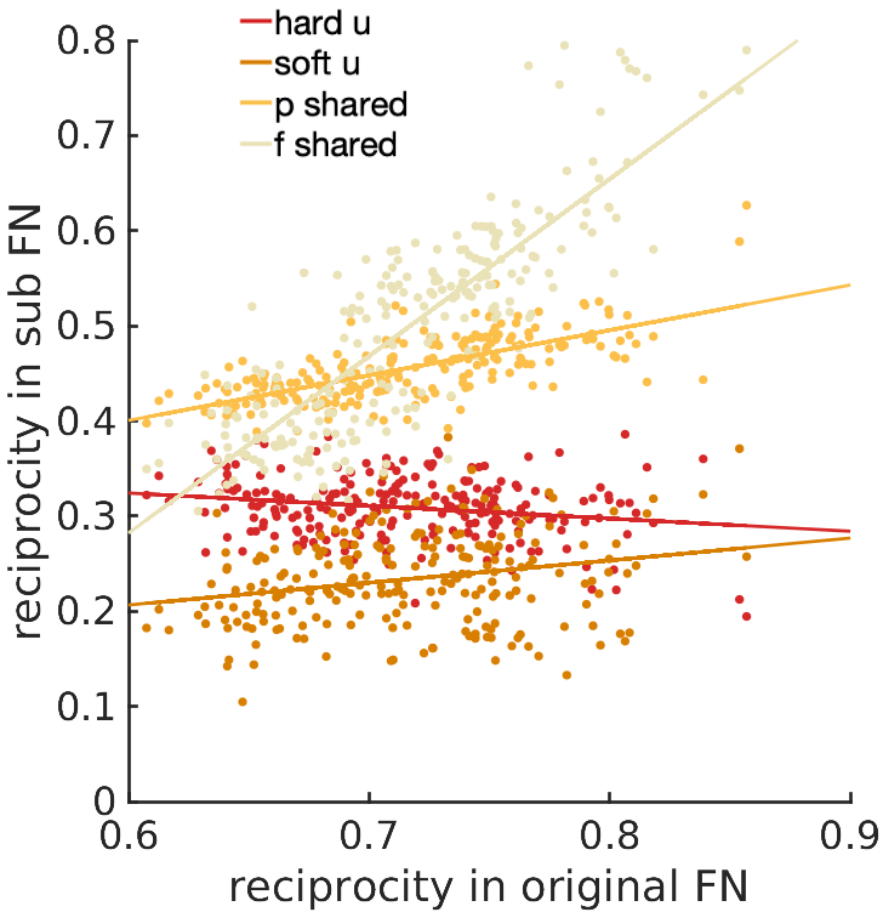

Figure 3.S1 – Unique sub-types display less reciprocal edges

Reciprocity is estimated as the number of bidirectional edges in an FN over the total number of non-zero edge. Line show he best linear fit.

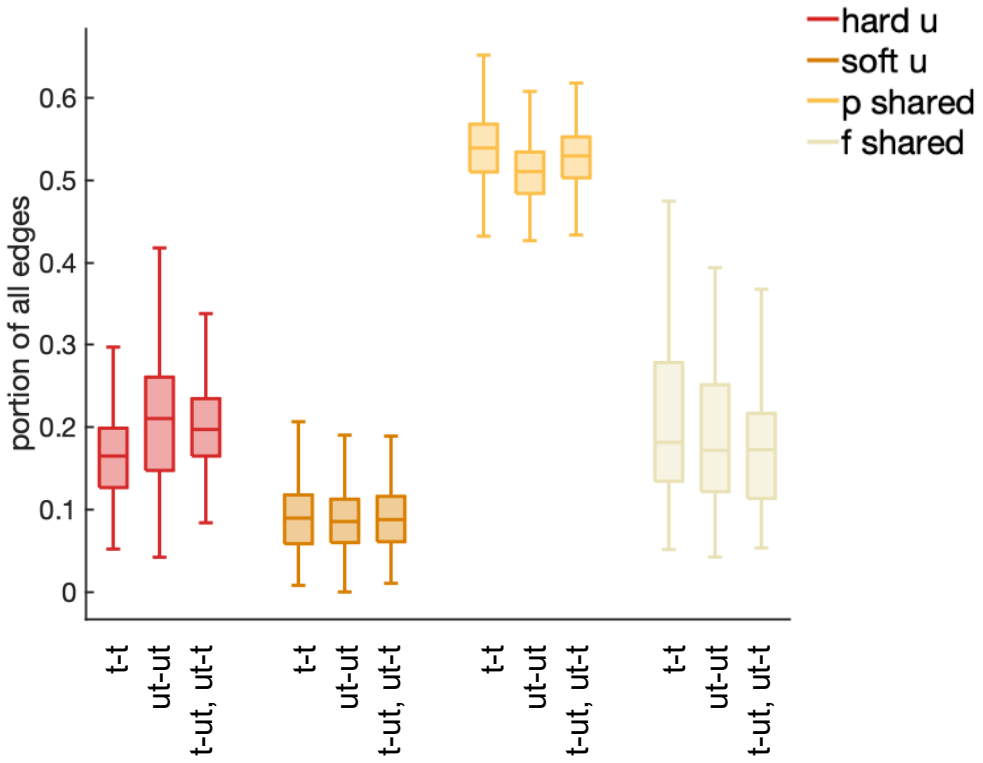

Figure 3.S2 – Unique sub-FNs do not include more tuned-tuned edges

Boxplots shows the portion of tuned-tuned (t-t), untuned-untuned (ut-ut) and mixed (t-ut, ut-t) edges sorted to each of the subtypes. Boxes represent the interquartile range and horizonal lines stand for the median.

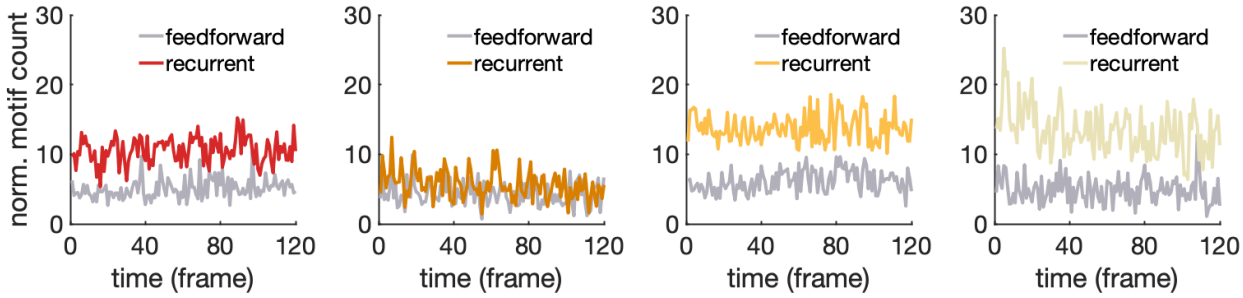

Fig 3.S3 – Recurrent and feedforward pattern counts over time

Pattern counts in example dataset. Patterns were divided by the number of expressed edges at each time point and then summed across all trials and all directions.

DISCUSSION

On the structure-function question

In chapter I the congruence between synaptic and functional networks was examined using multiple timescales and inference methods, each with its own statistical sensitivity (e.g. linear vs. non-linear). It is found that a regularization pipeline designed to account for negative as well as spurious interactions and a stacking algorithm optimized to weigh the different algorithm improve detection of synaptic connectivity. Nevertheless, Chapter I emphasizes the complex integration of activity into correlated dynamics. This complexity is evident in the percent of connections that are present in the recruitment network; neurons require many more than one input at a time to reach threshold, which was known since the early days of Neuroscience[1]. Moreover, neurons with many inputs in simulations do not necessarily have high firing rate[2,3]. That is, ground-truth degree is not correlated with activity. Approximately 50% recruitment across simulation time is consistent with the low instantaneous expression rates found in chapter III for FNs in vivo, considering that FN is not an anatomical network and is denser[4].

All inference methods point to statistical regularities in dynamics, i.e. functional connections, as being more reciprocal and clustered than the anatomical substrate that generates them. This is especially true for synchrony metrics, popular for functional interaction estimations[5,6], but is observed with lagged metrics as well. In this sense the structure of dynamics

[1] McCulloch and Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity."
[2] Vegué and Roxin, "Firing Rate Distributions in Spiking Networks with Heterogeneous Connectivity."
[3] Kuhn, Aertsen, and Rotter, "Higher-Order Statistics of Input Ensembles and the Response of Simple Model Neurons."
[4] Bojanek, Zhu, and MacLean, "Cyclic Transitions between Higher Order Motifs Underlie Sustained Asynchronous Spiking in Sparse Recurrent Networks."
[5] Lindsey et al., "Repeated Patterns of Distributed Synchrony in Neuronal Assemblies."
[6] Montani et al., "The Impact of High-Order Interactions on the Rate of Synchronous Discharge and Information Transmission in Somatosensory Cortex."

is an emergent phenomenon, possibly driven by non-linear integration[7,8]. Thus, to the extent that they display (dis)similar statistics, FNs are poor approximators of synaptic connectivity and one should exercise caution in interpreting them as such. A limitation of chapter I is that the synaptic connectivity structure used was an Erdos-Renyi graph, which is random with fixed probability. Yet it has been known since the days of Cajal that connectivity in the brain is not random. On the contrary, it is biased towards the prevalence of reciprocal connections and triplet motifs[9] and has a small-world architecture[10]. It thus may be that FNs inferred from real data more closely reflect the statistics of their underlying connectivity structure.

One way to approach this question is to design simulations with connectivity inspired by anatomy, and with biologically plausible integration mechanism and membrane time constants. It proved to be challenging to take into account all these factors and emulate the firing rates and irregularity seen in data. A study using similar simulations but with clustered neurons and the same inference method still found more reciprocal connections in the FNs compared to the synaptic networks[11], and another study has shown that the presence of hubs increases pairwise correlations in dynamics[12]. Hence, tracking the subthreshold activity may be required to address the structure-function question of how anatomical topology gives rise to considerably more correlated dynamics than expected. It was shown that correlation of subthreshold activity better reflects the synaptic

[7] Memmesheimer and Timme, "Non-Additive Coupling Enables Propagation of Synchronous Spiking Activity in Purely Random Networks."
[8] Rhodes, "Recoding Patterns of Sensory Input."
[9] Song et al., "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits."
[10] Perin, Berger, and Markram, "A Synaptic Organizing Principle for Cortical Neuronal Groups."
[11] Bojanek, Zhu, and MacLean, "Cyclic Transitions between Higher Order Motifs Underlie Sustained Asynchronous Spiking in Sparse Recurrent Networks."
[12] Pernice et al., "How Structure Determines Correlations in Neuronal Networks."

topology[13]. Indeed, access to subthreshold activity in simulation[14] and experiment[15] elucidated that synchrony between pairs of cells is due to common input and not a direct synaptic connection. Unfortunately, subthreshold activity is typically unavailable in population recordings.

The structure of functional networks themselves as a summary of dynamics and not as a reflection of underlying anatomy is highly relevant to coding. Chapter II found that the preference of a neuron to a stimulus feature on average, known as tuning, is corresponds to its location in the FN. In other words, activity across the population is organized such that cells which correlate with many other neurons in a clustered manner respond fairly evenly to multiple stimulus features. Chapter II thus expands on the results of Stevenson et al[16], where functional correlations were demonstrated to explain the spiking and tuning preferences of neurons in several brain regions in the monkey. In both studies, single cell properties can be viewed as a manifestation of network interactions. However, chapter II samples a complete and unbiased neuronal population, and for the first time characterizes pairwise correlations between neurons that are classically untuned.

On the neural code question

In a groundbreaking study, Dan et al[17] examined pairs of neurons in the cat lateral geniculate nucleus (LGN). They sorted spikes according to whether they were synchronous, that is, whether the two neurons spiked at the same time, or independent, where only one neuron spiked and the other one was silent. They showed that synchronous spikes carry more information about the brightness of pixels presented in comparison to the same number of independent spikes. By

[13] Kriener et al., "Correlations in Spiking Neuronal Networks with Distance Dependent Connections."
[14] Ostojic, Brunel, and Hakim, "How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains."
[15] Trong and Rieke, "Origin of Correlated Activity between Parasol Retinal Ganglion Cells."
[16] Stevenson et al., "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons."
[17] Dan et al., "Coding of Visual Information by Precisely Correlated Spikes in the Lateral Geniculate Nucleus."

filtering spikes Dan el al examined how correlations are instantiated, and were amongst the first to consider the perspective of a downstream element (see also [18]). A similar approach was employed in chapter III with several extensions: First, chapter III examined activity in mouse V1, which is less anatomically organized according to single-cell functional properties than cat LGN[19] and displays stronger pairwise correlations[20]. Nonetheless, the results echo Dan's classic paper. Second, pairs of spikes were inspected for their correspondence with a subset of correlations that are stimulus-specific. The stimulus specificity of correlations occurs regardless of the selectivity of neurons composing the pair, as reported in chapter II, and this was further confirmed in chapter III. Third, in this work dynamics were intersected with functional networks that are beyond pairwise.

Hence, while Dan et al study could offer conclusions on pairs alone, chapter III provides insights into the potential organization of the code across the population. In particular, stimulus-specific spatio-temporal patterns of activity are identified across the population. These patterns are decodable despite their variability at the single cell and population levels, implying that the code might rely on semantic similarity[21,22]. As Hebbian assemblies have received multiple interpretations in the past, this work refers to Harris'[23] four signatures of Hebbian assemblies to evaluate whether the proposed coding scheme fits the definition. The first signature is sequences of activity on temporal scales that are not included in the stimulus. Drifting gratings presented to the animals had a temporal frequency of 2Hz, which is faster than the mean firing rates observed in the population but slower than the timescale of pairs of spikes corresponding to edges.

[18] deCharms and Merzenich, "Primary Cortical Representation of Sounds by the Coordination of Action-Potential Timing."
[19] Kaschube, "Neural Maps versus Salt-and-Pepper Organization in Visual Cortex."
[20] Goris, Movshon, and Simoncelli, "Partitioning Neuronal Variability."
[21] Berry and Tkačik, "Clustering of Neural Activity."
[22] Ganmor, Segev, and Schneidman, "A Thesaurus for a Neural Population Code."
[23] Harris, "Neural Signatures of Cell Assembly Organization."

Furthermore, all stimuli shared this feature, while non-overlapping subsets of correlations were found for different stimuli. Second, Harris argues that trial-to-trial variability and irregular firing arise due to correlation in inputs[24,25]. The population of layer 2/3 cells analyzed in chapter III presumably receives correlated inputs from layer 4 and the LGN. Increased variability in sub-FN consistent spikes is thus not unexpected. As an input, identified correlated spikes are hypothesized to induce irregular firing in their target population. The third signature is organization on the population level and successful prediction of the activity of one neuron from its neighbors' activity. Previous analysis of the same datasets have shown non-random functional topology[26,27] and successful prediction of single-neuron, single-trial dynamics using couplings with other neurons[28,29]. Furthermore, in chapter III spikes are filtered according to four exclusive sub-FNs; these are four population-level collections of pairwise correlations. One sub-FN (soft u), in turn, summarizes spikes that are particularly informative, thus showing that organization of the neural population occurs in both the static and temporal domains. Finally, while chapter III does not demonstrate that some sub-FNs, specifically shared sub-types, are correlated with internal variables as required by the fourth signature, it lays the groundwork for future research on this question, as discussed later.

To summarize the picture of the code painted by chapter III, the author would like to offer an analogy. Musical metaphors have been used in Neuroscience to describe various

[24] Softky and Koch, "The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs."
[25] Stevens and Zador, "Input Synchrony and the Irregular Firing of Cortical Neurons."
[26] Levy, Sporns, and MacLean, "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding."
[27] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[28] Dechery and MacLean.
[29] Kotekal and MacLean, "Recurrent Interactions Can Explain the Variance in Single Trial Responses."

observations[30,31,32]. At the risk of sounding cliché, one can conceptualize population activity as a musical piece; what makes the piece, e.g. Beethoven's Moonlight Sonata, identifiable, are a subset of statistical dependencies between notes building up sequences, such as G followed by A followed by E. These core statistical dependencies that make up the melody are surrounded by other notes that make up the tempo and embellishments. For example, the Moonlight Sonata can still easily be identified even in a different musical style, such as metal or punk cover, and each of these styles has its own set of typical statistical dependencies between notes, played at the same time as the core notes for the Sonata melody.

Similarly, chapter III suggests that a subset of functional connections characterizes the stimulus, whereas others might code for other external variables such as modalities other than the one controlled by the experimenter or internal variables such as attention. In that regard, the neural code is not made up of the rate of single neurons, akin to the prevalence of specific notes in a piece, nor of the rate of many neurons as a population vector, analogous to the counts of many notes over the time course of the piece. Both of these schemes will potentially produce similar rates for multiple stimuli or music pieces, and will be variable across repetitions depending on uncontrolled variables (or the style in the musical analogy). The neural code is also not made up of precise timing of independent units. While some information about the melody can be extracted from the timing of single notes, they are only meaningful in the context of the notes they precede or follow. Even in a single dataset, which is a spatial sneak-peak into the brain, there are one to three orders of magnitude more neurons than musical notes, let alone in a whole region or the entire cortex, (9 to 11 orders of magnitudes more, depending on the organism). Revisiting this complexity that was

[30] Grinvald et al., "Neuronal Assemblies."
[31] Ikegaya et al., "Synfire Chains and Cortical Songs."
[32] Okun et al., "Diverse Coupling of Neurons to Populations in Sensory Cortex."

also discussed at the opening of this work, it stands to reason that a scalable coding scheme is highly desirable. Can the Moonlight Sonata scheme be scaled up?

Back to the philosophy of the code – a homunculus-inspired thought experiment

Most studies discussing the neural code assume an observer, decoder or interpreter, all of which refer to a readout entity. The work presented in chapter III is no exception. The extent to which the observer is ideal, optimal, has knowledge of internal states, and is (non)linear or probabilistic can be examined to determine how self-consistent and biologically feasible is the proposed coding scheme, as well as the efficiency and learnability of it, which are crucial for management of resources in the brain. These small thought-experiments set bounds on the code, and sometimes question its very existence.

Here the author would like to focus on the identity of the readout: Chapter III puts forward a coding scheme where almost all cells in the population are part of the assembly and spikes code for more than one variable, the value of which depends on the spike being part of a spiking pattern with at least one other spike from one other neuron. For ease of reference later, this scheme will be uncreatively named the Hebbian assembly plus scheme. What form of readout is required for such a coding scheme? Let i,j,k be neurons, and x_1 and x_2 two values of a feature (for example, motion at 30 and 60 degrees, but note that a similar argument applies for two variables in a multiplexed code, such as direction and contrast). Intuitively, a pair of neurons i,j that shows coincident spiking that codes for x_1 should project to target m, whereas a pair of neurons i,k that spike coincidently to code for x_2 should project to target n. Hence, this coding scheme depends on specific wiring patterns and the identity of the readouts.

What are m and n and what would their activity look like? It's unlikely that they are single neurons, as that goes against the observation that many neurons respond to the same feature value, as well as against the assumption of this work (and others) that neurons code in groups. m and n then must be populations, active only when the specific combinations of spikes are presented to them. This coding scheme then decorrelates these two downstream populations, which is ideal but might not be realistic. Additionally, if one speculates that just one coding scheme exists, namely the Hebbian assembly plus scheme, the code in the downstream population has the same format – that is, all neurons are part of the assembly, which necessitates that populations n and m overlap. Now the statistics of the readout layer resemble those of the original population, and a readout is required for the readout, ad infinitum. The homunculus strikes again.

One possible solution, the so-called anti-representationalist view, is that no readout is needed since there is no code. Activity propagates through all neurons, integrates new information and computes until it reaches neurons that project directly to the spinal cord and behavior is executed. Nevertheless, evidence points to different activity patterns for different stimuli, and spinal cord projections have to be somewhat specific otherwise all muscles will be recruited. Referring to these activity patterns and motor neurons as code and readouts, respectively, is thus purely a semantic debate (but see [33] for a thorough discussion of the coding metaphor).

A second way to resolve the fallacy is to deem the Hebbian assembly plus scheme incorrect. While it is possible, the proposed coding scheme is far from the only scheme to grapple with this difficulty. The Hebbian assembly plus scheme has some merit as discussed above and in chapter III, and can be amended. In particular, the assumption that the coding scheme found in excitatory layer 2/3 neurons of mouse V1 generalizes to other layers, regions and organisms is probably

[33] Brette, "Neural Coding: The Bureaucratic view of the brain"

wrong. Different layers and regions have diverse cellular makeups and wiring diagrams, supposedly to support distinct functions. The hypothesized readout need not then obey the same coding principles found in chapter III. This is not to say that the Hebbian assembly plus scheme is not scalable in the number of neurons within a population. The Hebbian assemble plus coding scheme is a puzzle piece. It is **a** viable coding scheme, rather than **the** neural code. The author hopes to remember this lesson for the rest of her career.

FUTURE DIRECTIONS

Higher-order interactions in spontaneous vs. stimulus-evoked activity

The brain is never silent. Neuronal activity is even initiated and propagated in slice and cell cultures, in the absence of any external stimulus. Numerous studies have asked whether intrinsically generated activity is similar to stimulus evoked dynamics. Broadly, comparable statistics, including spatio-temporal patterns, have been demonstrated in spontaneous and evoked dynamics[34,35,36,37]. It has further been suggested that dynamics that arise in spontaneous activity define the space that evoked activity can potentially visit[38]. Despite these similarities, Stringer et al[39] have showed that spontaneous (called 'ongoing' in their paper) and evoked activity are only similar along one dimension, or mode, of variability. Conversely, Jermakowicz et al[40], reported similar spike-train correlations magnitudes in V1 under stimuli and in the dark, but anesthesia as

[34] Afrashteh et al., "Spatiotemporal Structure of Sensory-Evoked and Spontaneous Activity Revealed by Mesoscale Imaging in Anesthetized and Awake Mice."
[35] Sarracino et al., "Predicting Brain Evoked Response to External Stimuli from Temporal Correlations of Spontaneous Activity."
[36] Miller et al., "Visual Stimuli Recruit Intrinsically Generated Cortical Ensembles."
[37] MacLean et al., "Internal Dynamics Determine the Cortical Response to Thalamic Stimulation."
[38] Luczak, Barthó, and Harris, "Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations."
[39] Stringer et al., "Spontaneous Behaviors Drive Multidimensional, Brainwide Activity."
[40] Jermakowicz et al., "Relationship Between Spontaneous and Evoked Spike-Time Correlations in Primate Visual Cortex."

well as strict neuron selection criteria might have impacted their results. The structure of pairwise correlations as a functional network summarizing dynamics have not been examined.

Several pieces of evidence that intrinsic activity might be characterized by more recurrent and higher-order interactions as compared to stimulus-evoked activity have accumulated while working on chapters II and III. First, functional networks inferred from OFF (gray) epochs have qualitatively different structure: while they are as sparse as stimulus FNs, they display more reciprocal connections than their stimulus-inferred counterparts, regardless of the inference algorithm used (note that reciprocal edges correspond to synchrony in the partial pairwise algorithm but not with confluent mutual information). They also have higher prevalence of three-nodes motifs (for motif descriptions see [41,42]), in particular increased isomorphic middle-man clustering as compared to stimulus FNs.

Second, shared sub-types in chapter III, which do not correspond to stimulus-specific correlations, also contain more reciprocal connections. The spikes corresponding to edges of these sub-types show more recurrent spatio-temporal patterns. Finally, f shared, the least specific sub-type of correlations, is enriched in both cycle and isomorphic motifs (Fig. 1), while unique sub-types show pretty low counts compared to density matched and degree preserving controls. Notably, functional higher-order interactions were described in alert monkey and evoked visual responses in anesthetized cat[43]. These findings are difficult to interpret since a state of alertness can still be thought of as an internal rather than external variable, and the statistics were not directly compared to true spontaneous activity.

[41] Chambers and MacLean, "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks."
[42] Dechery and MacLean, "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons."
[43] Yu et al., "Higher-Order Interactions Characterized in Cortical Activity."

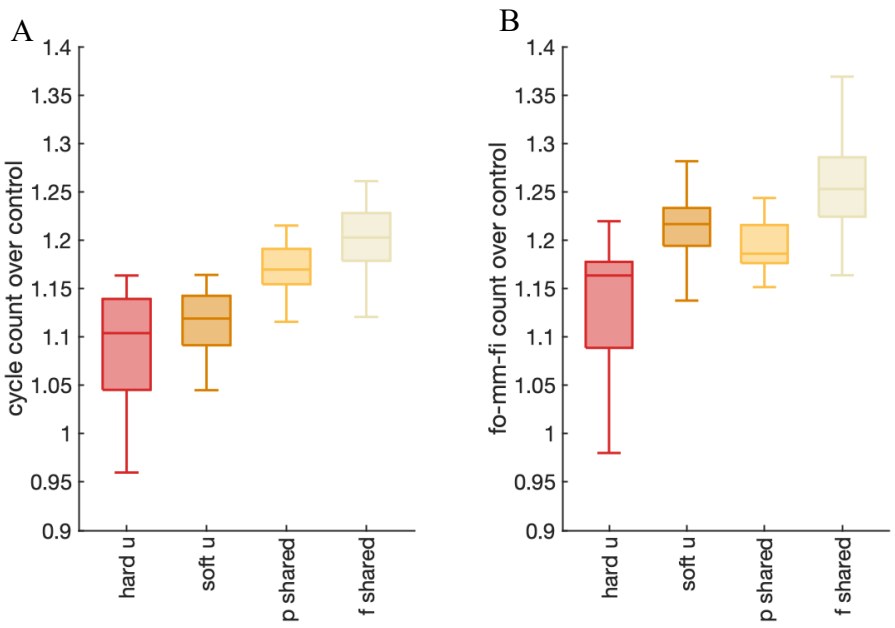

Figure 4.1 – F shared sub-FNs are enriched in triplet motifs

A. cycle motif counts over density and degree-sequence matched controls in the four sub-FNs from chapter III. Boxes and horizonal lines represent the interquartile range and median, respectively.

B. isomorphic motif counts over density and degree-sequence matched controls in the four sub-FNs. Boxes and lines the same as in A.

Similarly and unfortunately, there is no way to tell if activity during gray epochs or f shared consistent spiking is internally generated, and plenty of factors in the mouse environment can still drive these dynamics. Modelling spontaneous activity is challenging since the simulation needs to receive inputs to start it, and hence assumptions need to be made about how spontaneous activity is initiated. For example, an encoding model combined with a spiking neural network can be designed, with two initiation mechanisms: units activate according to a stimulus in their receptive field but in the absence of it and after a period of network quiescence, a subset of units will fire

with some probability. Imaging neuronal populations across lamina[44] in the dark and while engaging with a stimulus is thus warranted to evaluate this hypothesis.

Investigating multiplexing

In chapter III pairwise correlations in a functional network were allocated to four sub-FNs according to their stimulus specificity. Two of these sub-types (p shared and f shared) are composed of statistical regularities that occur regardless of the stimulus identity. For a long time, any shared variability that cannot be readily explained by factors within experimental control was deemed noise. However, neurons are biophysically capable of producing temporally precise spikes[45]. Additionally, shared sub-types showed increased edge expression, i.e. more spikes corresponded to these correlations, during presentation of a mean gray luminance screen, as compared with the unique sub-types. Subsequently, by examining the real-time expression of these sub-types alone one could tell when a drifting grating stimulus or a gray screen was presented to the mouse. Taken together, it appears that spikes that co-vary independently of the stimulus might code for something else, multiplexing representations. While chapter III strongly hints at this direction, it does not explicitly investigate it.

Multiplexing has been studied surprisingly little, and mostly within the same modality or task, where the experimenter varies a combination of features such as direction and contrast. This paradigm typically exposes multiplexing of both representation and format, for example, feature X is coded by rate while feature Y is coded by temporal sequences[46,47,48]. Chapter III proposes a

[44] Sakata and Harris, "Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex."
[45] Mainen and Sejnowski, "Reliability of Spike Timing in Neocortical Neurons."
[46] Ainsworth et al., "Rates and Rhythms."
[47] Lankarany et al., "Differentially Synchronized Spiking Enables Multiplexed Neural Coding."
[48] Kayser et al., "Spike-Phase Coding Boosts and Stabilizes Information Carried by Spatial and Temporal Spike Patterns."

different scheme, which is very biologically plausible, that multiple stimulus features and internal variables can be coded as spatio-temporal patterns, and interwoven patterns are separable as a function of the projection targets of correlated cells (Note: for reasons discussed in the subsection about intrinsic vs. evoked activity it is likely that shared sub-types in chapter III do not code for another stimulus feature, but rather for internal variable(s) such as attention or surprise. The argument presented here for multiplexing is more general).

This hypothesis can be directly tested experimentally by: 1) imaging populations of neurons in response to a controlled stimulus with several features. Then summarizing the activity as a functional network, and segmenting it separately to each feature according to the guidelines in chapter III. If one feature's trash (noise) is another feature's treasure (signal), an overlap and increased graph similarity should be observed between one feature soft u and another feature shared sub-types. Identification of sets of sub-FN consistent spikes should still be partially overlapping. The author believes some research groups led by readers of this document already have data appropriate for such analysis. And 2) since the FN paradigm preserves neuron labels, tracers could be injected to pairs of neurons with correlations that were sorted into specific sub-types for different stimulus features, to elucidate (some of) the readouts. A potential difficulty of this approach is that every neuron projects to multiple targets[49], rendering execution and interpretation challenging. Network simulations, where the connectivity is fully accessible, can supplement experiments. Simulations can be modelled after FNs from real data, and optimal readout topology and weights can be learned.

[49] Kasthuri et al., "Saturated Reconstruction of a Volume of Neocortex."

Inhibitory neurons

All three chapters of this work investigated only excitatory neurons; in chapter I, inhibitory neurons with realistic wiring probabilities were used to generate dynamics without runaway excitation, but were not analyzed for their functional connections. In chapters II and III only pyramidal neurons of layer 2/3 expressed the calcium indicator and therefore no inhibitory neurons were imaged and negative correlations were set to 0.

Despite early work by Hartline[50] characterizing an inhibitory circuit, inhibition was considered for a long time to act as a global mechanism, balancing neural networks to prevent over-excitation and induce (a)synchronization[51,52,53]. Exciting advances in recent years have changed that perception; on the anatomical level, an unprecedented number of inhibitory (and excitatory) cell types have been observed in the brain[54]. The wiring of these cell types have been shown to be organized into anatomical motifs in multiple cortical areas, and together with their morphology and cellular properties these motifs underlie well-defined computations and cortical functions[55,56,57] as demonstrated by optogenetic stimulation and suppression[58,59]. On the modeling side, plasticity of inhibitory synapses with realistic and local learning rules (like homeostatic and spike-timing-dependent-plasticity) have been demonstrated to give rise to neurons' receptive field properties, such as tuning[60]. Finally, employing similar methods to those described in this work, it

[50] Hartline, Wagner, and Ratliff, "Inhibition in the Eye of Limulus."
[51] Barral and D Reyes, "Synaptic Scaling Rule Preserves Excitatory–Inhibitory Balance and Salient Neuronal Network Dynamics."
[52] Landau et al., "The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks."
[53] Brunel, "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons."
[54] Markram et al., "Reconstruction and Simulation of Neocortical Microcircuitry."
[55] Wei and Feller, "Organization and Development of Direction-Selective Circuits in the Retina."
[56] Oviedo, "Connectivity Motifs of Inhibitory Neurons in the Mouse Auditory Cortex."
[57] Feldmeyer et al., "Inhibitory Interneurons and Their Circuit Motifs in the Many Layers of the Barrel Cortex."
[58] Cone et al., "Different Inhibitory Interneuron Cell Classes Make Distinct Contributions to Visual Contrast Perception."
[59] Adesnik et al., "A Neural Circuit for Spatial Summation in Visual Cortex."
[60] Agnes, Luppi, and Vogels, "Complementary Inhibitory Weight Profiles Emerge from Plasticity and Allow Flexible Switching of Receptive Fields."

was recently shown that inhibitory neurons form rich-clubs in functional networks in mouse orbitofrontal cortex[61]. Inhibitory neurons are hard to ignore nowadays and an account of how anatomical connectivity is related of functional connectivity is incomplete without addressing inhibition.

Incorporating inhibitory interactions in a functional network is thus a natural extension of this work, yet poses several challenges: experimentally, simultaneous labelling and 2-photon calcium imaging of excitatory and inhibitory neurons is possible either with a fully (triple) transgenic or a combined transgenic and viral approaches[62]. However, generating FNs from a mixed population is not straightforward. Mutual information is positive by definition, but as done in chapters I and III can be multiplied by the sign of the Pearson correlation between the pair.

This was the technique used in preliminary work by the author, with network of adaptive exponential integrate-and-fire (AdeX) neurons that show low-rate irregular activity. After signing lagged mutual information values, Dale's low was enforced: excitatory / inhibitory neurons were only allowed to have outgoing positive / negative functional edges, respectively. Firing rate was not correlated with degree in excitatory neurons, which is the number of edges a neuron has. This is in agreement with FNs from real dynamics in mouse V1. Conversely, in degree (incoming edges) was correlated with firing rate in inhibitory neurons, whereas their out degree (outgoing edges) was inversely correlated with firing rate. Furthermore, when considering only excitatory / inhibitory sub-networks, in a similar vein to chapter II, excitatory, but not inhibitory, neurons are more clustered amongst themselves, indicating more coordinated spiking in the excitatory pool.

[61] Hafizi et al., "Inhibition-Dominated Rich-Club Shapes Dynamics in Cortical Microcircuits in Awake Behaving Mice."
[62] Najafi et al., "Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning."

This preliminary finding is incongruent with Hafizi et al [63], and might be due to the edges omitted to observe Dale's law in this work but not theirs.

Another challenge is that many complex networks do not include negative links, with the exception of predator-pray and gene-regulation networks. Consequently, many analytic tools like graph alignment and Page rank are not designed for negative edges. Exploring FNs with both excitatory and inhibitory neurons hence calls for conceptual as well as technical innovation.

[63] Hafizi et al., "Inhibition-Dominated Rich-Club Shapes Dynamics in Cortical Microcircuits in Awake Behaving Mice."

REFERENCES

Adesnik, Hillel, William Bruns, Hiroki Taniguchi, Z. Josh Huang, and Massimo Scanziani. "A Neural Circuit for Spatial Summation in Visual Cortex." *Nature* 490, no. 7419 (October 10, 2012): 226–31. https://doi.org/10.1038/nature11526.

Afrashteh, Navvab, Samsoon Inayat, Edgar Bermudez Contreras, Artur Luczak, Bruce L. McNaughton, and Majid H. Mohajerani. "Spatiotemporal Structure of Sensory-Evoked and Spontaneous Activity Revealed by Mesoscale Imaging in Anesthetized and Awake Mice." *BioRxiv*, May 25, 2020, 2020.05.22.111021. https://doi.org/10.1101/2020.05.22.111021.

Agnes, Everton J., Andrea I. Luppi, and Tim P. Vogels. "Complementary Inhibitory Weight Profiles Emerge from Plasticity and Allow Flexible Switching of Receptive Fields." *Journal of Neuroscience* 40, no. 50 (December 9, 2020): 9634–49.

Ainsworth, Matt, Shane Lee, Mark O. Cunningham, Roger D. Traub, Nancy J. Kopell, and Miles A. Whittington. "Rates and Rhythms: A Synergistic View of Frequency and Temporal Coding in Neuronal Networks." *Neuron* 75, no. 4 (August 23, 2012): 572–83. https://doi.org/10.1016/j.neuron.2012.08.004.

Barral, Jérémie, and Alex D Reyes. "Synaptic Scaling Rule Preserves Excitatory–Inhibitory Balance and Salient Neuronal Network Dynamics." *Nature Neuroscience* 19, no. 12 (October 17, 2016): 1690–96. https://doi.org/10.1038/nn.4415.

Berry, Michael J. II, and Gašper Tkačik. "Clustering of Neural Activity: A Design Principle for Population Codes." *Frontiers in Computational Neuroscience* 14 (2020). https://doi.org/10.3389/fncom.2020.00020.

Bojanek, Kyle, Yuqing Zhu, and Jason MacLean. "Cyclic Transitions between Higher Order Motifs Underlie Sustained Asynchronous Spiking in Sparse Recurrent Networks." *PLOS Computational Biology* 16, no. 9 (September 30, 2020): e1007409. https://doi.org/10.1371/journal.pcbi.1007409.

Brette, Romain. "Neural Coding: The Bureaucratic Model of the Brain." *The Behavioral and Brain Sciences* 42 (November 1, 2019): e243. https://doi.org/10.1017/s0140525x19001997.

Brunel, Nicolas. "Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons." *Journal of Computational Neuroscience* 8, no. 3 (May 1, 2000): 183–208. https://doi.org/10.1023/A:1008925309027.

Chambers, Brendan, and Jason N. MacLean. "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks." Edited by Jochen Triesch. *PLOS Computational*

Biology 12, no. 8 (August 19, 2016): e1005078. https://doi.org/10.1371/journal.pcbi.1005078.

Cone, Jackson J., Megan D. Scantlen, Mark H. Histed, and John H. R. Maunsell. "Different Inhibitory Interneuron Cell Classes Make Distinct Contributions to Visual Contrast Perception." *ENeuro* 6, no. 1 (March 11, 2019): ENEURO.0337-18.2019. https://doi.org/10.1523/ENEURO.0337-18.2019.

Dan, Yang, Jose-Manuel Alonso, W. Martin Usrey, and R. Clay Reid. "Coding of Visual Information by Precisely Correlated Spikes in the Lateral Geniculate Nucleus." *Nature Neuroscience* 1, no. 6 (October 1998): 501–7. https://doi.org/10.1038/2217.

deCharms, R. Christopher, and Michael M. Merzenich. "Primary Cortical Representation of Sounds by the Coordination of Action-Potential Timing." *Nature* 381, no. 6583 (June 1996): 610–13. https://doi.org/10.1038/381610a0.

Dechery, Joseph B., and Jason N. MacLean. "Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons." Edited by Jeff Beck. *PLOS Computational Biology* 14, no. 5 (May 4, 2018): e1006153. https://doi.org/10.1371/journal.pcbi.1006153.

Feldmeyer, Dirk, Guanxiao Qi, Vishalini Emmenegger, and Jochen F. Staiger. "Inhibitory Interneurons and Their Circuit Motifs in the Many Layers of the Barrel Cortex." *Neuroscience*, Barrel Cortex Function, 368 (January 1, 2018): 132–51. https://doi.org/10.1016/j.neuroscience.2017.05.027.

Ganmor, Elad, Ronen Segev, and Elad Schneidman. "A Thesaurus for a Neural Population Code." *Elife* 4 (2015): e06134.

Goris, Robbe L. T., J. Anthony Movshon, and Eero P. Simoncelli. "Partitioning Neuronal Variability." *Nature Neuroscience* 17, no. 6 (June 2014): 858–65. https://doi.org/10.1038/nn.3711.

Grinvald, Amiram, Amos Arieli, Misha Tsodyks, and Tal Kenet. "Neuronal Assemblies: Single Cortical Neurons Are Obedient Members of a Huge Orchestra." *Biopolymers* 68, no. 3 (2003): 422–36. https://doi.org/10.1002/bip.10273.

Hafizi, Hadi, Sunny Nigam, Josh Barnathan, Naixin Ren, Ian H. Stevenson, Sotiris C. Masmanidis, Ehren L. Newman, Olaf Sporns, and John M. Beggs. "Inhibition-Dominated Rich-Club Shapes Dynamics in Cortical Microcircuits in Awake Behaving Mice." *BioRxiv*, May 9, 2021, 2021.05.07.443074. https://doi.org/10.1101/2021.05.07.443074.

Harris, Kenneth D. "Neural Signatures of Cell Assembly Organization." *Nature Reviews Neuroscience* 6, no. 5 (May 2005): 399–407. https://doi.org/10.1038/nrn1669.

Hartline, H K., Henry G Wagner, and Floyd Ratliff. "Inhibition in the Eye of Limulus." *Journal of General Physiology* 39, no. 5 (May 20, 1956): 651–73. https://doi.org/10.1085/jgp.39.5.651.

Ikegaya, Yuji, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste. "Synfire Chains and Cortical Songs: Temporal Modules of Cortical Activity." *Science* 304, no. 5670 (April 23, 2004): 559–64. https://doi.org/10.1126/science.1093173.

Jermakowicz, Walter J., Xin Chen, Ilya Khaytin, A. B. Bonds, and Vivien A. Casagrande. "Relationship Between Spontaneous and Evoked Spike-Time Correlations in Primate Visual Cortex." *Journal of Neurophysiology* 101, no. 5 (May 1, 2009): 2279–89. https://doi.org/10.1152/jn.91207.2008.

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

Kasthuri, Narayanan, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongil Lee, et al. "Saturated Reconstruction of a Volume of Neocortex." *Cell* 162, no. 3 (July 2015): 648–61. https://doi.org/10.1016/j.cell.2015.06.054.

Kayser, Christoph, Marcelo A. Montemurro, Nikos K. Logothetis, and Stefano Panzeri. "Spike-Phase Coding Boosts and Stabilizes Information Carried by Spatial and Temporal Spike Patterns." *Neuron* 61, no. 4 (February 2009): 597–608. https://doi.org/10.1016/j.neuron.2009.01.008.

Kotekal, Subhodh, and Jason N. MacLean. "Recurrent Interactions Can Explain the Variance in Single Trial Responses." *PLOS Computational Biology* 16, no. 1 (January 30, 2020): e1007591. https://doi.org/10.1371/journal.pcbi.1007591.

Kriener, Birgit, Moritz Helias, Ad Aertsen, and Stefan Rotter. "Correlations in Spiking Neuronal Networks with Distance Dependent Connections." *Journal of Computational Neuroscience* 27, no. 2 (October 2009): 177–200. https://doi.org/10.1007/s10827-008-0135-1.

Kuhn, Alexandre, Ad Aertsen, and Stefan Rotter. "Higher-Order Statistics of Input Ensembles and the Response of Simple Model Neurons." *Neural Computation* 15, no. 1 (2003): 67–101.

Landau, Itamar D., Robert Egger, Vincent J. Dercksen, Marcel Oberlaender, and Haim Sompolinsky. "The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks." *Neuron* 92, no. 5 (December 2016): 1106–21. https://doi.org/10.1016/j.neuron.2016.10.027.

Lankarany, Milad, Dhekra Al-Basha, Stéphanie Ratté, and Steven A. Prescott. "Differentially Synchronized Spiking Enables Multiplexed Neural Coding." *Proceedings of the National Academy of Sciences* 116, no. 20 (May 14, 2019): 10097–102.

Levy, Maayan, Olaf Sporns, and Jason N. MacLean. "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding." *Cell Reports* 31, no. 2 (April 14, 2020): 107483. https://doi.org/10.1016/j.celrep.2020.03.047.

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. https://doi.org/10.1152/jn.1997.78.3.1714.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. "Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations." *Neuron* 62, no. 3 (May 14, 2009): 413–25. https://doi.org/10.1016/j.neuron.2009.03.014.

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 2005): 811–23. https://doi.org/10.1016/j.neuron.2005.09.035.

Mainen, Z. F., and T. J. Sejnowski. "Reliability of Spike Timing in Neocortical Neurons." *Science* 268, no. 5216 (June 9, 1995): 1503–6. https://doi.org/10.1126/science.7770778.

Markram, Henry, Eilif Muller, Srikanth Ramaswamy, Michael W. Reimann, Marwan Abdellah, Carlos Aguado Sanchez, Anastasia Ailamaki, et al. "Reconstruction and Simulation of Neocortical Microcircuitry." *Cell* 163, no. 2 (October 8, 2015): 456–92. https://doi.org/10.1016/j.cell.2015.09.029.

McCulloch, Warren, and Walter Pitts. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *Bulletin of Mathematical Biophysics* 5 (1943): 19.

Memmesheimer, Raoul-Martin, and Marc Timme. "Non-Additive Coupling Enables Propagation of Synchronous Spiking Activity in Purely Random Networks." *PLOS Computational Biology* 8, no. 4 (April 19, 2012): e1002384. https://doi.org/10.1371/journal.pcbi.1002384.

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 (September 23, 2014): E4053–61.

Montani, F., R. A. A. Ince, R. Senatore, E. Arabzadeh, M. E. Diamond, and S. Panzeri. "The Impact of High-Order Interactions on the Rate of Synchronous Discharge and Information Transmission in Somatosensory Cortex." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367, no. 1901 (August 28, 2009): 3297–3310. https://doi.org/10.1098/rsta.2009.0082.

Najafi, Farzaneh, Gamaleldin F. Elsayed, Robin Cao, Eftychios Pnevmatikakis, Peter E. Latham, John P. Cunningham, and Anne K. Churchland. "Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning." *Neuron* 105, no. 1 (January 8, 2020): 165-179.e8. https://doi.org/10.1016/j.neuron.2019.09.045.

Okun, Michael, Nicholas A. Steinmetz, Lee Cossell, M. Florencia Iacaruso, Ho Ko, Péter Barthó, Tirin Moore, et al. "Diverse Coupling of Neurons to Populations in Sensory Cortex." *Nature* 521, no. 7553 (May 2015): 511–15. https://doi.org/10.1038/nature14273.

Ostojic, Srdjan, Nicolas Brunel, and Vincent Hakim. "How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains." *Journal of Neuroscience* 29, no. 33 (August 19, 2009): 10234–53.

Oviedo, Hysell V. "Connectivity Motifs of Inhibitory Neurons in the Mouse Auditory Cortex." *Scientific Reports* 7, no. 1 (December 5, 2017): 16987. https://doi.org/10.1038/s41598-017-16904-2.

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 (March 29, 2011): 5419–24.

Pernice, Volker, Benjamin Staude, Stefano Cardanobile, and Stefan Rotter. "How Structure Determines Correlations in Neuronal Networks." Edited by Olaf Sporns. *PLoS Computational Biology* 7, no. 5 (May 19, 2011): e1002059. https://doi.org/10.1371/journal.pcbi.1002059.

Rhodes, Paul A. "Recoding Patterns of Sensory Input: Higher-Order Features and the Function of Nonlinear Dendritic Trees." *Neural Computation* 20, no. 8 (August 2008): 2000–2036. https://doi.org/10.1162/neco.2008.04-07-511.

Sakata, Shuzo, and Kenneth D. Harris. "Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex." *Neuron* 64, no. 3 (November 12, 2009): 404–18. https://doi.org/10.1016/j.neuron.2009.09.020.

Sarracino, A., O. Arviv, O. Shriki, and L. de Arcangelis. "Predicting Brain Evoked Response to External Stimuli from Temporal Correlations of Spontaneous Activity." *Physical Review Research* 2, no. 3 (September 2, 2020): 033355. https://doi.org/10.1103/PhysRevResearch.2.033355.

Softky, W. R., and C. Koch. "The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs." *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." Edited by Karl J. Friston. *PLoS Biology* 3, no. 3 (March 1, 2005): e68. https://doi.org/10.1371/journal.pbio.0030068.

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. https://doi.org/10.1038/659.

Stevenson, Ian H., Brian M. London, Emily R. Oby, Nicholas A. Sachs, Jacob Reimer, Bernhard Englitz, Stephen V. David, et al. "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons." *PLOS Computational Biology* 8, no. 11 (November 15, 2012): e1002775. https://doi.org/10.1371/journal.pcbi.1002775.

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. "Spontaneous Behaviors Drive Multidimensional, Brainwide Activity." *Science* 364, no. 6437 (April 19, 2019). https://doi.org/10.1126/science.aav7893.

Trong, Philipp Khuc, and Fred Rieke. "Origin of Correlated Activity between Parasol Retinal Ganglion Cells." *Nature Neuroscience* 11, no. 11 (November 2008): 1343–51. https://doi.org/10.1038/nn.2199.

Vegué, Marina, and Alex Roxin. "Firing Rate Distributions in Spiking Networks with Heterogeneous Connectivity." *Physical Review E* 100, no. 2 (August 12, 2019): 022208. https://doi.org/10.1103/PhysRevE.100.022208.

Wei, Wei, and Marla B. Feller. "Organization and Development of Direction-Selective Circuits in the Retina." *Trends in Neurosciences* 34, no. 12 (December 2011): 638–45. https://doi.org/10.1016/j.tins.2011.08.002.

Yu, Shan, Hongdian Yang, Hiroyuki Nakahara, Gustavo S. Santos, Danko Nikolić, and Dietmar Plenz. "Higher-Order Interactions Characterized in Cortical Activity." *Journal of Neuroscience* 31, no. 48 (November 30, 2011): 17514–26. https://doi.org/10.1523/JNEUROSCI.3127-11.2011.