

Information processing in micro and meso-scale circuits during
normal and disease states

by

Francisco Luongo

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

**Information processing in micro and meso-scale neural circuits during normal
and disease states**

Copyright 2015
by
Francisco J. Luongo

To julio

Acknowledgments

I have many people to thank for making all of this work possible. I would like to thank Vikaas for giving me the opportunity to work in his lab and having the patience to let me figure things out on my own. I would like to thank my committee composed of Michael Stryker, Anatol Kreitzer, and Philip Sabes who provided very helpful comments and advice on my projects along the way. I would also like to thank Mark Goldman for serving as the external examiner during my defense as well as being the the organizer and instructor of the woods hole computational neuroscience course which introduced me to wide array of computational techniques that will no doubt prove invaluable in future scientific endeavors. I would like to thank the many members of the Sohal lab for their continued support both social and experimentally. Tosha Patel especially deserves credit for managing mice and juggling many tasks related to managing such a huge lab. I would like to thank Kathleen Cho for entertaining my cynicism at times and also for being a partner in crime during conferences. I would also like to thank Lowry Kirkby for being a great collaborator on the DARPA SUBNETS work and taking command of the github account so I don't have to. I also owe a huge debt of gratitude to Liqun Luo for giving an undergraduate a chance to work in his lab and be exposed to neuroscience and Tom Clandinin for allowing said undergrad to continue working in his lab for a few years after graduation. I owe alot to the two of them for providing me the opportunity to initially work in neuroscience and for writing me letters that were good enough to get me into graduate school.

I would like to thank Biz whose positive mental attitude (PMA) has been the best thing

to happen to me. She is my life partner and together we have managed to explore, struggle, and learn together. We crack jokes and have so much fun together that it is now hard to imagine my life before her. She is truly my best friend but more importantly my life partner and I look forward to our times together in the future. I also would like to thank my various friends throughout San Francisco who have kept me sane by not discussing science on the weekends. Especially my current roommates Ian, Justine, Allegra, and Eugene who have been both great friends as well as flatmates. I would also like to thank Gerry Shih for being a great buddy to experience with whom to experience San Francisco in my 20s. I would like to thank all of my various music and non-music friends including but not limited to Lily, Alex Blackstock, Petko, Stan, Vesela, Iris, Kristy, Ian, Ryan Dalton, Ryan Basso, Kat, Will Dalton, Ben Kallman, Vanessa, Dina, Patrick, Leah, Amrit, Will, Ashish, Shaanan, Maly, Toby, Quinn, Josh, and of course many others who have made my many years here fulfilling.

Saving the most important for last, I would like to thank my family especially my parents, Cesar and Maria. They have both been incredibly supportive over the years and are the major reason that I am where I am today. They have witnessed me go through many life changes and decisions and have always been supportive of my choices. They instilled in me a drive to succeed that 'ganas' to want to get something done and achieve it for one's own satisfaction. Even at times when I was having a hard time coping with failure, they were always there for me and it was a reassuring thought that I could count on them for support. I also want to thank my sister who started her PhD studies at the same time that I started, it has been great to have another person who understands the trial and tribulations brought on by graduate studies. I want to thank my extended family, my aunt Luz and her boyfriend Giovanni who provide loads of laughs and fun whenever I see them. My uncle Carlos and

his wife Sabrina and my cousins Sebastian, Christina, and Jillian, whose house we invade for family gatherings. I would like to thank my abuelo Julio who while no longer with us, was always a positive thought in my mind and even into old age had an amazing sense of humor. I would like to thank abuelo Efrain who is the academic of his family who also understands what graduate school is like to go through. I would also like to thank my abuela Gaby who although we have major ideological differences she must have done something right because my sister and I turned out fine. On a business trip to Japan a long time ago, my dad brought me back a daruma doll in which you paint one eye and only paint the other once your goal has been achieved. After completing graduate school I can confidently paint that other eye.

Abstract

Information processing in micro and meso-scale neural circuits during normal and disease states

by

Francisco J. Luongo

Doctor of Philosophy in Neuroscience

University of California, San Francisco

Professor Vikaas Sohal, Chair

Neural computation can occur at multiple spatial and temporal timescales. The sum total of all of these processes is to guide optimal behaviors within the context of the constraints imposed by the physical world. How the circuits of the brain achieves this goal represents a central question in systems neuroscience. Here I explore the many ways in which the circuits of the brain can process information at both the micro and meso scale. Understanding the way information is represented and processed in the brain could shed light on the neuropathology underlying complex neuropsychiatric diseases such as autism and schizophrenia. Chapter 2 establishes an experimental paradigm for assaying patterns of microcircuit activity and examines the role of dopaminergic modulation on prefrontal microcircuits. We find that dopamine type 2 (D2) receptor activation results in an increase in spontaneous activity while dopamine type 1 (D1) activation does not. Chapter 3 of this

dissertation presents a study that illustrates how cholinergic activation normally produces what has been suggested as a neural substrate of attention; pairwise decorrelation in microcircuit activity. This study also shows that in two etiologically distinct mouse models of autism, FMR1 knockout mice and Valproic Acid exposed mice, this ability to decorrelate in the presence of cholinergic activation is lost. This represents a putative microcircuit level biomarker of autism. Chapter 4 examines the structure/function relationship within the prefrontal microcircuit. Spontaneous activity in prefrontal microcircuits is shown to be organized according to a small world architecture. Importantly, this architecture is important for one concrete function of neuronal microcircuits; the ability to produce temporally stereotyped patterns of activation. In the final chapter, we identify subnetworks in chronic intracranial electrocorticographic (ECoG) recordings using pairwise electrode coherence and dimensionality reduction techniques. We show that we can further reduce the dimensionality of these networks by identifying 'key-interactions' that are informative of the overall subnetwork state at any given point in time. This study highlights that redundancy in ECoG data can be exploited to identify low-dimensional representation of brain-wide subnetworks. Taken together, these studies represent the development of multiple technological and analytical techniques aimed at understanding how information is processed and modulated at emergent circuit and network levels as well as understanding their dysfunction in a neuropsychiatric disease state.

Contents

Contents	viii
List of Figures	xi
List of Supplementary Figures	xii
List of Tables	xiv
1 Introduction	1
2 Modulation of prefrontal networks by dopamine	6
2.1 Introduction	6
2.2 Results	8
2.3 Discussion	12
2.4 Material and Methods	15
2.5 Figures	21
3 A microcircuit phenotype in autism models	29
3.1 Introduction	29
3.2 Results	32
3.3 Discussion	36
3.4 Material and Methods	41
3.5 Figures	47
4 Small world networks and patterned activity	53
4.1 Introduction	53
4.2 Results	56
4.3 Discussion	69
4.4 Material and Methods	75
4.5 Figures	83
5 Extracting network structure from ECOG data	100
5.1 Introduction	100
5.2 Results	102
5.3 Discussion	110
5.4 Materials and Methods	114

5.5	Figures	118
5.6	Tables	129
6	Concluding remarks	138

Bibliography

List of Figures

2.1	Single photon GCaMP imaging resolves simultaneous activity from many neurons in prefrontal slices	21
2.2	Spontaneous network activity contains both stereotyped patterns and a scale free structure	22
2.3	Activating D2 receptors increases variable activity and motifs	23
3.1	Single photon GCaMP imaging resolves simultaneous activity from many neurons in prefrontal slices	47
3.2	Cholinergic modulation decorrelates microcircuit activity in wild-type mice, but not in models of autism	48
3.3	Autism models, but not DISC1 mutant or fluoxetine-treated mice, exhibit abnormally elevated correlations in carbachol	49
4.1	Single photon imaging of GCaMP signals resolves simultaneous activity from many neurons in prefrontal slices	83
4.2	Spontaneous prefrontal network activity is enriched in positive correlations and events in which multiple neurons are co-active.	84
4.3	Prefrontal microcircuits spontaneously generate many more stereotyped sequences of activity than expected by chance.	85
4.4	Generation of surrogate datasets with various functional organizations	86
4.5	Small world networks outperform clustered and random networks, and reproduce the levels of sequences observed in actual data	87
5.1	Extracting network wide structure in ECoG data use pairwise coherence and dimensionality reduction.	119
5.2	Analyzing principal component projections and extracting lower dimensional representations.	120
5.3	Key interactions present across patients.	121
5.4	Subnetworks identified via PCs are frequency specific	122
5.5	PCA and ICA yield similar subnetworks.	123

List of Supplementary Figures

2.1	Standard deviation projections from sample experiment	24
2.2	GCaMP3 signal with corresponding whole cell recording	25
2.3	Summary of percent time each cell is active	26
2.4	Dependence of correlations on distance	27
2.5	Cumulative distribution plot for the intervals between the first and last occurrences of each motif	28
2.6	Time alone does not reproduce the effects of quinpirole on network activity	28
3.1	Differences in activity levels do not explain the inability of cholinergic modulation to decorrelate activity in autism models	50
3.2	The abnormally increased correlations observed in two models, relative to controls, in carbachol, are not an artifact of differences in activity levels	51
3.3	VPA-exposed datasets continue to exhibit an increase in strong correlations, even after matching the higher levels of activity in saline-exposed datasets	52
4.1	. The small-world organization of correlations is robust to changes in the event detection threshold	88
4.2	The small world organization is maintained even after determining a unique threshold for strong correlations for each cell.	89
4.3	Prefrontal microcircuits are more clustered than expected based on a random, distance-dependent coupling rule	90
4.4	Path lengths and clustering coefficients are both independent of network size	91
4.5	Durations of identified multineuron sequences	91
4.6	Sample activity rasters from surrogate datasets.	92
4.7	Example correlation matrices from surrogate datasets.	93
4.8	Distribution of correlations between neurons in surrogate datasets and the most similar neurons in actual datasets.	94
4.9	Path lengths and clustering coefficients for surrogate datasets	95
4.10	Small world surrogate datasets based on correlation matrices from one experiment and activity rasters from another outperform clustered and random networks.	96
4.11	Number of patterns increases as a function of small-worldness.	97
4.12	Distribution of number of sequence length vs. number of repeats across all experiments.	98
4.13	Datasets exposed to active ACSF instead of carbachol also exhibit a small-world organization.	99

5.1	Method for generating shuffled surrogates.	124
5.2	Eigenvalue distributions of real vs. shuffled data across patients.	124
5.3	Descriptive statistics for patient EC71 alpha band data.	125
5.4	Example alpha band PC data from EC77.	125
5.5	Example alpha band PC data from EC80.	126
5.6	Example alpha band PC data from EC81.	126
5.7	Explained variance from top 3 key interactions in each principal component across frequencies and patients.	127
5.8	Key interactions present across multiple frequencies.	128

List of Tables

5.1	Summary of coverage in each patient	130
5.2	Depression/Anxiety indices for each patient	131
5.3	Key interactions: Delta (1-3Hz) band	132
5.4	Key interactions: Theta (4-7Hz) band	133
5.5	Key interactions: Alpha (8-15Hz) band	134
5.6	Key interactions: Beta (15-30Hz) band	135
5.7	Key interactions: low Gamma (30-70Hz) band	136
5.8	Key interactions: high Gamma (70-120Hz) band	137

Chapter 1

Introduction

Neuroscience has come a long way from the days of Hodgkin and Huxley who were performing seminal recordings on the squid giant axon. Single unit recordings have endowed us the ability to probe and understand that neurons fundamentally communicate using electrical signals. The challenge moving forward lies in understanding the thesaurus of this communication between neurons. In recent years, the ability to record from larger numbers of neurons has become a reality. (Buzsáki, 2004; Grienberger and Konnerth, 2012; Stirman et al., 2014)

We also now know that neurons come in many different varieties of neurotransmitter released, morphologies, receptor composition, and genetic markers. (Markram et al., 2004; McConnell, 1991) This diversity in the cell-autonomous properties are only further compounded when one considers the networks in which these neurons function. The input matrix or the sum total of all cell-types that a cell receives can have just as much of an effect on its activity as other cell-intrinsic factors. In addition, the responses of cells even to a standard input can be dynamically affected by either state or neuromodulatory changes adding further complexity to the system. (Arnsten et al., 2012; Hasselmo, 1995)A greater understanding

of these mechanisms would also allow insight into the neuropathologies underlying complex neuropsychiatric diseases such as autism, depression, and schizophrenia.

Fundamental to understanding the mechanisms of information processing is the ability to read out and perturb the activity from ensembles of neurons within the context of a circuit. Over the past several years, there have been tremendous strides made in the ability to record from multiple neurons using either optical or electrical signal readouts. (Buzsáki, 2004; Freeman, 2015; Grienberger and Konnerth, 2012) These methods can allow for the simultaneous recording of 100 neurons across multiple cortical/subcortical areas. Understanding how patterns at an emergent microcircuit and network level are produced represents a fundamental question in systems neuroscience.

The studies presented in this dissertation represent my efforts over the past years to learn about the mechanisms of information processing at the emergent level of neural circuits and networks by exploiting high density optical and electrical recording techniques. The goal of all of these projects is to explore the statistical patterns present at an emergent network level, explore how these patterns are produced by the intrinsic properties of networks, and lastly use this knowledge to explore novel biomarkers for neuropsychiatric disease conditions.

The first study presented in chapter 2 represents experiments that laid the ground work for thinking about some of these questions of emergent statistical properties at the level of microcircuits as well as troubleshooting for the methods to be used in future studies. This study was designed to explore the emergent representations of the application of Dopamine Type 1 (D1) and Dopamine Type 2 (D2) receptor agonists to acute prefrontal slices. In

order to study the microcircuit level representation of these neuromodulators requires the ability to assay multiple neurons within the network simultaneously. To achieve this, we developed a single-photon GCaMP imaging approach which allowed us to record the activity of 70-100 neurons in a slice over the course of an hour. We also implemented automated analysis techniques that allowed us to automatically segment these cells, extract the periods of activity, and generate rasters of network activity for further analysis. While this study produced mildly interesting results, it really laid the experimental ground work for the work carried out in chapters 3 and 4.

Chapter 3 explores the control of correlations between neurons in the microcircuit by cholinergic modulation and identified a convergent deficit in this mechanism that is conserved across etiologically distinct models of autism. In this study, we identified a cholinergically mediated decorrelation of pairwise correlations in prefrontal microcircuits. This deficit results in an increase of functional correlations in autism models compared to WT controls under normal cholinergic conditions. Importantly this cholinergic deficit is specific to autism models and not general prefrontal pathologies, highlighting the potential of this deficit to represent a microcircuit level biomarker of autism.

In chapter 4 we attempt to elucidate the structure/function relationship within a prefrontal microcircuit. In this study we identified that functional correlations between prefrontal neurons follow a small-world organization during spontaneous activity in isolated prefrontal microcircuits *in-vitro*. Small-world networks exhibit higher-order structure similar to clustered networks in the correlation matrix but importantly retain the ability to

pass information between neurons in an efficient manner analogous to random networks. We then developed a novel analytical method for generating datasets with arbitrary network organization of the correlation structure in order to assay the relative effect of small-world networks on producing patterns of neural activity. We find that small-world networks outperform both clustered and random networks in producing these patterns, highlighting the importance of a small-world organization on the production of neuronal patterns.

In chapters 2-4, we explore the statistics and neuromodulation within microcircuits of rodent prefrontal cortex. In chapter 5 we transition to examining the network-wide patterns of activity present in human intracranial ECoG datasets. In this study we attempt to identify network-wide interactions between multiple regions in the mesolimbic circuit of humans. We use dimensionality reduction on pairwise coherences between regions to identify patterns of interactions across multiple regions that account for large amounts of the variance. We term these patterns subnetworks and show that these subnetworks exhibit redundancy which allows the activity within subnetworks to be predicted using only a few key interactions between regions. These key interactions could represent either subnetwork readouts or drivers of activity across multiple mesolimbic brain regions and thus represent foci of interest for the potential treatment of depression and anxiety using invasive brain stimulation.

Taken together these studies represent many years of work towards understanding the mechanisms of information processing in neuronal microcircuits and networks. As technologies for recording and perturbing neuronal activity allow experiments involving more neurons and more regions, an important question will be how to design models or exper-

iments to properly assay the emergent properties of neuronal networks. (Marder, 2015)

The experiments and analyses outlined here will hopefully contribute towards this common goal of understanding the role of neuronal microcircuits and networks in fundamentals of information processing and how these mechanisms go aberrant in neuropsychiatric disease states.

Chapter 2

Patterned activity in prefrontal networks and modulation by dopamine

2.1 Introduction

A basic question about neuronal circuits is what patterns of activity do they produce, and how are these altered under various conditions that are relevant to behavior. In particular, recurrent networks in deep layers of the neocortex give rise to internally generated activity (Luczak et al., 2009; Sakata and Harris, 2009; Sanchez-Vives and McCormick, 2000; Shu et al., 2003; Takeuchi et al., 2011). Such internally generated activity in the prefrontal cortex (PFC) may form the basis of persistent firing during the delay period of working memory

tasks (Fuster and Alexander, 1971) and contribute to other PFC-dependent cognitive functions such as set-shifting and decision making (Curtis and Lee, 2010; Durstewitz et al., 2010). However several questions about internally generated activity in the PFC and elsewhere remain. First, are local prefrontal circuits sufficient to generate structured patterns of activity suitable for the sorts of functions outlined above? A previous study observed persistent firing in prefrontal slices (Sanchez-Vives and McCormick, 2000), however, this occurred in the context of network-wide UP states, leaving unclear whether isolated prefrontal networks generate the sorts of sparse patterns that might subserve functions such as working memory. A later study observed repetitive sequences of sparse activity in prefrontal slices, however this was in very young (P13-P22) mice (Ikegaya et al., 2004), leaving unclear whether mature circuits produce structured activity.

Second, a major hypothesis, based on cellular, clinical, and behavioral studies, is that dopamine D2 receptors (D2Rs) make internally generated prefrontal activity more variable in ways that can facilitate behavioral adaptation and, under pathological conditions, contribute to schizophrenia (Durstewitz and Seamans, 2008; Winterer and Weinberger, 2004). Specifically, D2R activation may enable prefrontal circuits to explore more possible configurations of network activity. However, no study has yet demonstrated that these effects actually occurs, nor defined exactly how such effects might come about. Here, we imaged neuronal activity in prefrontal slices to answer three questions: (1) can isolated prefrontal networks generate structured patterns of activity? (2) If so, what is the nature of this structure? (3) How do D2Rs alter these patterns?

2.2 Results

We used single-photon, wide field imaging to capture fluorescent signals from the genetically encoded Ca²⁺ indicator GCaMP3 (Tian et al., 2009a) in acute brain slices from young adult (P41-57) mice. We recorded sparse, robust GCaMP3 signals from neurons in layer 5 (L5) of the medial prefrontal cortex (mPFC; Fig. 2.1; Supplementary Fig. S2.1). All experiments included a low concentration of the cholinergic agonist carbachol (2 μ M) in the bath to model basal cholinergic tone in vivo and promote spontaneous network activity. Using a combination of independent component analysis and image segmentation (Mukamel et al., 2009a), we identified the locations of neurons, measured their GCaMP3 signals, and detected events corresponding to increased activity in these neurons (Fig. 2.1A-C). As in previous studies using GCaMP3, spiking above a threshold rate produces an approximately linear increase in GCaMP3 fluorescence which eventually saturates (Fig. 2.1D; Supplementary Fig. S2.2) (Tian et al., 2009a; Yamada and Mikoshiba, 2012). We further confirmed that during carbachol-induced spontaneous activity, detectable increases in GCaMP3 signals correspond to bursts of spikes in individual prefrontal neurons (Fig. 2.1D).

Each experiment recorded signals from 50-80 active L5 neurons over $\tilde{1}$ hour. Activity in each neuron varied widely over the course of an experiment, with periods of high activity lasting tens of seconds interspersed with periods of quiescence (Fig. 2.1C; Supplementary Fig. S2.3). To determine whether this activity is randomly distributed across the network and in time, versus organized according to some structure, we measured structure in three ways: computing correlations between neurons, detecting synchronous bursts of activity

across the network, and searching for small groups of cells that repeatedly become co-active over tens of minutes.

For each experiment, we computed correlations between activity in different neurons, and found that the distribution of correlations included many more strong correlations than would be expected by chance (Fig. 2.2A-B). We compared these correlations to those obtained after shuffling our data. Notably, neurons exhibit long-lasting periods of increased activity interspersed with periods of quiescence such temporal structure could strongly influence correlations, because neurons whose periods of high activity happen to coincide will have high correlations. Typical ways of shuffling, e.g. rearranging inter-event intervals, would destroy this temporal structure and thus underestimate the number of strong correlations expected by chance. Therefore, to preserve this temporal structure in shuffled data, we simply shifted large chunks of each neurons event train in time (Methods). As shown in Figs. 2.2A-B, we found that real data contained an excess of strong correlations between different neurons as compared to this shuffled data. Strong correlations were present even for pairs of neurons separated by large distances (Supplementary Fig. S2.4).

Whereas correlations measure structure at the level of neuron pairs, we have measured simultaneous activity from many (≈ 50) neurons, and therefore, sought to characterize activity at the level of the full network. First, we simply divided each experiment into 1 sec windows, detected network events, in which more than one neuron was active in the same window, and defined the event size as the number of active neurons. We found that the distribution of event sizes followed a power law. Shuffled data consistently contained many fewer large

events (Fig. 2.2C), and the frequency with which events occurred decayed more slowly as a function of event size, for real vs. shuffled data (Fig. 2.2D). Thus network activity has a scale free structure, i.e. there is no characteristic size for network events.

A common hypothesis about recurrent excitatory networks in L5 of the PFC and in other brain regions is that they generate stereotyped patterns of activity that may serve to represent specific decisions or mnemonic information. To test this hypothesis, we divided each hour long movie into 1 sec windows, and looked for combinations of neurons that were co-active during the same window. For each combination, we counted the number of times these neurons became co-active in the same 1 sec window. We found that every experiment contained many such motifs, small groups of neurons that repeatedly become co-active over tens of minutes (see example in Fig. 2.2E) 9. We compared the numbers of times that motifs of various sizes recurred in real and shuffled data, and found that most experiments in control ACSF + carbachol (6/7) contained many motifs that occurred more often than expected by chance (Fig. 2.2F,G). Moreover, these motifs persist over extremely long timescales, >10 min (Supplementary Fig. S2.5).

Finally, we found that after applying the D2R agonist (-)quinpirole (10 μ M) for 10 min, network activity increased (Fig 2.3A,B; n=9 experiments). This increase in the activity of individual cells was accompanied by a corresponding increase in the number of large network events, defined as above (Fig. 2.3C). No such increase occurred for experiments maintained in control ACSF for a similar duration (Supplementary Fig. S2.6; n=7 experiments). Besides increasing the amount of network activity, quinpirole increased the rate at which this

activity changes, as indicated by an increase in the rate of decay of the autocorrelation (the normalized dot product between network activity at one time and activity after a delay) (Fig. 2.3D). It turns out that the specific manner in which quinpirole increases network activity plays a critical role in accelerating the rate at which this activity evolves. Specifically, network activity could increase either because the probability that an inactive neuron becomes active increases, and/or because the duration of each active state increases. Quinpirole selectively increased the probability that a neuron becomes active, but did not increase the duration of active states (Fig. 2.3B). In fact, quinpirole produced a very small, but statistically significant reduction in the duration of active states (Fig. 2.3B). This selectivity is important, as the following analysis demonstrates.

For each experiment, we created two surrogate datasets based on the pattern of network activity recorded in control conditions. Each surrogate dataset was created, as described below, to match the level of activity observed in each neuron after applying quinpirole (Methods). For the first surrogate dataset, we simply prolonged each active state observed in control conditions by a fixed percentage. For the second surrogate dataset, we introduced additional active states into the control pattern of activity. In this case, we also matched the distribution of active state durations observed in quinpirole. Remarkably, increasing network activity in these two distinct ways produced markedly different effects on the stability of network activity, again measured using the autocorrelation. Prolonging each active state markedly slowed the decay of the autocorrelation (green line in Fig. 2.3D), corresponding to a stabilization of network activity. By contrast, increasing the number of active states

accelerated the decay of the autocorrelation, reproducing the effect of quinpirole (blue line in Fig. 2.3D). Finally, we measured the effects of quinpirole on motifs. As shown in Fig. 2.3E and 2.3F, the occurrence of motifs increased dramatically after applying quinpirole ($p < 10^{-5}$; $n=9$ experiments). Notably, modifying control datasets to match the level of activity observed in quinpirole by either prolonging active states or randomly adding additional active states does not reproduce this quinpirole-induced increase in the occurrence of motifs (Fig. 2.3F; $p < 0.001$; $n=9$ experiments). Thus, this effect of quinpirole is not simply a consequence of increasing the level of network activity. As a further test, we created a surrogate dataset based on recordings in quinpirole, by dividing the quinpirole dataset into 30 second windows, identifying all of the active states and quiescent periods within each window, then rearranging the active states and quiescent periods in each window. Thus, firing rates, distributions of active state durations, and distributions of quiescent period durations, will all be preserved for each neuron and each 30 second window. The number of occurrences of motifs was much greater for quinpirole datasets than for these surrogate dataset (purple line in Fig. 2.3F; $p < 10^{-7}$; $n=9$ experiments). Thus, the quinpirole-induced increase in motifs must reflect coordinated fluctuations in activity across subsets of neurons on relatively rapid timescales < 30 sec.

2.3 Discussion

In summary, we report three main findings. First, single-photon, wide-field imaging of GCaMP3 signals can simultaneously measure activity in > 50 neurons. This approach re-

solves activity in similar numbers of cells to what has been reported using two-photon imaging (Borghuis et al., 2011). Second, isolated prefrontal networks generate spontaneous activity that contains both stereotyped patterns and a scale free structure. Importantly, isolated prefrontal networks suffice to generate the sorts of sparse, metastable patterns hypothesized to mediate functions such as working memory and decision making. These patterns of network activity, in which groups of neurons exhibit periods of co-activity interspersed with periods of quiescence, resemble results from recent modeling studies (Ashok and Doiron, 2012), as well as the sparse, stereotyped sequences of activity previously observed in very young tissue (Ikegaya et al., 2004). In addition, we demonstrated that these repetitive motifs of activity are embedded within larger patterns that have a scale-free structure. This links our study with a large body of work on cortical networks operating at criticality (Stewart and Plenz, 2006). Importantly, whereas those studies observed neuronal avalanches in LFP recordings, here we observe similar phenomena (network events) in recordings of simultaneous activity from ≈ 50 neurons.

Finally, we demonstrated that, as has long been hypothesized, D2R activation increases the variability of prefrontal activity, and may enable prefrontal networks to more effectively sample many possible configurations of network activity. Furthermore we defined exactly what this means, and how it comes about. D2R activation results in (1) an increase in the rate at which network activity changes, and (2) a dramatic increase in the number of motifs, small sets of neurons that are repeatedly coactive. Notably, these effects do not occur simply because D2R activation increases noise, i.e. randomly distributed activity.

Rather D2R activation specifically increases the rate at which cells become active, and these increases in activity are coordinated across neurons on rapid timescales so as to increase the occurrence of motifs. This elucidates specific mechanisms through which D2Rs modify network activity in ways that are likely important for both normal and pathological behavior.

Thus, we have demonstrated a new, widely applicable, approach for analyzing patterns of activity generated by specific circuits. Specifically, the ability to measure simultaneous activity in large numbers of neurons, even with limited temporal resolution, can yield important insights about the structure of activity at the network level.

2.4 Material and Methods

All experiments were conducted in accordance with procedures established by the Administrative Panels on Laboratory Animal Care at the University of California, San Francisco.

Slice preparation

Wild-type P26 P33 C57BL/6 mice (Charles River) were injected unilaterally with 1 μ l of AAV5/2-CaMKII::GCaMP3 (UNC) at the coordinates (in mm) 1.7 anterior-posterior (AP), 0.3 mediolateral (ML), and -2.2 dorsoventral (DV). 350 micron thick coronal slices were prepared from these animals 15-27 days after injection. Slices preparation followed our previously described protocol 18. The one notable deviation from the published protocol was that immediately after brain slices were prepared, they were transferred to an N-Methyl-D-Glucamine (NMDG)-based recovery solution for 10 min before being transferred to ACSF for the remainder of their recovery. The NMDG-based solution was maintained at 32° C, and consisted of the following (in mM): 93 N-Methyl-D-Glucamine (NMDG), 93 HCl, 2.5 KCl, 1.2 NaH₂PO₄, 30 NaHCO₃, 25 glucose, 20 HEPES, 5 Na-ascorbate, 5 Na-pyruvate, 2 thiourea, 10 magnesium sulfate, 0.5 calcium chloride. This NMDG preparation method was used to improve the overall health of adult slices to ensure sufficient amounts of activity for analysis. ACSF contained the following (in mM): 126 NaCl, 26 NaHCO₃, 2.5 KCl, 1.25 NaH₂PO₄, 1 MgCl₂, 2 CaCl, and 10 glucose. All recordings were at 32.5 \pm 1° C. All drugs were dissolved in water to make 10mM and 20mM stock aliquots for the (-) quinpirole and carbachol respectively, which were thawed and diluted into ACSF before each experiment.

Application of quinpirole began 10-30 min after the start of an experiment and all analysis was performed on data acquired after quinpirole had been applied for at least 10 minutes.

Intracellular and cell attached recording

We obtained somatic whole-cell and cell attached recordings using differential contrast video microscopy on an upright microscope (BX51WI; Olympus). Recordings were made using a Multiclamp 700A (Molecular Devices). Patch electrodes (tip resistance 26 MOhms) were filled with the following (in mM): 130 K-gluconate, 10 KCl, 10 HEPES, 10 EGTA, 2 MgCl, 2 MgATP, and 0.3 NaGTP (pH adjusted to 7.3 with KOH). Spiking was evoked during current clamp recordings via depolarizing current pulses (250ms, -250pA to +250pA).

Imaging

GCaMP3 imaging was performed on a Olympus BX51 upright microscope outfitted with a 20x 1.0NA water immersion lens with 0.5x reducer (Olympus) and ORCA-ER CCD Camera (Hamamatsu Photonics). Illumination was delivered using a Lambda DG4 arc lamp (Sutter Instruments). Light path was delivered through a 472/30 excitation filter, 495nm single band dichroic, and a 496nm long pass emission filter (Semrock). GCaMP3 movies consisted of 36000 frames acquired at 10Hz (1 hr) with 4x4 sensor binning yielding a final resolution of 256 x 312 pixels. Light power during imaging was 100 - 500 μ W/mm². The open source Micro Manager software suite (v1.4, NIH) was used to control all camera parameters and acquire movies Any movies that had significant drift, movement, or lacked significant amounts of

activity were excluded from further analysis. Specifically, when movies that exhibited drift greater than 0.25 soma diameters, the experiment was terminated prematurely, and any data collected was excluded from analysis. Significant movement could also be detected during ICA by the appearance of elliptical rather than circular segments. Experiments fell into two categories: most experiments (16/21) contained spontaneous activity in which at least 50 neurons were active during 1-3% of all frames. The remaining experiments (5/21) exhibited at least an order of magnitude less activity, i.e. most detected neurons were active during $\leq 0.1\%$ of all frames. Slices that fell into the latter category were presumed to be unhealthy, and indeed, the morphologies of many cells in these slices had an unhealthy appearance that was evident under DIC optics. Thus, we excluded experiments in the latter category from further analysis.

Signal extraction

All analyses and signal extraction was performed using MATLAB (Mathworks). Locations of cells were automatically identified using a modified version of the published CellSort 1.1 toolbox [13]. Signals were extracted from movies and the baseline fluorescence function, F_0 , was calculated for every trace using the mode of the kernel density estimate over a 100s rolling window, implemented via the MATLAB function `ksdensity` following the procedure outlined in [20]. All signal traces shown represent normalized versions of the $\frac{F-F_0}{F_0}$ trace.

Threshold based event detection was performed on the traces by detecting increases in $\frac{F-F_0}{F_0}$ exceeding 2.5σ over one second, and then further thresholding these events by keeping

only those events which exceeded a 4σ increase over two seconds. σ is the standard deviation of $\frac{F-F_0}{F_0}$, calculated over the entire movie. Thus all detected events have a deviation of at least 4σ from baseline. After identifying these events in the calcium signal from a cell, the cell was considered active during the entire period from the beginning to the peak of the event. The beginning of the event was defined as the first point for which $\frac{F-F_0}{F_0}$ increases by 2.5σ within 1 second and by 4σ within 2 seconds. The peak of the event was defined as the local maximum of the entire event, from the beginning of the event until $\frac{F-F_0}{F_0}$ returns to the same baseline value. We then created a matrix in which each row corresponds to a neuron, and each column corresponds to a frame. Entries in this matrix were 1 if a given neuron was active during a given frame, and 0 otherwise. All subsequent analyses were performed on this two-dimensional representation of network activity over time (c.f. Fig. 2.1C). Correlations between cells were calculated between the event trains corresponding to those two cells after subtracting the mean level of activity from each event train. Thus the correlation between two event trains and is where is a rolling mean of the event train calculated by convolving with a Gaussian kernel with $\sigma = 50$ sec.

The standard deviation projection in Figure 2.1 was obtained as follows. For each pixel, we computed the standard deviation of $\frac{F-F_0}{F_0}$ over 100 second intervals throughout the movie, then plotted the maximum value of these standard deviations. In Figure S2.1, we simply plotted the standard deviation at each pixel over the indicated 30 second time window from a sample movie.

Detecting network events/motifs

Network events were detected using a version of the network raster that had been downsampled by a factor of 10, i.e., going from 36000 frames at 10 Hz to 3600 frames at 1 Hz. Network events were detected simply by measuring how many cells were active at any given point in this downsampled network raster. Comparisons were made with shuffled data, which was shuffled as described in the main text. Briefly, experiments were divided into 100 chunks, and these chunks were then shifted relative to one another by random amounts, thus preserving the fine temporal structure while breaking temporal structure on longer timescales. Motifs represent unique configurations of neurons within network events. These motifs were counted by taking into consideration the unique cell identity of every network event and also taking into account all of the possible subsets within the network event as well. For example each occurrence of a specific 5 cell network event would contain within it one 5 cell motif, five 4 cell motifs, and ten 3 cell motifs. In this iterative way all motifs in a given dataset were counted and significance was calculated based on comparisons to shuffled versions of the same dataset.

Calculating the autocorrelation

The autocorrelation function was calculated as the normalized dot product between network activity at one time and activity after a delay. I.e. if the vector x_t represents network activity at time t , then the value of the autocorrelation function for a specific delay τ , is given by:

$$A(\tau) = \frac{x_t x_{(t+\tau)}}{\|x_t\| \|x_{(t+\tau)}\|}$$

Statistical analysis

We used Students t tests to compare pairs of groups, ANOVA to compare multiple groups, and the two-tailed Kolmogorov-Smirnov (KS) test to compare pairs of distributions. Error bars where shown indicate standard error of the mean.

2.5 Figures

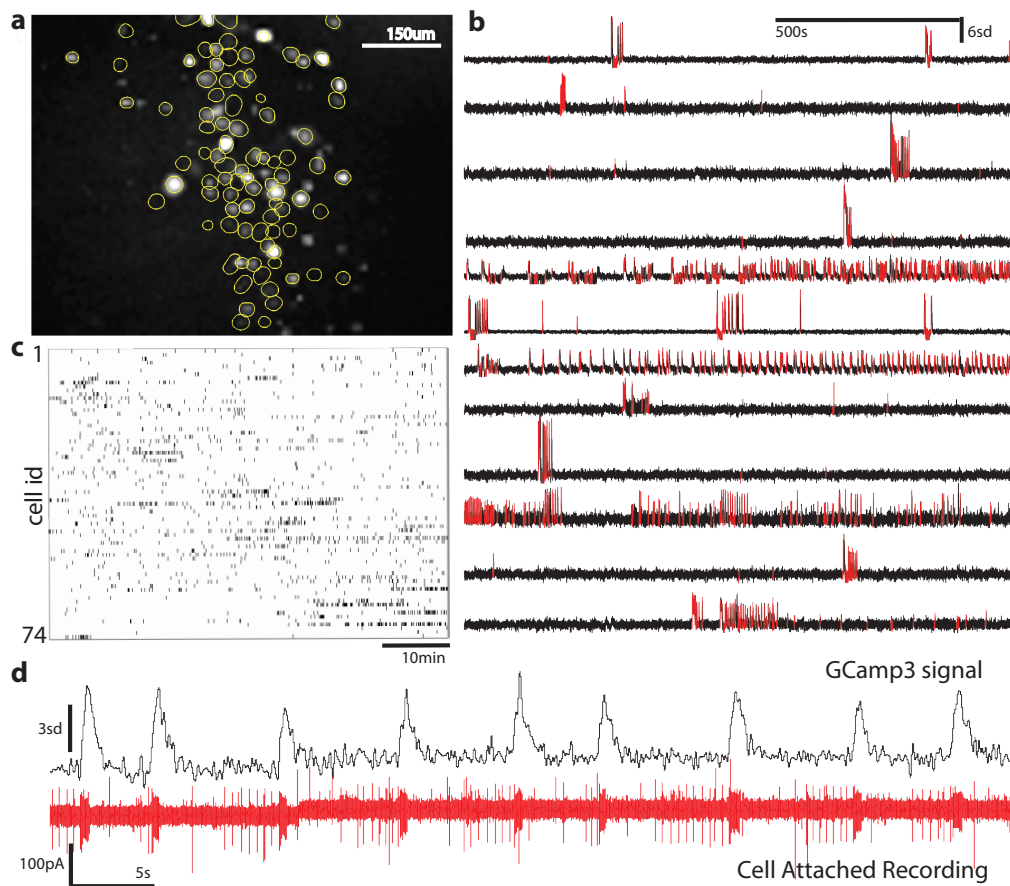


Figure 2.1: **Single photon GCaMP imaging resolves simultaneous activity from many neurons in prefrontal slices.**

a Regions of Interest (ROIs) obtained by an automated algorithm, superimposed on an image showing the maximum standard deviation projection of each pixel during a movie of GCaMP3 fluorescence. (Methods) **b** Sample GCaMP3 signals from 12 neurons. The overlaid red lines indicated times when we detected that each neuron was active. **c** Example raster of spontaneous network activity for a single 60 min experiment with 74 neurons. **d** Simultaneous recording of the GCaMP3 signal (top trace) and spiking measured in cell-attached mode (bottom trace) for a single neuron during spontaneous network activity showing that periods during which the GCaMP3 signal rise correspond to bursts of action potentials.

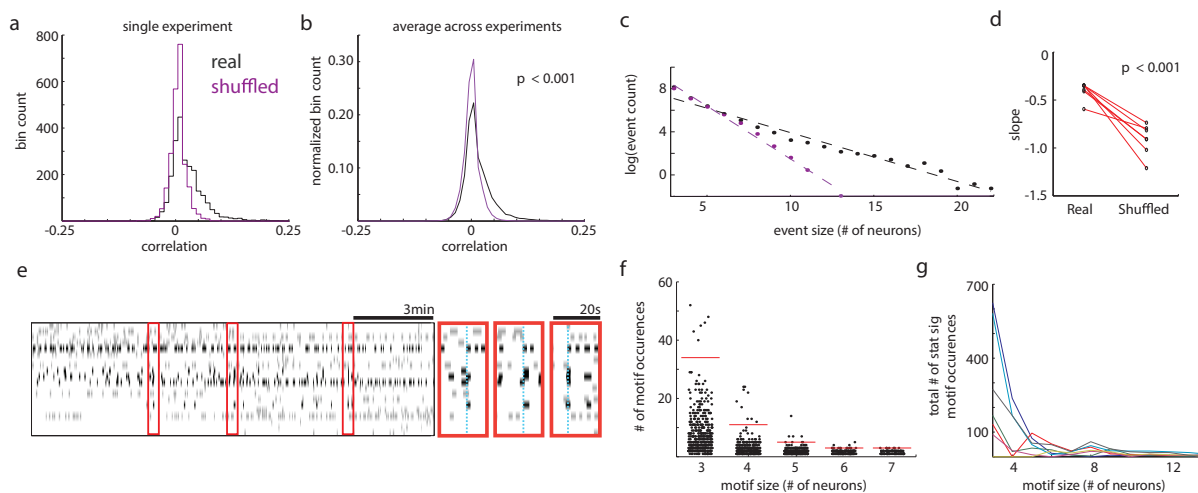


Figure 2.2: **Spontaneous network activity contains both stereotyped patterns and a scale free structure.**

a-b Distribution of correlations between neurons for a single experiment **a**, and the average across all control experiments **b** i.e. experiments maintained for 60 min in control ACSF with carbachol (2 μ M). Real data has fewer zero correlations and is enriched in positive correlations compared to shuffled data ($n=7$ experiments; $p < 0.001$ by two-tailed KS test). **c** Semilog plot showing how often network events, in which several neurons are active simultaneously, occur as a function of their size in real (black) and shuffled (purple) data, again averaged over all control experiments ($n=7$). **d** Population data comparing the slopes of the semilog plot of event frequency vs. event size for each real vs. shuffled control dataset ($n=7$). **e**, Example of a motif, i.e. 5 neurons that are repeatedly co-active. Left: raster of network activity in which the 5 co-active neurons are shown in black. Red boxes outline three time windows during which these 5 neurons are co-active. Right: These three time windows are shown on an expanded timescale, and vertical dotted lines indicate times when the 5 neurons are co-active. **f** How often different motifs occur, as a function of their size (in neurons), for one experiment. Red lines indicate thresholds for significance ($p < 0.01$) based on shuffled datasets. I.e. if the threshold for N -neuron motifs equals X occurrences, then in 99% of shuffled datasets, the maximum number of occurrences of an N -neuron motif was less than or equal to X . **g**, total number of occurrences of statistically significant motifs as a function of motif size. Each trace corresponds to a different experiment ($n=7$).

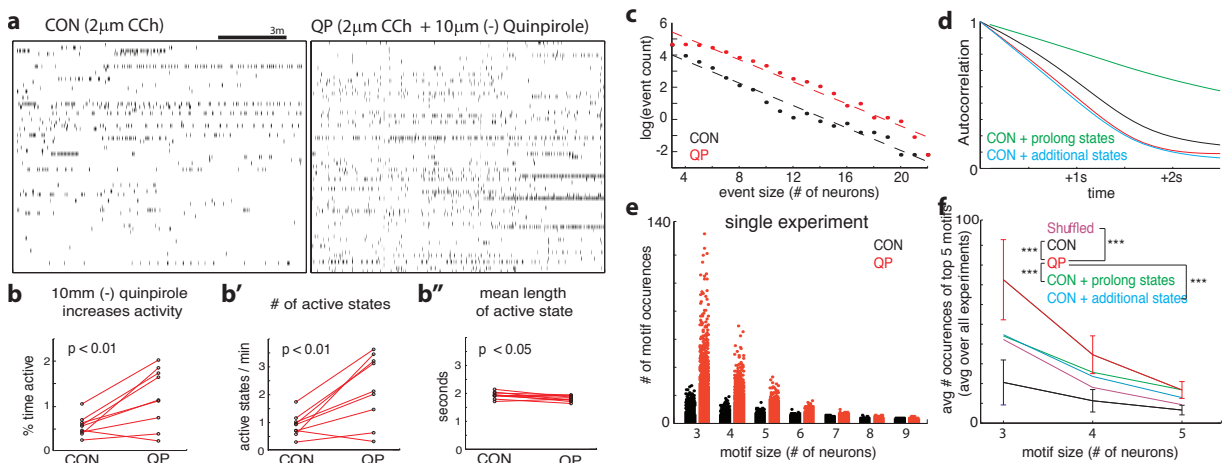




Figure 2.3: Activating D2 receptors increases variable activity and motifs.

a Rasters showing 15 min of network activity in control (CON) and quinpirole (QP) conditions for one experiment. **b** Levels of neuronal activity increase in quinpirole and this reflects an increase in the number of active states per neuron, rather than an increased duration of each active state. **c** Semilog plot showing how often network events occur as a function of their size in control (black) and quinpirole (red) conditions. **d** The autocorrelation function, which measures how quickly network activity evolves, is plotted for control (black) and quinpirole conditions (red), as well as for control datasets that have been modified by prolonging each active state (green), or randomly adding additional active states (blue). **e** How often different motifs occur, as a function of their size (in neurons), for one experiment, in both control conditions (black) and quinpirole (red). **f** For each experiment and condition (control or quinpirole), we identified the five most frequently occurring motifs composed of 3, 4, or 5 neurons. This plot shows the average number of occurrences of these top five motifs, as a function of size (in neurons) for control datasets (black), quinpirole datasets (red), control datasets that have been modified by prolonging (green) or adding (blue) active states, and quinpirole datasets in which active and quiescent states have been shuffled (purple). The mean number of occurrences of these motifs was higher for the quinpirole dataset than for all other datasets by ANOVA using experiment, motif size, and condition as factors ($p < 0.001$, $n = 9$). *** = $p < 0.001$.

Time lapse representation using standard deviation projections of the example movie S1

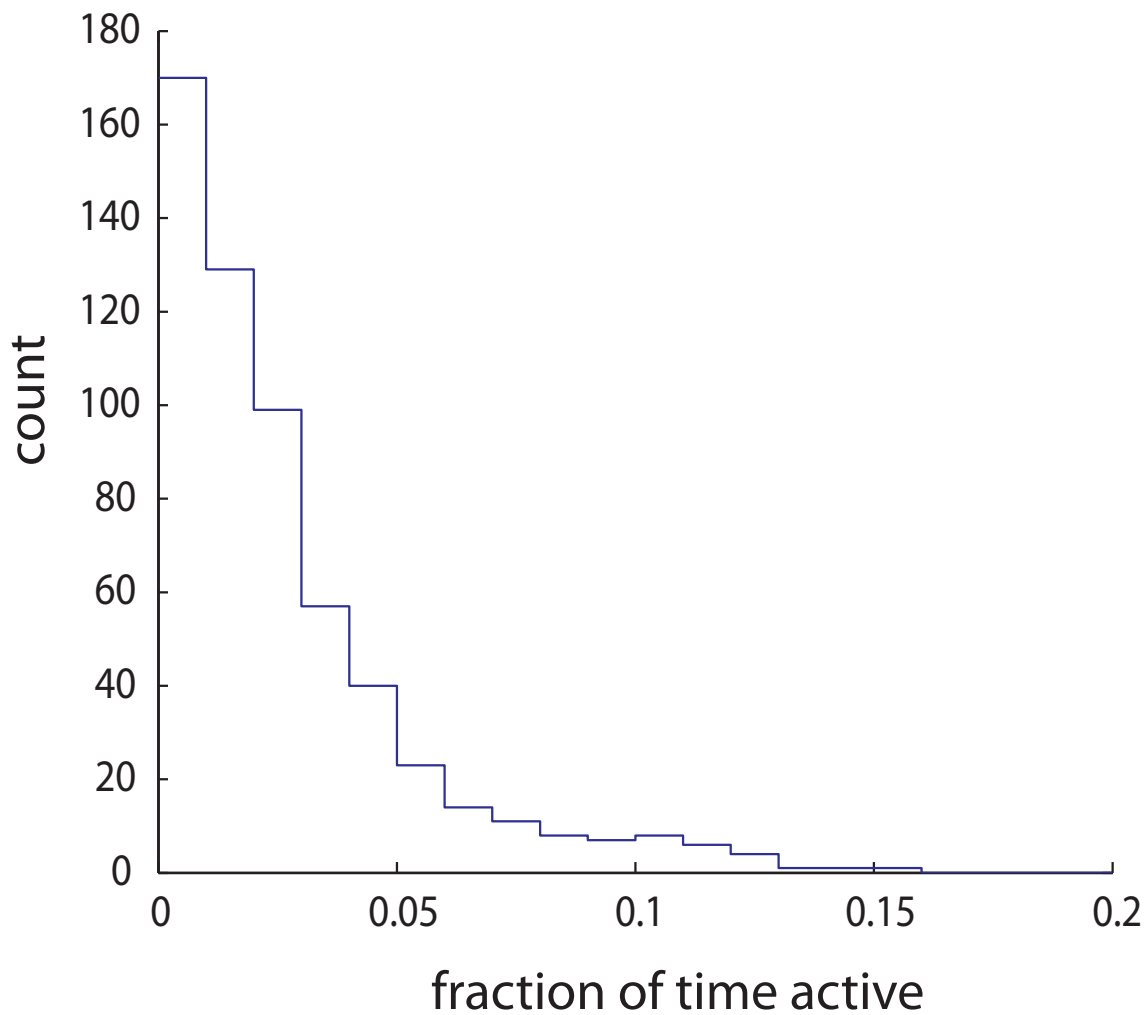


Supplementary Figure 2.1: **Standard deviation projections from sample experiment.** Still images at various timepoints from the standard deviation projection of a movie of GCaMP3 fluorescence in layer 5 neurons in the mPFC. These images illustrate how GCaMP3 signals originating from distinct neurons fluctuate over time.



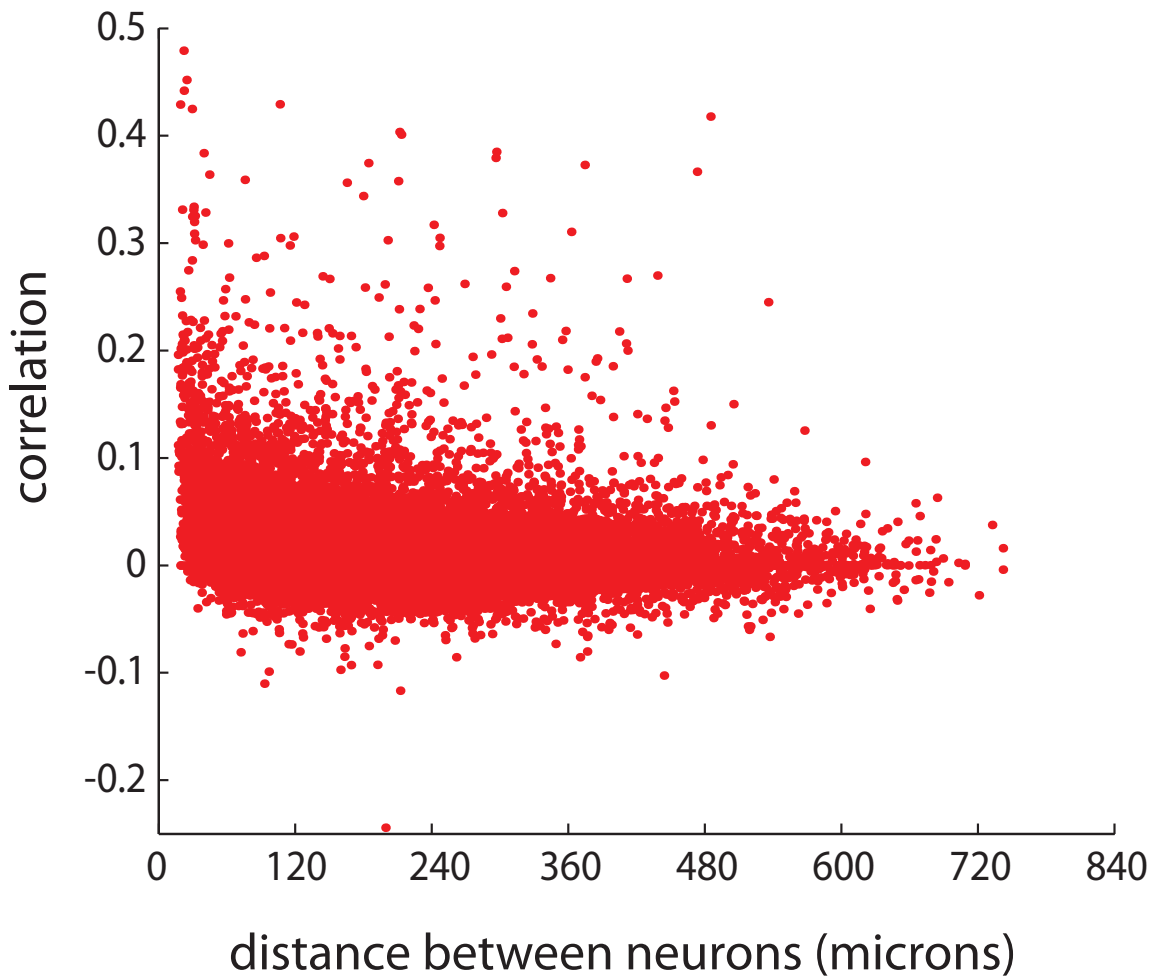
Supplementary Figure 2.2: **GCaMP3 signal with corresponding whole cell recording.** The GCaMP3 signal (top trace) measured from a neuron during a whole cell current clamp recording (bottom trace). Injecting depolarizing current pulses elicits spiking, and spiking above a threshold rate elicits GCaMP3 signals that increase with the spike rate.

Distribution of % time that each neuron was active, totalled over all control experiments (n=7)



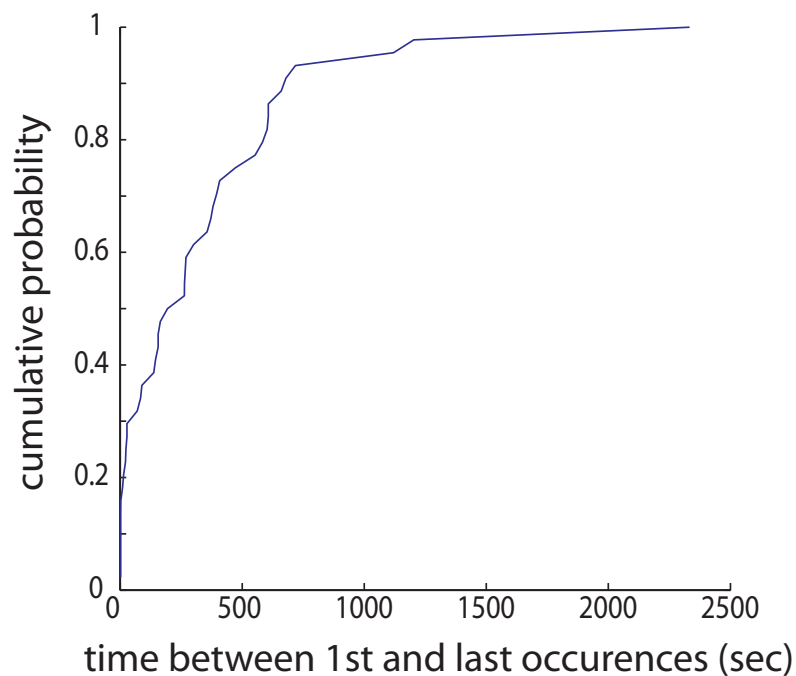
Supplementary Figure 2.3: **Summary of percent time each cell is active.**

Cumulative density function showing that most cells are active for less than 5% of the duration of each experiment.




Supplementary Figure 2.4: **Dependence of correlations on distance.**

Correlations between neurons, plotted as a function of distance between those neurons. The correlation between activity in two neurons is weakly related to the distance between them, but high correlations can be seen even at distances of several hundred microns.



Supplementary Figure 2.5: **Cumulative distribution plot for the intervals between the first and last occurrences of each motif.**



Supplementary Figure 2.6: **Time alone does not reproduce the effects of quinpirole on network activity.**

Same as Fig. 2.3b, but none of the changes seen after applying quinpirole are reproduced by comparing the early and late portions of datasets that were maintained in control ACSF for the duration of the experiment ($n=7$).

Chapter 3

A microcircuit phenotype in autism models

3.1 Introduction

Autism reflects disparate genetic and environmental causes, suggesting that common behavioral phenotypes may reflect convergent defects at the level of neuronal circuits controlling these behaviors. Consistent with this hypothesis, recent analyses have revealed that several genes linked to autism are co-expressed within deep layer microcircuits in the prefrontal cortex (PFC) (Chang et al., 2015; Willsey et al., 2013). Identifying convergent abnormalities located within these microcircuits would yield attractive targets for future circuit-based therapeutic interventions. Until recently, a major barrier to identifying such circuit-level abnormalities has been recording from large numbers of neurons simultaneously. However,

advances in imaging and electrophysiological techniques have made it possible to characterize patterns of circuit-level activity by, for example, calculating correlations between neurons (Runfeldt et al., 2014; Sadoovsky and MacLean, 2013). Despite these advances, cognitive processes are associated with complex neural dynamics embedded within a high dimensional state space. Thus, in order to reveal defects in circuit-level activity that are associated with disorders such as autism, it is valuable to first identify specific neuronal correlates or signatures for relevant cognitive processes. Identifying these signatures makes it possible to evaluate whether they are altered in the setting of diseases such as autism. Any such alterations represent putative pathophysiological mechanisms contributing to cognitive dysfunction in these disease states.

Here, we set out to identify possible microcircuit-level abnormalities associated with autism. Our approach was to look for convergent abnormalities that are 1) conserved across multiple, etiologically distinct models of autism, and 2) impact neural correlates of cognitive processes that are known to be abnormal in autism. In particular, two recent studies identified a possible neural signature of attention using multi-neuronal recording in monkeys performing a visuospatial attention task (Cohen and Maunsell, 2009; Mitchell et al., 2009). Both studies observed the same result: during the attended portion of the task, the pairwise correlations between cortical neurons decreased. This decorrelation would reduce the overall noise of an output signal that was composed of a sum across the population, in fact, this pairwise decorrelation accounted for 80% of the total improvement in the signal-to-noise ratio (Cohen and Maunsell, 2009). This phenomenon is hypothesized to reflect the decorrelation

of spontaneous network activity by neuromodulation. In particular, cholinergic modulation plays a well-established role in attention and cortical decorrelation in vivo (Goard and Dan, 2009). Nevertheless, it remains unknown whether cholinergic modulation can induce such decorrelations by acting directly on cortical microcircuits and if so, whether defects in this mechanism might be present in autism or other conditions with a high comorbidity of attentional deficits. Notably, in Fragile X syndrome (the most common known single gene cause of autism), 70% of affected children meet criteria for an attentional deficit (Cornish et al., 2004). Overall, approximately 50% of children with autism also meet criteria for attention deficit hyperactivity disorder (Boyle and Kaufmann, 2010; Gadow et al., 2005, 2004). Deficits in attention have also been observed in FMR1 KO mice (Casten et al., 2011; Dickson et al., 2013; Moon et al., 2006) and other mouse models of autism (M et al., 2013).

To examine how autism might affect a neural correlate of attention, we studied how cholinergic modulation affects spontaneous activity generated by deep layer prefrontal microcircuits under normal conditions and in mouse models of autism. We focused on microcircuits in deep layers of the PFC because cholinergic modulation within the PFC has been directly implicated in attention (Dalley et al., 2004; Parikh and Sarter, 2008), and abnormalities associated with autism are likely to be intrinsic to deep layer prefrontal microcircuits (Willsey et al., 2013). Therefore, we isolated these microcircuits in brain slices, and used single-photon, wide field GCaMP imaging (Tian et al., 2009b) to measure spontaneous activity in many neurons (50-100) at once. As described below, we found that cholinergic modulation can indeed act directly on cortical microcircuits to decorrelate spontaneous ac-

tivity, mimicking the neural signature previously linked to attention in vivo. Furthermore, this decorrelation, which represents a possible neural substrate for attention, is defective in two etiologically distinct mouse models of autism.

3.2 Results

We recorded sparse, robust GCaMP signals from neurons in layer 5 (L5) of the medial prefrontal cortex (mPFC; Fig. 3.1) in acute brain slices from late adolescent (P41-57) mice ($n = 95$). Using a combination of independent component analysis (ICA) and image segmentation (Mukamel et al., 2009b), we located neurons, measured their GCaMP signals, and detected events corresponding to increased activity in these neurons (Fig. 3.1A-C). Each experiment recorded from 50-100 active L5 neurons for 1 hour. Table 1 lists all of our experiments, performed in control mice or mice modeling autism or other manipulations, and using either carbachol or high K^+ ACSF to elicit spontaneous activity. Table 1 also includes summary statistics for each set of experiments.

We first sought to demonstrate the presence of correlated activity within our datasets. Experiments contained more strong correlations than shuffled datasets, in which the event train for each neuron is shifted in time by a different random amount, or scrambled datasets, in which the neuronal identities associated with each event are randomly reassigned (Fig. 3.1D). Shuffled datasets preserve the temporal structure of activity within each neuron, whereas scrambled datasets preserve the number of neurons active at any given point in time, thus maintaining the temporal structure of activity at the network level. We set an arbitrary

threshold of 0.15, and quantified the fraction of correlations exceeding this threshold. As shown in Fig. 3.1D, correlations exceeding this threshold were present at negligible levels in shuffled or scrambled datasets, so we used this threshold to define *strong* correlations, i.e. correlations which exceed those expected simply by chance.

Given the posited link between cholinergic modulation and cortical circuit decorrelation, we next compared spontaneous prefrontal microcircuit activity elicited by 2 μM carbachol to activity in hi K^+ , low Ca^{2+} active ACSF (Zhao et al., 2011), which contains relatively high K^+ and low Ca^{2+} . Active (or high K^+) ACSF elicits similar levels of spontaneous activity as 2 μM carbachol (avg % time active = 4.2 ± 0.3 in carbachol vs. 4.5 ± 0.5 in high K^+), making it possible to compare the properties of such activity when cholinergic modulation is present or absent, independent of changes in activity levels. These experiments, comparing recordings in carbachol to those in high K^+ , were all done using WT mice on an FVB background, which served as controls for experiments with FMR1 KO mice described below. Consistent with our hypothesis that cholinergic modulation decorrelates prefrontal microcircuit activity, the distribution of pairwise correlations was shifted to the right in high K^+ compared to carbachol (Fig. 3.2A). This resulted in a marked increase in the fraction of 'strong' correlations (correlations ≥ 0.15), from 1.2 ± 0.3 % in carbachol to 4.0 ± 0.9 % in high K^+ (Fig. 3.2A inset).

Notably the decrease in correlations observed in carbachol was not an artifact of differences in activity. Plotting the fraction of strong correlations against the mean % time activity for each experiment in carbachol or high K^+ shows that the fraction of strong cor-

relations is increased in high K^+ compared to carbachol, for all activity levels (Fig. 3.2D). To quantitatively and explicitly account for possible confounding effects of activity on the prevalence of strong correlations, we performed linear regression (Fig. 3.2D), and computed the difference between the observed fraction of strong correlations and the number expected based simply on the level of activity. These residual values are shown in Fig. 3.2E, and confirm that even when activity levels are taken into account, correlations in high K^+ are stronger than those in carbachol. Next, we examined whether this carbachol-induced decorrelation, which represents a possible biological substrate for attention, might be disrupted in two etiologically dissimilar mouse models of autism. First, we studied mice exposed to valproic acid (VPA) in utero at E10.5 (Methods). VPA is an anticonvulsant and mood stabilizer, and there is a markedly elevated rate of autism in the children of mothers treated with VPA (but not other anticonvulsants) during pregnancy (Moore et al., 2000; Rasalam et al., 2005). Rodents exposed to VPA in utero exhibit numerous autism-like phenotypes (Schneider and Przewocki, 2005; Schneider et al., 2008). We also studied FMR1 KO mice, which model Fragile X syndrome, the most common known genetic cause of autism. In stark contrast to the carbachol-induced decorrelation described earlier, in both VPA-exposed and FMR1 KO slices, carbachol failed to elicit a leftward shift of the correlation distribution (Fig. 3.2B,C). In fact, carbachol increased correlations in the VPA-exposed microcircuits (Fig. 3.2C). A potentially complicating factor is that unlike the case in wild-type slices, in both VPA-exposed and FMR1 KO slices, levels of activity were lower in carbachol compared to high K^+ ACSF. We calculated correlations using mean subtracted event trains, which

should minimize activity level-dependent effects. Indeed, as Fig. 3.2D shows, there is at best a very weak relationship between activity levels and correlations. We also carried out additional analyses described in the Supplemental Information which confirm that the absence of a carbachol-induced decorrelation in autism models was not simply an artifact of differences in activity levels (Fig. S3.1).

We reasoned that this loss of the carbachol-induced decorrelation should cause correlations during carbachol-induced activity to be higher in models of autism compared to controls. We compared VPA-exposed mice to age-matched controls whose mothers had been injected with saline at the same timepoint. For FMR1 KO mice, we used age-matched FMR1 WT controls. In carbachol, both autism models exhibit a significant increase in the fraction of strong correlations relative to controls (Fig. 3.3A,B). Notably, levels of carbachol-induced activity were essentially identical in FMR1 WT and KO mice ($4.2 \pm 0.3\%$ and $4.3 \pm 0.4\%$ in WT and KO respectively), demonstrating that the abnormally high correlations we observed during carbachol-induced activity in FMR1 KO mice are not simply an artifact of differences in activity. By contrast, levels of carbachol-induced activity were different for VPA-exposed mice and their wildtype controls, but once again, accounting for these differences could not explain the abnormally increased strong correlations observed in VPA-exposed mice (Fig. S3.2C; SFig. 3.3).

Finally, we evaluated whether this abnormal increase in correlations during carbachol-induced activity possesses specificity for autism. For this, we examined genetic and pharmacologic perturbations that affect the PFC but do not model autism: mice expressing a

dominant negative version of DISC1 in excitatory neurons (CaMKII-TTA / TetO-dnDISC1), and mice chronically treated with the commonly prescribed selective-serotonin reuptake inhibitor fluoxetine. DISC1 disruption models aspects of schizophrenia and depression, including behaviors related to PFC dysfunction, e.g. impaired working memory and decreased social interaction (Clapcote et al., 2007; Hikida et al., 2007; Pletnikov et al., 2008). Chronic treatment with fluoxetine alters cell proliferation in the mPFC (Kodama et al., 2004) in ways that may contribute to its effects on depression-like and repetitive behaviors. If the changes we saw in VPA-exposed and FMR1 KO mice were simply nonspecific consequences of perturbing PFC microcircuits, then we would expect these other mice to exhibit a similar increase in correlations during carbachol-induced activity. However, we observed no such increase in correlations during carbachol-induced activity in either of these cases compared to controls (Fig. 3.3C, D).

3.3 Discussion

In this study we have identified a putative microcircuit level-phenotype associated with autism: a defect in the ability of cholinergic modulation to decorrelate spontaneous cortical circuit activity. We show that this decorrelation, previously observed *in vivo*, can also be observed at the level of isolated cortical microcircuits, and that it is lost in two mouse models of autism. Lastly, we show that this defect is at least somewhat specific for autism, as it is not present following two other manipulations that affect prefrontal cortex: chronic treatment with fluoxetine, and expression of a dominant-negative form of DISC1. This suggests that

the loss of cholinergic-dependent decorrelation observed might represent a microcircuit-level phenotype with specificity for autism, and thus could constitute an attractive target for novel circuit level-therapies.

Relationship to in vivo studies

While we have shown the cholinergic decorrelation to be present in vitro, it will be important to compare our findings to features of activity occurring in vivo. That being said, in vivo recordings address fundamentally different questions from those addressed here. Activity in vivo is strongly influenced by synaptic inputs originating outside the local microcircuit as well as factors such as behavior state. Thus, while in vivo studies are uniquely able to identify network phenomena, e.g., decorrelations recruited by attention (Cohen and Maunsell, 2009; Mitchell et al., 2009) and acetylcholine (Goard and Dan, 2009), they are not well suited for determining how these phenomena relate to local microcircuit mechanisms. For example, even if it were possible to precisely control levels of neuromodulation in vivo, doing so would almost certainly elicit confounding changes in feedforward and feedback inputs. Here, by studying isolated microcircuits within brain slices, we demonstrate unequivocally that the decorrelations observed in vivo can emerge as a result of cholinergic modulation acting directly on microcircuits. (Interestingly, whereas we studied ongoing patterns of spontaneous network activity similar to those believed to mediate attentional effects in vivo, another recent in vitro study (Runfeldt et al., 2014) showed that acetylcholine decorrelates the latency to the first spike in somatosensory neurons responding to thalamic input). We go on to

demonstrate that the carbachol-induced decorrelation of spontaneous network activity we found is disrupted in models of autism. If one observed abnormal correlations in models of autism in vivo, it would be unclear whether these abnormalities reflect the inability of cholinergic modulation to decorrelate local circuits (as we have found) vs. more macroscopic deficits, e.g. the inability of these animals to engage the appropriate behavioral state. Our approach is specifically intended to isolate and examine the integrity of specific microcircuit mechanisms, in a way that complements in vivo work. Our focus on mechanisms that are intrinsic to deep layer prefrontal circuits is particularly relevant given that these circuits represent the single strongest locus of convergence for autism genes (Willsey et al., 2013).

Relevance to autism

Great progress has been made in identifying genes and environmental factors that contribute to autism. However, it has been challenging to understand how all of these 'add up' to alter circuit activity in ways that could impact behavior. Cortical microcircuits, within which many cell types and synapses interact to generate patterns of neural activity, represent an attractive locus at which at which many different genetic, environmental, and developmental lesions could converge to elicit common phenotypes. Thus, many studies have begun exploring possible circuit-level endophenotypes for complex neuropsychiatric disorders including autism (Gibson et al., 2008; Goncalves et al., 2013; Markram et al., 2007). As discussed above, attentional deficits are a major feature of autism generally, and Fragile X syndrome in particular. The failure of cholinergic modulation to decorrelate PFC microcircuit activity

in autism models holds face validity as a possible contributor to these deficits. The presence of this same abnormality in two models of autism with very different etiologies (but not in association with other manipulations of prefrontal function) provides additional empirical validity for the hypothesis that this abnormality plays a role in the pathophysiology of autism. Importantly, this circuit-level abnormality may come about as a result of multiple, dissimilar cellular and synaptic effects in FMR1 KO and VPA-exposed mice. Indeed, as stated earlier, numerous autism-related genes are convergently expressed within the deep layer PFC circuits we have shown are abnormal (Willsey et al., 2013). This highlights the importance of identifying microcircuit-level phenotypes, such as the loss of cholinergic decorrelation, that could represent convergent consequences of disruptions in multiple genes, and which could plausibly contribute to common aspects of autism, e.g. attentional deficits.

Possible circuit mechanisms

Of course, it will still be important to identify the cellular and/or synaptic actions of cholinergic modulation which normally help desynchronize microcircuit activity, but are abnormal in VPA-exposed and/or FMR1 KO mice. Various studies have shown that cholinergic modulation can increase the excitability of interneurons which express somatostatin or vasoactive intestinal peptide (Kawaguchi, 1997; Xiang et al., 1998), hyperpolarize fast-spiking interneurons (36), regulate short-term plasticity at both excitatory and inhibitory synapses (Pafundo et al., 2013), preferentially elicit persistent firing in subcortically-projecting (vs. callosally-projecting) L5 pyramidal neurons (Dembrow et al., 2010), and suppress intracortical (but

not thalamocortical) excitatory synapses (Gil et al., 1997). Of course, it remains unclear whether each of these effects is present in the prefrontal circuits we have studied, contributes to the cholinergic-induced decorrelation we have found, and/or is altered in FMR1 KO and VPA-exposed mice. Rather than to explore each of these effects individually, we instead took the unconventional approach of developing a new assay, using it to examine whether all the effects of cholinergic modulation add up to exert some net influence on circuit activity, and then evaluating whether that net circuit-level influence is altered in a consistent way across VPA-exposed and FMR1 KO mice. Examining individual cellular and synaptic effects of cholinergic modulation would be time consuming, and might miss the 'big picture' of how diverse abnormalities can add up to elicit a common, circuit-level phenotype. By contrast, the approach demonstrated here can explore how various factors shape pairwise decorrelation and other emergent network properties. For example, one could combine optogenetic inhibition of genetically targeted cell populations with GCaMP imaging to explore how specific cell types influence network correlation / decorrelation.

Conclusions

By describing a potential new microcircuit-level phenotype for autism, our study builds a critical foundation for future work to elucidate the detailed pathways, leading from the FMR1 gene or VPA exposure, to the abnormal response of prefrontal microcircuits to cholinergic modulation we have found. One could also evaluate potential therapies by measuring how well they improve the ability of cholinergic modulation to decorrelate these circuits,

this approach might reveal novel therapeutic strategies that bypass the molecular pathways disrupted in autism, and instead exploit alternative mechanisms to normalize overall circuit function.

3.4 Material and Methods

All experiments were conducted in accordance with procedures established by the Administrative Panels on Laboratory Animal Care at the University of California, San Francisco.

Subjects

P26-33 mice of either sex (Charles River) were injected unilaterally with 500 nL of AAV5/2::synapsin::GCaMP6s (UNC viral vector core) at the coordinates (in mm): 1.7 anterior-posterior (AP), 0.3 mediolateral (ML), and -2.2 dorsoventral (DV). Experiments studying the valproic acid (VPA) model of autism used C57BL/6 mice whose pregnant mothers had been injected with a single dose of VPA (500 mg/kg i.p.) at E10.5. For these experiments, control mice were C57BL/6 mice whose pregnant mothers had been injected with saline at E10.5. VPA solution was prepared by dissolving VPA in 0.9 % saline to a final concentration of 150 mg/mL. Experiments studying a mouse model of Fragile X syndrome used male FMR1 WT or KO mice on a FVB background (Jackson labs). In some cases (active ACSF cohorts), these FMR1 WT and KO mice were littermates, while others (carbachol cohorts) were not. For some experiments, C57Bl/6 mice were treated with fluoxetine based on a previously described protocol (Welch et al., 2007). Fluoxetine was administered (5 mg kg⁻¹, I.P.) once

daily for 6 days prior to imaging with the final injection coming 24-48 hours before imaging. Dominant negative DISC1 mutant mice were generated by crossing B6-CamKII::TtA (JAX: 00310) mice with tetO-DISC1dn (JAX: 008790) to yield mice expressing dominant negative DISC1 in neocortical pyramidal cells.

Slice preparation

In all cases, 350 micron thick coronal slices were prepared from these animals 15-27 days after injection (6-8 weeks of age). Slices preparation followed our previously described protocol (17) with one deviation: immediately after brain slices were prepared, they were transferred to an N-Methyl-D-Glucamine (NMDG)-based recovery solution for 10 min before being transferred to ACSF for the remainder of their recovery (18). The NMDG-based solution was maintained at 32° C, and consisted of the following (in mM): 93 N-Methyl-D-Glucamine (NMDG), 93 HCl, 2.5 KCl, 1.2 NaH₂PO₄, 30 NaHCO₃, 25 glucose, 20 HEPES, 5 Na-ascorbate, 5 Na-pyruvate, 2 thiourea, 10 magnesium sulfate, 0.5 calcium chloride. This NMDG preparation method was used to improve the overall health of adult slices to ensure sufficient amounts of activity for analysis. ACSF contained the following (in mM): 126 NaCl, 26 NaHCO₃, 2.5 KCl, 1.25 NaH₂PO₄, 1 MgCl₂, 2 CaCl, and 10 glucose. All recordings were at 32.5 ± 1°C. "Active" ACSF was identical to normal ACSF except with elevated KCl (3.5mM vs. 2mM) and reduced CaCl (1.2mM vs. 2mM).

Imaging

GCaMP imaging was performed on an Olympus BX51 upright microscope with a 20x 1.0NA water immersion lens, 0.5x reducer (Olympus), and ORCA-ER CCD Camera (Hamamatsu Photonics). Illumination was delivered using a Lambda DG4 arc lamp (Sutter Instruments). Light was delivered through a 472/30 excitation filter, 495nm single band dichroic, and 496nm long pass emission filter (Semrock). All movies that were analyzed consisted of 36000 frames acquired at 10Hz (1 hr) with 4x4 sensor binning yielding a final resolution of 256 x 312 pixels. Light power during imaging was 100 - 500 $\mu\text{W}/\text{mm}^2$. The Micro Manager software suite (v1.4, NIH) was used to control all camera parameters and acquire movies. Any movies that had significant drift (greater than 0.25 soma diameters), movement, or lacked significant amounts of activity were excluded from further analysis. Significant movement could be detected during independent components analysis (ICA) by the appearance of elliptical rather than circular segments. We observed that active, GCaMP-expressing neurons were found within a discrete layer (c.f. Fig. 3.1A) consistent with the location of layer 5 in medial prefrontal cortex.

Signal extraction

All analyses and signal extraction was performed using MATLAB (Mathworks). Locations of cells were automatically identified using a modified version of the published CellSort 1.1 toolbox (Mukamel et al., 2009b). In particular, a factor, μ , specifies the weight between spatial and temporal sparseness: $\mu = 0$ is purely spatial and $\mu = 1$ is purely temporal. All of

our analyses used $\mu = 0.2$ which was within the optimal range outlined in (Mukamel et al., 2009b). Signals were extracted from movies and the baseline fluorescence function, F_0 , was calculated for every trace using the mode of the kernel density estimate over a 100s rolling window, implemented via the MATLAB function `ksdensity` following the procedure outlined in (O'Connor et al., 2010). All signal traces shown represent normalized versions of the $\frac{F-F_0}{F_0}$ trace.

Threshold based event detection was performed on the traces by detecting increases in $\frac{F-F_0}{F_0}$ exceeding 2.5σ over one second, and then further thresholding these events by keeping only those events which exceeded a 4σ increase over two seconds. σ is the standard deviation of $\frac{F-F_0}{F_0}$, calculated over the entire movie. Thus all detected events have a deviation of at least 4σ from baseline. After identifying these events in the calcium signal from a cell, the cell was considered "active" during the entire period from the beginning to the peak of the event. The beginning of the event was defined as the first point for which $\frac{F-F_0}{F_0}$ increases by 2.5σ within 1 second and by 4σ within 2 seconds. The peak of the event was defined as the local maximum of the entire event, from the beginning of the event until $\frac{F-F_0}{F_0}$ returns to the same baseline value. We then created a matrix in which each row corresponds to a neuron, and each column corresponds to a frame. Entries in this matrix were 1 if a given neuron was active during a given frame, and 0 otherwise. All subsequent analyses were performed on this two-dimensional representation of network activity over time (c.f. Fig. 3.1C). Correlations between cells were calculated between the binary event trains corresponding to those two cells after subtracting the mean level of activity from each event train. The standard deviation

projection in Figure 1 was obtained as follows. For each pixel, we computed the standard deviation of $\frac{F-F_0}{F_0}$ over 30 second intervals throughout the first 10 minutes of a movie, then plotted the maximum value of these standard deviations.

Statistical analysis

Unless otherwise noted, we used the Mann Whitney U-test to compare pairs of groups, repeated measures ANOVA to compare multiple groups, and the two-tailed Kolmogorov-Smirnov (KS) test to compare pairs of distributions. To compare the number of strong correlations between conditions, we treated the fraction of strong correlations in each slice as an observation. Error bars where shown indicate standard error unless otherwise noted. Cell identification, signal extraction and normalization, event detection, and all other data analysis was done using fully automated routines that were independent of the investigator and thus blinded.

Matching levels of activity between datasets

To reduce activity in one group of datasets in order to match levels of activity in another group of datasets (c.f. Fig. S3.1C,D; Fig. S3.2C), we randomly deleted epochs of activity until the % time active in a given dataset matched the mean level of activity in the second group. For example, to generate Fig. S3.2C, we deleted randomly selected epochs of activity from saline-exposed datasets until they reached a mean level of activity of 1.4%, corresponding to the mean level of activity in VPA-exposed datasets. All data shown repre-

sents an average based on 5 instantiations of this epoch removal algorithm. The upsampling procedure performed in Figure S3.3 was achieved by overlaying multiple datasets to produce a single, virtually upsampled dataset. For example, we combined 2, 3, or 4 saline-exposed datasets or 3, 5, or 7 VPA-exposed datasets to achieve a higher-than-normal density of events, as indicated in the Fig. S3.3 legend. The correlation distributions represent the means, computed over all possible combinations of N datasets drawn from the total number of experiments for each condition. The datasets were combined by simply computing the union of the two datasets to produce a new dataset with a number of virtual neurons equal to the lesser of the numbers of cells in the two experiments. E.g. overlaying two VPA datasets, with 82 neurons and 84 neurons, yields a combined dataset with 82 virtual neurons. In the combined dataset, virtual neuron i is active at time t if neuron i is active at time t in either of the original datasets.

3.5 Figures

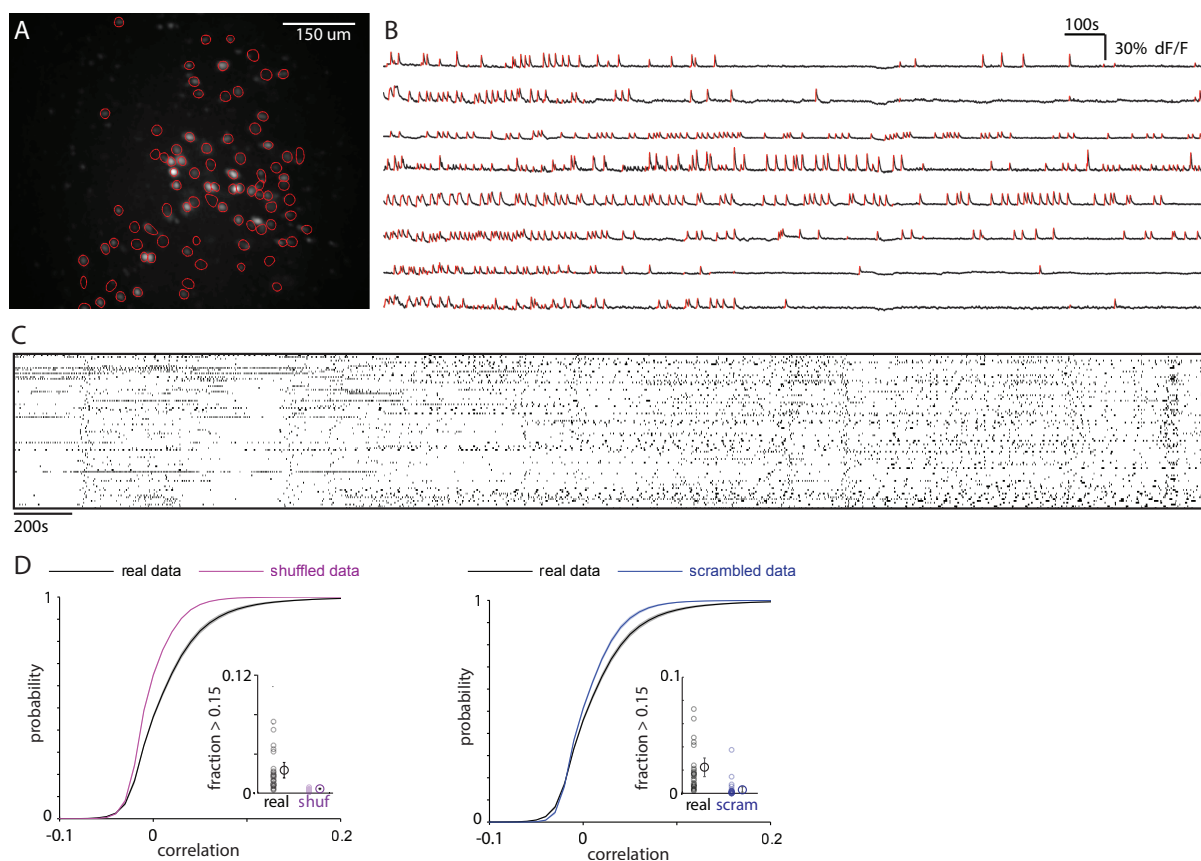


Figure 3.1: **Single photon GCaMP imaging resolves simultaneous activity from many neurons in prefrontal slices.**

a, Regions of Interest (ROIs) obtained by an automated algorithm, showing the locations of neurons, superimposed on the maximum standard deviation projection of a GCaMP6s movie. **b**, Sample GCaMP6s signals from 8 neurons (out of 84 total neurons imaged). Overlaid red lines indicate times when we detected that a neuron was active. **c**, Example raster of spontaneous network activity for a single 60 min experiment with 84 neurons. **d**, Cumulative probability distributions for correlations in real datasets (black line; $n = 29$ experiments from wild-type mice) vs. those observed in data shuffled by shifting each neuron's event train by a different random amount (shuffled, purple line) or by randomly reassigning the neuronal identity associated with each event (scrambled, blue line). Inset: the fraction of correlations ≥ 0.15 are shown for real and shuffled / scrambled datasets. In each case, both the individual datapoints as well as the means and standard errors are shown.

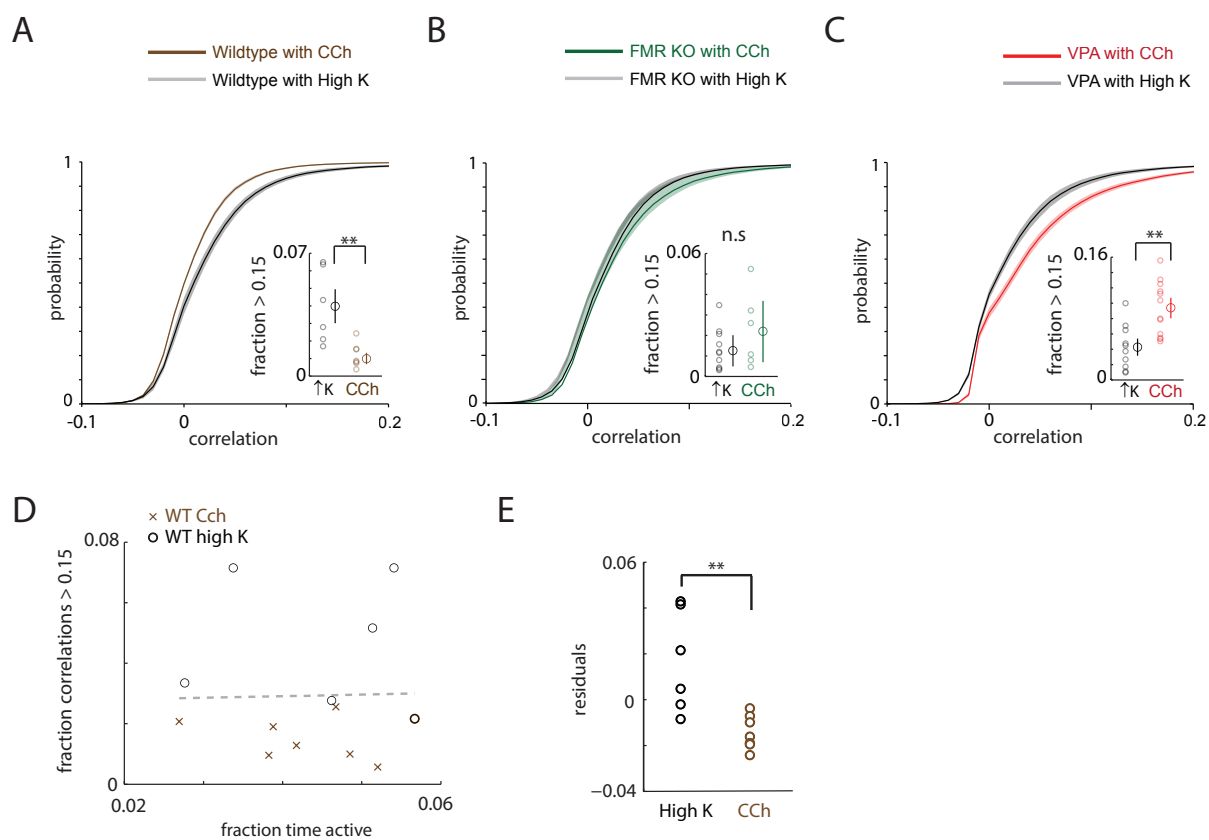


Figure 3.2: **Cholinergic modulation decorrelates microcircuit activity in wild-type mice, but not in models of autism**

a, Cumulative distribution of pairwise correlations for FMR1 WT datasets using either 2 μM carbachol ($n = 7$) or high K+ ACSF to elicit activity ($n = 6$). Inset: In carbachol, there are fewer strong correlations (>0.15) compared to high K+ ACSF ($p < 0.01$). **b**, Cumulative distribution plot for correlations recorded from FMR1 KO mice in carbachol (green line) or high K+ ACSF (black line). Inset: In carbachol, there is no reduction in strong correlations (>0.15), compared to high K+ ACSF. Both individual datapoints as well as the means and standard errors are shown. **c**, Cumulative distribution plot for correlations recorded from VPA-exposed mice in carbachol (red line) or high K+ ACSF (black line). Inset: In carbachol, there are more strong correlations (>0.15), compared to high K+ ACSF. Both individual datapoints as well as the means and standard errors are shown. **d**, Scatterplot of the fraction of strong correlations as a function of the fraction of time active for WT datasets in carbachol (black O symbols) or high K+ ACSF (brown X symbols). The gray dashed line represents a linear fit of all points (carbachol and high K+ ACSF). **e**, Residual values for the fractions of strong correlations, i.e. the difference between the actual values, and the fraction expected based on a linear relationship between activity and strong correlations, for carbachol (black) and high K+ ACSF (brown) datasets. Even after accounting for a possible relationship between activity and strong correlations, the fraction of strong correlations in high K+ ACSF is still significantly greater than in carbachol. ** $p < 0.01$ by Wilcoxon rank sum test

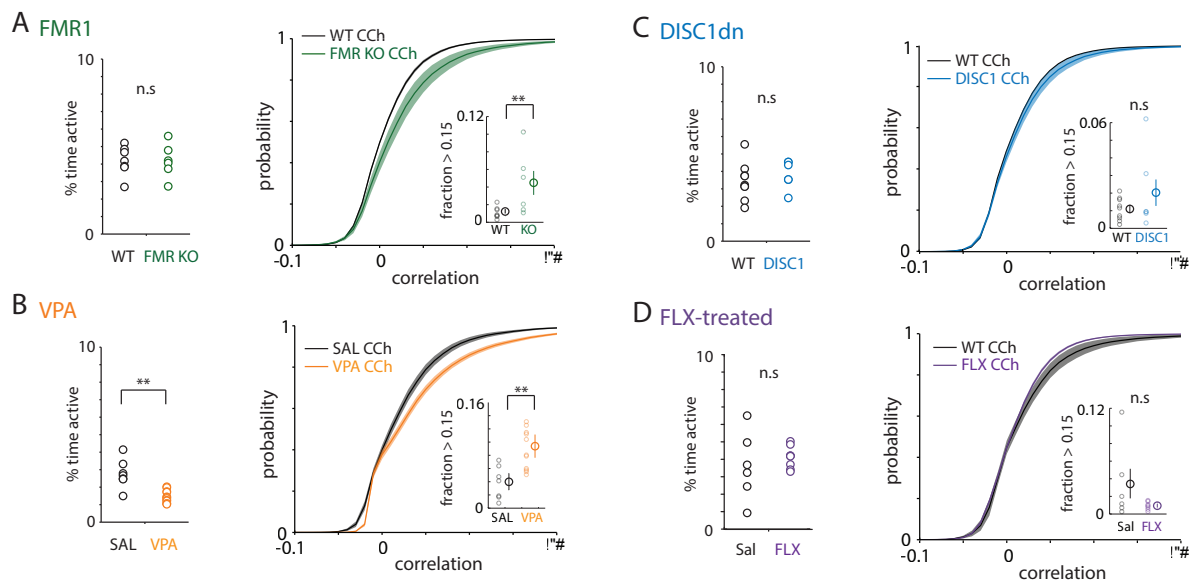
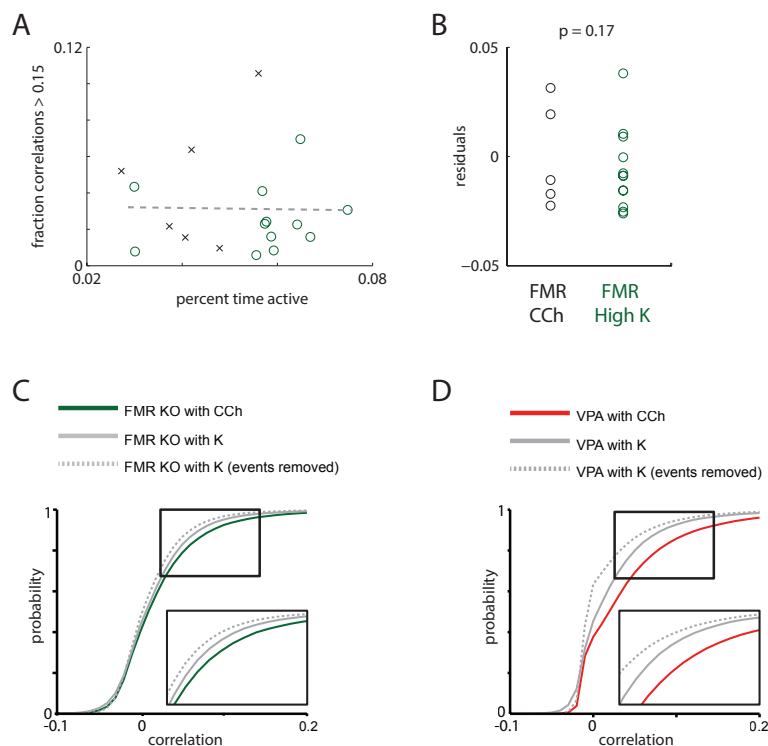


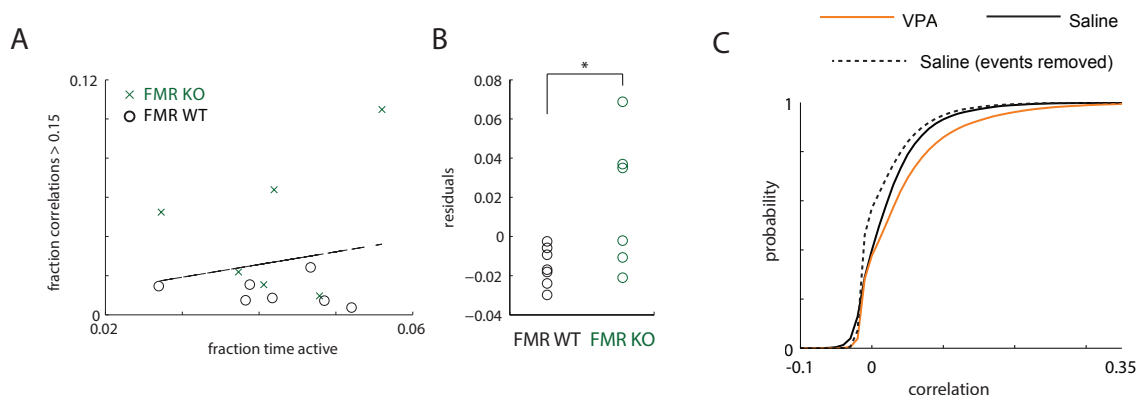
Figure 3.3: **Autism models, but not DISC1 mutant or fluoxetine-treated mice, exhibit abnormally elevated correlations in carbachol**

a, Left: Plot of the average amount of time each neuron was active for FMR1 WT (filled black circles, left; $n=7$) or FMR1 KO slices (filled green circles, right; $n=6$). **Right:** Cumulative distribution of correlations, averaged over all experiments in FMR1 WT (black; $n=7$) or FMR1 KO slices (green; $n=6$). Inset shows fraction of strong correlations, i.e. values ≥ 0.15 for each condition, and both individual datapoints as well as the means and standard errors are plotted. **b,** Similar to **a**, but compares experiments in saline-exposed (black; $n=9$) or VPA-exposed slices (orange; $n=12$). **c,** Similar to **a**, but compares experiments in control (black circles, left; $n=8$) vs. dominant negative DISC1 mutant mice (blue circles, right; $n=6$). **d,** Similar to **a**, but compares experiments in saline-treated (black circles, left; $n=6$) compared to FLX-treated mice (purple circles, right; $n=7$). ** $p < 0.01$ by Wilcoxon ranksum test.




Supplementary Figure 3.1: Differences in activity levels do not explain the inability of cholinergic modulation to decorrelate activity in autism models

a, Scatterplot showing the fraction of strong correlations as a function of the mean % time active in FMR1 KO mice in carbachol (black) or high K+ ACSF (green). The gray dashed line represents a linear fit of a linear model to all points (carbachol + high K+ ACSF). **b**, Residual values for the fraction of strong correlations in carbachol (black) and high K+ ACSF (green), i.e. the difference between the actual values and the number expected based on a linear relationship between the level of activity and the prevalence of strong correlations. Even after accounting a possible relationship between activity and correlations, there is no carbachol-induced decorrelation in FMR1 KO circuits. **c**, Cumulative distribution plot for correlations recorded from FMR1 KO mice in carbachol (green line) or high K+ ACSF (gray line). The cumulative distribution plot is also shown for datasets recorded in high K+ ACSF that have had epochs of activity removed to match activity levels observed in carbachol (dotted gray line). These plots show that removing epochs of activity tends to decrease, rather than increase, correlations, suggesting that abnormally high correlations in FMR1 KO mice in carbachol are not an artifact of decreased activity levels. **d**, An analysis based on linear regression and residual values was not possible for the VPA datasets, because activity levels were largely non-overlapping for the carbachol and high K+ conditions, making it impossible for linear regression to dissociate activity-driven differences in correlations from those driven by condition (carbachol vs. high K+). Instead, we again removed epochs of activity from VPA-exposed datasets in high K+ ACSF to match the levels of activity in carbachol. Cumulative distribution plots are shown for correlations recorded from VPA-exposed mice in carbachol (red line), high K+ ACSF (gray line), and high K+ datasets that have had epochs of activity removed to match activity levels observed in carbachol (dotted gray line). Our finding that VPA-exposed circuits exhibit a carbachol-induced increase (rather than decrease) in strong correlations cannot be explained simply by the fact that activity levels are lower in carbachol than in high K+, since lower activity levels tend to decrease, rather than increase, correlations.



Supplementary Figure 3.2: **The abnormally increased correlations observed in two models, relative to controls, in carbachol, are not an artifact of differences in activity levels**

a, Scatterplot showing the fraction of strong correlations as a function of the mean % time active for FMR1 WT mice in carbachol (black O symbols) or FMR1 KO mice in carbachol (green X symbols). The gray dashed line represents a linear fit of all points. **b**, Mean activity levels were almost identical for FMR1 WT (black) and FMR1 KO (green) datasets in carbachol. Nevertheless, we analyzed residual values based on the linear regression shown in **a** to confirm that the increased correlations in FMR1 KO mice in carbachol were not an artifact of differences in activity. We plotted residual values for the fraction of strong correlations in FMR1 WT (black) and FMR1 KO (green) datasets in carbachol, i.e. the differences between the actual values and the values expected based on a linear relationship between the fraction of strong correlations and mean % time active. Even after accounting for a possible relationship between correlations and activity, residual correlations in FMR1 KO mice are still significantly greater than those in FMR1 WT mice. **c**, Again, an analysis based on linear regression and residual values was not possible for the VPA-exposed datasets and saline-exposed controls, because activity levels were largely non-overlapping for these two conditions, making it impossible for linear regression to dissociate activity-driven differences in correlations from those driven by condition (VPA-exposed vs. saline-exposed). Instead, we created surrogate datasets to artificially match levels of activity between these two conditions by either removing epochs of activity (here), or by overlaying multiple datasets to increase levels of activity (in Fig. 3.3). Here we have plotted the cumulative distributions for correlations recorded in carbachol from VPA-exposed (red line) or saline-exposed mice (gray line). Reducing activity in control datasets, by removing epochs of activity to match the activity levels observed in VPA-exposed datasets (dotted gray line), tends to decrease correlations and exacerbate the difference between these two populations. Thus, the increased strong correlations in VPA-exposed datasets, compared to controls, does not appear to be an artifact of differences in activity levels. * $p < 0.01$ by Wilcoxon rank sum test



Supplementary Figure 3.3: **VPA-exposed datasets continue to exhibit an increase in strong correlations, even after matching the higher levels of activity in saline-exposed datasets.**

Generating surrogate datasets by overlying multiple actual datasets to artificially increase levels of activity (described in the Methods), shows that regardless of the level of activity, correlations are consistently increased in VPA-exposed datasets in carbachol (orange) compared to saline-exposed controls in carbachol (black). The legend shows the mean % time active for each surrogate dataset generated by overlaying multiple VPA-exposed datasets or saline-exposed controls. This provides additional evidence against the idea that the increased correlations observed in VPA-exposed datasets are simply an artifact of differences in activity levels.

Chapter 4

Correlations between prefrontal neurons form a small world network that optimizes the generation of multineuron sequences of activity

4.1 Introduction

Many *in vivo* studies have identified sequential patterns of activity within the prefrontal cortex and other associational cortices during working memory (Baeg et al., 2003; Fujisawa et al., 2008; Harvey et al., 2012; Seidemann et al., 1996). The visual cortex generates similar patterns both in response to visual stimulation as well as spontaneously (Luis et al., 2015).

However, it remains unclear how such types of patterns are generated. Specifically, do these sequences depend on a specific functional organization within cortical microcircuits, or emerge simply as a byproduct of random interactions between neurons? Many studies have used theory and simulations to show how neural networks resembling local neocortical circuits can generate the sorts of stereotyped trajectories or sequential patterns of activity believed to mediate a variety of functions including motor planning and working memory (Ashok and Doiron, 2012; Ganguli et al., 2008; Goldman, 2009; Hennequin et al., 2014; Laje and Buonomano, 2013; Sussillo and Abbott, 2009). Nevertheless, it has not yet been possible to demonstrate that the functional organization of an actual neocortical microcircuit contributes to its ability to generate meaningful patterns of multineuron activity.

Recent advances in recording techniques have made it possible to collect large datasets and make important observations, e.g., neurons within isolated cortical networks exhibit correlations that follow a small world organization and can fire in a temporally precise manner (Sadovsky and MacLean, 2013). However a fundamental unsolved challenge is developing analytic methods, beyond simulations, that can directly assess how the functional organization of a circuit contributes to the production of emergent patterns of activity. Notably, one study used sophisticated imaging methods to show that isolated neocortical microcircuits in brain slices spontaneously generate sequential patterns of activity (Ikegaya et al., 2004), however, later studies showed that even networks of independently firing neurons or with randomly arranged correlations between neurons could generate similar levels of sequences (Roxin et al., 2008). Most studies to date have compared emergent patterns observed exper-

imentally to those expected based on various models using simulations (Roxin et al., 2008; Sadvovsky and MacLean, 2013). Using simulations in this way is complicated, because unlike shuffling experimental data, it requires making assumptions so that simulations will approximately, though not exactly, match experimentally observed levels and temporal patterns of activity. Conversely, altering the circuit organization in simulations will invariably alter levels and temporal patterns of activity, making it impossible to disambiguate the individual contributions of these three factors to the observed patterns of activity.

Maximum entropy models have been used to directly assess how the organization of a circuit influences its ability to generate patterns of activity. These models predict instantaneous patterns of network activity by fitting the mean firing rate for each neuron, as well as pairwise functional interactions between neurons, without making additional assumptions (Schneidman et al., 2006). Maximum entropy models can reproduce multineuron patterns of activity within the isolated retina and cultures of dissociated cortical neurons, and have shown that in these systems, emergent patterns of multineuron activity do indeed reflect pairwise functional interactions between neurons. However, these models have been less successful for fitting multineuron patterns of activity from cortical networks in brain slices and in vivo (Ohiorhenuan et al., 2010). In particular, maximum entropy models do not model sequential patterns of activity observed in cortical networks (Ohiorhenuan et al., 2010), because they predict the distribution of instantaneous patterns of activity. This underscores the need for new methods to assess how sequential patterns of activity in experimental datasets depend on the underlying functional organization of a circuit.

Here, we outline one such method, which can dissociate the relative contributions of activity levels, temporal patterns of activity, pairwise correlations, and higher order structure such as clustering, to key features of emergent network-level behavior. We use this method, together with new experimental approaches, to show that the small world functional organization of local prefrontal microcircuits optimizes the production of stereotyped, sequential patterns of activity believed to mediate important cognitive functions in vivo. This reveals a specific, concrete, and important function for small world networks, which are found ubiquitously in the nervous system and throughout nature, and demonstrates a new method for addressing the critical question of whether knowing the functional organization underlying a neuronal system helps to understand how it works (Marder, 2015).

4.2 Results

We initially used single-photon, wide field imaging to capture fluorescent signals from the genetically encoded Ca^{2+} indicator GCaMP6s (Chen et al., 2013) in acute brain slices from late adolescent (P41-57) mice ($n = 29$). We recorded sparse, robust GCaMP signals from neurons in deep layers of the medial prefrontal cortex (mPFC; Fig. 4.1). We image GCaMP signals at 10 Hz. Thus, while GCaMP sensors are in principle capable of resolving single spikes (Chen et al., 2013), the events detected here reflect increases in spike rate rather than single spikes. We used a low concentration of the cholinergic agonist carbachol (2 μM) in the bath to model basal cholinergic tone in vivo and promote spontaneous network activity (Fellous and Sejnowski, 2000). Using a combination of independent component

analysis (ICA) and image segmentation (Mukamel et al., 2009a), we identified the locations of neurons, measured their GCaMP signals, and detected events corresponding to increased activity in these neurons (Fig. 4.1). Each individual experiment recorded from 58-94 active neurons over 1 hour. Some of these represent control experiments from a recent study of neuronal correlations in mice which model autism (Luongo et al., 2015), however that study only looked at pairwise correlations and did not examine the sequential patterns of activity which are the subject of this study.

Prefrontal microcircuit activity contains more correlations than expected by chance

First, we determined whether prefrontal microcircuits generate spontaneous activity that is randomly distributed versus organized according to some structure, by computing correlations between activity in different neurons. We found that the distribution of correlations included more strong correlations than would be expected by chance (Fig. 4.2A). We compared these correlations to those obtained after shuffling our data. Notably, neurons exhibit long-lasting periods of increased activity interspersed with periods of quiescence. To preserve the autocorrelated temporal structure in shuffled data, we first simply shifted large chunks of each neurons event train in time (Methods). Shuffling data in this way preserves the ISI distribution of each individual neuron (autocorrelation) but should destroy any structure between neurons. As shown in Fig 4.2A, we found that real data contained an excess of strong correlations between different neurons compared to this shuffled data.

To further disentangle the effect of this temporal structure, we also compared the correlations observed to those found in a scrambled version of each dataset. Unlike the shuffled datasets, these scrambled datasets account for any network-wide non-stationarity that may be present. These represent an important control as correlation measures have been shown to be particularly sensitive to violations of stationarity (Brody, 1999).

To understand how we generated these scrambled datasets, first note that each original dataset is composed of a set of epochs of activity, where each epoch is defined by which neuron was active, when it became active, and how long it remained active. To scramble a dataset, we simply shuffled the list of neuron identities associated with each epoch. As a result, each neuron has exactly the same number of epochs of activity in the original and scrambled datasets, i.e. the scrambled dataset is effectively rate-matched. These scrambled datasets, which maintain both the temporal structure of total network activity and the distribution of neuronal rates, nevertheless contain significantly fewer correlations than actual datasets (Fig. 4.2B). Notably, in actual datasets, strong correlations were present even for pairs of neurons separated by large distances (Fig. 4.2C).

Prefrontal microcircuit activity has a small world organization

Next, to analyze the organization of activity in prefrontal microcircuits we asked are correlations between neurons arranged randomly or do they follow a specific pattern? To address this question, we examined the structure of a network in which each neuron represents a node, and the strong correlations between neurons represent edges. We assumed an edge is

present between two neurons if the correlation between them exceeds the noise level, and absent otherwise. Networks constructed in this way, using the correlations between neurons, describe the functional organization of a neural circuit, rather than the organization of physical, synaptic connections. Such functional networks have been widely studied and shown to be powerful tools for exploring the emergent behavior of neuronal circuits (Abeles et al., 1993; Diesmann et al., 1999; Sadvovsky and MacLean, 2013; Schneidman et al., 2006; Watts and Strogatz, 1998) and in many cases reflect important aspects of physical connectivity (Aviel et al., 2003; Bassett and Bullmore, 2006; Sporns and Zwi, 2004).

We next calculated clustering coefficients and path lengths for each functional network. These measures are used to distinguish between randomly wired networks, small world networks, and clustered networks (Yu et al., 2008). These three classes of networks can all have the same number of edges (strong correlations), but differ in how those edges are arranged. In randomly wired networks, edges (correlations) between neurons are distributed randomly, whereas in clustered networks, edges are distributed topographically, such that each neuron is connected (strongly correlated) with other nearby neurons, which also connect to each other, forming clusters. Small world networks are formed when a small fraction of the edges in a clustered network are randomly rewired, and are thought to exhibit desirable features of both clustered and random networks. For example, clustered, but not random networks, have distributed structure, which can be useful, e.g. for generating stereotyped patterns of activity as we will show. By contrast, random networks have a short path length, i.e. the average distance (along edges) from one node to another is relatively short, facilitating

information flow through the network. Real world networks, e.g. social networks, often possess a small world organization which maintains some structure while also having short path lengths typical of random networks.

Clustering coefficients are defined as follows. Suppose node A is connected to nodes B and C then the clustering coefficient is just the probability that B and C are themselves connected. We find that our datasets all contain (1) clustering coefficients that are markedly higher than those found in random networks (Fig. 4.2E; actual clustering coefficients: 0.30 ± 0.02 vs. 0.07 ± 0.01 for random networks, $p < 0.001$), and (2) low path lengths, similar to those in random networks (Fig. 4.2D; actual path lengths: 3.0 ± 0.2 vs. 3.1 ± 0.4 for random networks, $p=0.78$). These results do not change if we vary the threshold for detecting events within Ca^{2+} signals (Fig. S4.1), or if we determine the threshold for defining strong correlations on an individual cell basis, instead of using a single threshold for the entire network (Fig. S4.2).

Our datasets also contain substantially more clustering than is expected simply based on a random, distance-dependent rule for coupling (Fig. S4.3). Finally, both path length and clustering were independent of network size (Fig. S4.4). Thus, within prefrontal microcircuits, correlations are arranged non-randomly, forming small world networks.

Prefrontal microcircuits spontaneously generate sequential patterns of activity. Given the bevy of modeling studies that have described ways in which the functional organization of recurrent networks can shape patterns of emergent activity, we decided to measure the capacity of prefrontal microcircuits to generate sequential patterns of activity (Abeles et al.,

1993; Diesmann et al., 1999). We used a template matching algorithm similar to the one used previously (Ikegaya et al., 2004) to count the number of occurrences of each potential sequence of activity. As diagrammed in Figure 4.3A, this algorithm identifies every pattern of activation within a 5 second window following the onset of activity in a neuron, counts the number of times the same pattern recurs (allowing for a jitter of ± 1 frames, i.e. 0.1 sec), and counting all sub-patterns within the original pattern. A 5 second window was chosen for consistency with Ikegaya et al. 2004, however most patterns observed had much shorter, more biologically plausible durations, less than 0.5 or 1 second (Fig. S4.5). In order to determine whether these datasets contain more sequential patterns of activity than expected by chance, we initially made comparisons to the two types of control datasets described earlier. First, we calculated the number of sequential patterns in datasets that had been shuffled by shifting each neurons activity train in time. Second, we calculated the number of sequential patterns in a scrambled version of each dataset, in which the cell identities corresponding to each epoch of activity had been randomly reassigned, while maintaining the temporal pattern of total network activity and the distribution of activity levels across cells. Every actual dataset contained more sequential patterns of activity than the corresponding shuffled (Fig. 4.3B, D) or scrambled datasets (Fig. 4.3C, D). Thus, actual datasets generate more stereotyped sequential patterns of activity (i.e., sequences that were observed multiple times) than expected simply by chance.

Understanding how the small world organization of prefrontal microcircuits contributes to sequence generation

Of course, the critical question we set out to address is whether the functional organization of prefrontal microcircuits facilitates sequence generation. As described in the Introduction, previous studies have, in most cases, failed to observe sequential patterns of microcircuit activity exceeding the number expected based on randomly arranged correlations between neurons. Naively, one would expect the small world organization we found to facilitate sequence generation, because the presence of clustering in a network implies that specific groups of neurons are repeatedly co-active, which should lead to the repeated occurrence of sequences composed of these neurons. Thus, we would first like to confirm this intuition, that small world networks generate more sequences than networks in which correlations are randomly arranged.

We would also like to address two additional questions whose answers are less obvious. First, does increasing clustering beyond the levels observed in actual datasets further increase sequence generation, i.e. would maximally clustered networks, in which correlations are arranged topographically, generate even more sequential patterns of activity than small world networks? The previously laid out intuition would suggest that more clustered networks should generate more sequences. On the other hand, stronger clustering implies that once one or two cells are known to be co-active, the remaining co-active cells should become more predictable. This might reduce the number of unique sequences that occur in more clustered networks.

A second question is whether the temporal pattern of total network activity (i.e. the number of neurons active as a function of time) together with the functional organization of the network, i.e. the (undirected) pairwise correlation matrix, are sufficient to explain the number of sequential patterns of activity we observed? I.e. does a model which includes the temporal pattern of total network activity and the pairwise correlation matrix, but makes no additional assumptions, reproduce the levels of sequences observed in actual data? Or do additional aspects of network activity, e.g. the detailed statistics of activity within individual neurons or knowledge about whether one neuron tends to be active before vs. after another, contribute to sequence generation?

To address the preceding questions, we would ideally be able to compare patterns of network activity that are similar overall, but reflect specific changes in the arrangement of correlations. Therefore, we devised a novel method to create surrogate datasets that all have exactly the same amount and temporal pattern of activity as our original dataset; i.e. if N neurons were active at time t in the original dataset, then N neurons would be active at time t in each surrogate dataset as well. However, these surrogate datasets either preserve or alter the correlations between neurons, and can thereby evaluate how specific changes in the organization of pairwise correlations would affect the occurrence of sequential patterns of activity. Note that maximum entropy models can be used in a similar way, to test whether knowing the mean activity level for each neuron and the pairwise functional interactions between neurons is sufficient to predict patterns of network activity. However, as noted earlier, maximum entropy models only predict instantaneous patterns of activity,

whereas we are interested in sequential patterns of network activity. Therefore, our method (described below) generates surrogate datasets that capture the evolution of network activity over time.

To understand how we generated these surrogate datasets, again consider that each original dataset can be defined as a set of epochs of activity, where each epoch is defined by which neuron was active, when it became active, and how long it remained active. Each dataset is also associated with a set of correlations. As described below, to generate each surrogate dataset, we simply re-assigned the neuron associated with each epoch of activity, in order to achieve a new, target set of correlations (this can be thought of as non-randomly shuffling the cell identities). In this way, each surrogate datasets preserved the overall temporal structure of network activity because the timing and duration of every epoch of activity was the same in the original dataset and in all the surrogate datasets. Thus, the total number of neurons active at any given point in time would be the same for all of these datasets. However, the particular combination of cells active at a given point in time would differ across datasets, and as a result, each surrogate dataset also had a different set of correlations.

The target correlation matrix could be chosen to preserve the pattern of strong correlations present in our original datasets. Alternatively, the target correlation matrix could randomly rearrange these correlations, preserving the number of strong correlations but disrupting their organization. The iterative optimization procedure by which we reassigned the cell identities in order to achieve the target correlation matrix is diagrammed in Fig. 4.4A. First, we randomly selected an epoch of activity, i , whose cell identity had not yet

been re-assigned. Second, we found all the other epochs of activity j which (1) overlap in time with epoch i and (2) already have had a new cell identity assigned. Third, for each cell assigned to an epoch j which overlaps with epoch i , we computed the difference between the target correlation vector corresponding to that cell, and the current value of that correlation vector based on the partially constructed surrogate dataset. (This step can be thought of as guessing which cell should be assigned to a particular epoch of activity by first figuring out what other cells are active at the same time, then choosing cells which are strongly correlated with these known active cells). Next, we multiplied each difference vector by the amount of overlap between epoch j and epoch i and computed a weighted sum of all the difference vectors. (This step can be thought of as tallying up all of the guesses about which cells should be assigned to the epoch of activity being considered). Finally, we selected a new cell identity for epoch i by simply choosing the cell corresponding to the maximum value of this weighted difference vector, excluding all cells that were already active during any part of epoch i (i.e. we chose the cell that represents the consensus based on tallying up all of the guesses). If the difference vector was empty, e.g. because there were no overlapping epochs of activity that have had new cell identities assigned, then we simply used the cell identity of this epoch of activity from the original dataset.

Of course, although the target matrices were used to generate the surrogate datasets, the resulting surrogate datasets had slightly different correlation matrices (this reflects both the stochastic nature of the algorithm and the constraint that the temporal pattern of activity is fixed). Examples of correlation matrices for surrogate datasets are shown alongside their

target correlation matrices in Fig. 4.4B. As Fig. 4.4B shows, the surrogate datasets had the desired properties, i.e., correlation matrices for the surrogate datasets were very similar to the target matrices. For example, the correlation matrices for surrogate datasets designed to preserve the arrangement of correlations present in the original datasets were very similar to the original correlation matrices, as evidenced by a normalized dot product of 0.96 ± 0.02 (mean \pm std dev). This original arrangement of correlations was lost in the surrogate datasets generated from correlation matrices that had been randomly rearranged, as the normalized dot product in this case was markedly lower: 0.21 ± 0.14 ($p < 0.001$). Further examples of activity rasters and correlation matrices are shown in supplementary Figures S4.6 and S4.7. Notably, for the randomly rearranged case, the surrogate datasets did have correlation matrices very similar to their (randomly rearranged) target matrices, as evidenced by a normalized dot product of 0.95 ± 0.01 . Thus, our method creates surrogate datasets that conform to the training correlation matrices.

Small world networks generate more sequences than either random or clustered networks

For each of our 29 datasets, we used the method described above to generate surrogate datasets which either preserved the small world organization of pairwise correlations, randomly rearranged these correlations, or rearranged these correlations topographically, resulting in a maximally clustered network (c.f. Fig. 4.4B). We confirmed that the network properties of each surrogate dataset matched the respective network type—clustered, ran-

dom, or small world (Fig. S4.9). Then we counted the number of sequential patterns of activity (as a function of sequence length) for each actual dataset, shuffled dataset, scrambled dataset, surrogate dataset based on randomly rewiring the correlation matrix, surrogate dataset which preserved the original small world correlation matrix, and surrogate dataset based on a maximally clustered correlation matrix (Fig. 4.5A). Fig. 4.5B summarizes the total number of sequential patterns of activity for each case. As shown in Fig. 4.5, surrogate datasets that preserve the original small world pattern of correlations contain the same number of sequences as our original datasets, and generate significantly more sequential patterns of activity than either random or maximally clustered networks. These findings directly address the two questions raised earlier. First, a model which preserves the overall pattern of network activity (i.e. the timing and duration of each epoch) and the pairwise correlation matrix, but makes no additional assumptions is sufficient to reproduce the levels of sequential patterns of activity observed in actual data. Second, the small world organization of pairwise correlations we found in prefrontal microcircuits optimizes the production of these sequential patterns.

An important question is whether our method simply renames the cells. I.e., perhaps small world surrogate datasets match the performance of the actual datasets because each neuron in a surrogate dataset effectively mimics the activity pattern of a neuron in the original dataset? To address this question, we calculated the correlation between the event train of each surrogate dataset neuron and the most similar event train in the original dataset. The correlation was 0.2 for random/clustered surrogate datasets, and 0.3 for small world

surrogate datasets (Fig. 4.8A). Thus, cells in the small world surrogate datasets are not simply renamed versions of cells in the original datasets. Furthermore, this slight increase in correlation for small world surrogate datasets relative to random/clustered surrogate datasets could be eliminated simply by generating surrogate datasets using an activity pattern from one experiment together with a correlation matrix from a different experiment. In this case, the correlation between each surrogate neuron and the most similar neuron in each original dataset was 0.1 for random, clustered, and small world surrogate datasets (Fig. 4.8B). Furthermore, even when we generated surrogate datasets in this way, using the activity pattern from one experiment together with the correlation matrix from a different experiment, small world surrogate datasets continued to reproduce the number of sequences observed in actual datasets, and to outperform random or clustered datasets (Fig. 4.10).

Having shown that small world networks generate more sequential patterns of activity than random or clustered ones, we next explored sequence generation for surrogate datasets with intermediate degrees of small-worldness. We created these networks by starting with the original (small-world) or maximally clustered correlation matrices, then randomly rearranging 10% of the strong correlations. We then quantified the small-worldness of each of these two types of intermediate networks, as well as our three original types of surrogate datasets (random, maximally clustered, or small-world), using the widely used metric

$$sw = \frac{C}{C_{er}} \frac{L_{er}}{L}$$

, where C and L are the clustering coefficient and path length of the network, and C_{er} and L_{er} are the average clustering coefficient and path length for erdos-renyi random networks

(Humphries and Gurney, 2008).

We observed a strong correlation between the number of sequential patterns of activity generated by a network and its small-worldness ($p < 0.001$; Fig. S4.11). Finally, although our scrambled datasets represent random datasets that match both the distribution of activity levels across cells and the temporal distribution of overall activity levels, we also explored sequence generation within networks composed of rate-matched homogeneous or inhomogeneous Poisson neurons. Networks composed of Poisson neurons lack the large numbers of short sequences (3 neurons) that recur many times (>10 -15 times / hour) in actual datasets and small-world surrogate datasets (Fig. S4.12). They also lack the intermediate-sized sequences (4-6 neurons) which recur >5 times / hour in actual datasets and small-world surrogate datasets. By contrast, networks of Poisson neurons generate extremely large sequences (>10 neurons) which repeat a small number of times but contain numerous subsequences. Such large sequences are never observed in actual datasets or small-world surrogate datasets.

4.3 Discussion

This study evaluated how the functional organization of prefrontal microcircuits, i.e. the small world correlation matrix, contributes to an important function of these circuits: their ability to generate multineuron sequences of activity. We found that pairwise correlations between prefrontal neurons form a small world network that optimizes the production of a diverse set of such sequences. Conversely, this small world organization is sufficient to

explain the occurrence of sequential patterns of activity in these circuits. We also outline a general method for generating surrogate datasets that makes minimal assumptions about the structure of the underlying data, maintains the temporal structure of total network activity, and either preserves the pattern of pairwise correlations or alters it in specific ways. This method will complement other approaches, e.g. theory, simulations, and maximum entropy models, for linking the functional organization of a circuit to the emergence of complex multineuronal patterns of activity.

Small world neural networks

Small world networks were first formalized to describe diverse real world networks that combine desirable features of both highly structured and randomly connected networks clustering and short path lengths, respectively (Watts and Strogatz, 1998). These features have been hypothesized to reflect an ability of small world networks to process information on both local and global scales. Small world topologies have been described for inter- and intra-areal functional connectivity in cat, monkey, and human cortex (Bassett and Bullmore, 2006; Sporns and Zwi, 2004). Small-world topologies have also been identified for microcircuits in neuronal culture (Downes et al., 2012), rat and cat visual cortex (Yu et al., 2008), deep layers of monkey cortex (Gerhard et al., 2011), and various cortical microcircuits of mice (Sadovsky and MacLean, 2013). Introducing higher order interactions such as clustering into a model of otherwise randomly connected cortical neurons nicely recapitulates the doubly stochastic nature of cortical spike trains, further supporting the presence of small world

topologies within cortical microcircuits (Ashok and Doiron, 2012).

Our results not only support the idea that small world networks represent a common organizing principle for neural systems, but more importantly demonstrate a specific and concrete way in which such an organization actually impacts neural circuit function. Small world networks have been hypothesized to effectively balance local and global modes of computation, but until now, applications of these ideas to neural networks have been largely abstract. Here we show exactly how this balance enables prefrontal microcircuits to generate stereotyped temporal patterns of activity, while also maximizing the diversity of such patterns. This solves a challenging tradeoff between the tendencies of random networks to produce multitudes of patterns that fail to repeat, and highly clustered networks to produce a small number of extremely repetitive patterns. Neural circuits that need to generate these patterns in order to reliably and flexibly support a diverse repertoire of behaviors must balance these two regimes, and here we show that the small world functional organization of actual prefrontal microcircuits achieves this balance. In this way, prefrontal microcircuits can contribute to the stereotyped patterns of activity that have been associated with flexible behavior in vivo (Baeg et al., 2003; Fujisawa et al., 2008; Harvey et al., 2012; Seidemann et al., 1996).

Remarkably, surrogate datasets based on a small world network of pairwise correlations between neurons and the temporal pattern of total network activity we observed are sufficient to largely reproduce the levels of sequential patterns observed in actual recordings. Naively, one might have expected undirected pairwise correlations to be extremely limited

in their ability to capture temporally structured activity, and that adequately describing sequential patterns of activity would require knowledge about the underlying, asymmetric pattern of neural connections, or at least which cell in a correlated pair tends to be active first. However, we find that at the level of cortical microcircuits, the undirected pairwise correlation matrix, with its small world topology, captures the key features of network organization involved in sequence generation. This is reminiscent of other studies that have used coupling between neurons to explain complex features of activity in actual neural circuits (Stevenson et al., 2012). Interesting, a distance-dependent, random coupling rule was not sufficient to reproduce the level of small-worldness we found in prefrontal microcircuits. Consistent with previous studies of cultured neurons (Downes et al., 2012), this implies that additional mechanisms, besides a local random connection rule, give rise to the small world organization present in actual neural networks.

Relationship to in vivo studies

Of course, it will be important to compare our findings to features of activity occurring in vivo. That being said, in vivo recordings would address fundamentally different questions from those addressed here. Activity in vivo is strongly influenced by synaptic inputs originating outside the PFC. Thus, if we found that correlations follow a small world organization in vivo, it would be unclear whether this reflects the organization of the local circuits vs. the pattern of incoming input. Furthermore, one major goal of our study was to determine whether prefrontal microcircuits are themselves sufficient to generate sequential patterns of

activity at meaningful levels. This question, which must be answered to understand whether these sequences emerge at the microcircuit level vs. as a result of functional interactions between brain regions, can only be addressed by studying isolated microcircuits, e.g. in brain slices. Two other recent studies similarly used brain slices to establish the presence of small world topologies and temporally precise firing in cortical microcircuits (Sadovsky and MacLean, 2013).

The importance of higher order interactions in modeling neuronal ensembles

Unsurprisingly, neural representations involving populations of cortical neurons are more robust and carry more information than individual spike trains (Averbeck et al., 2006). While most studies have focused on the information gained by population averaging to remove noise, recent studies suggest that some information may be encoded by the actual ensembles of neurons that are activated (Fujisawa et al., 2008; Harvey et al., 2012; Luis et al., 2015). However, understanding the factors that lead to the emergence of neuronal ensembles requires methods that are able to determine whether specific factors, e.g. the organization of pairwise correlations between neurons, are necessary and/or sufficient to explain the observed levels and complexity of these ensembles. Maximum entropy models have proven extremely useful for addressing this question in the context of instantaneous patterns of multineuron activity observed within the retina or cultures of dissociated cortical neurons (Schneidman et al., 2006), but do not always reproduce sequential patterns of activity in cortical networks.

Furthermore, we explored how hypothetical forms of network organization, e.g. random or maximally clustered networks, would impact neuronal ensembles. Similar analyses would be challenging with widely used versions of maximum entropy models, in which the relationship between pairwise functional interactions and network activity is highly nonlinear. As a result, rearranging pairwise functional interactions to model different network topologies would likely alter overall levels of activity, confounding any attempt to dissociate changes in ensembles due to alterations in network organization from those due to altered activity levels.

As more studies record patterns of activity from large populations of neurons, either in vivo or in reduced preparations, a critical question will be whether those patterns are simply a byproduct of randomly arranged correlations between neurons, or whether they depend on more complex patterns of functional interactions. A closely related question is whether knowing the functional organization of a network, helps to understand how that circuit actually works (Marder, 2015). The computational method presented here represents a powerful new tool for answering both of these questions. This method reassigns the cell identities associated with activity in the original dataset, rather than simulating new patterns of activity. This is critical, because this method preserves the exact pattern of total network activity present in the original dataset, making it possible to directly probe and dissociate the relative contribution of activity levels, temporal patterns of activity, pairwise correlations, and higher order structure such as clustering, to key features of emergent network-level behavior.

Conclusions

Understanding the relationship between organization and function in neural networks is a central challenge of modern neuroscience, but until now has been addressed largely using theory and simulations. Here we investigate this relationship in actual cortical microcircuits and find that correlations between prefrontal neurons form a small world network that optimizes the ability of prefrontal microcircuits to generate sequential patterns of activity which have been linked to important behavioral functions. This reveals an important function for small world neural networks, and illustrates a novel approach for linking the functional organization of a neural network to the patterns of activity it generates.

4.4 Material and Methods

All experiments were conducted in accordance with procedures established by the Administrative Panels on Laboratory Animal Care at the University of California, San Francisco.

Slice preparation

Wild-type P26 P33 C57BL/6 mice (Charles River) were injected unilaterally with 1 μ l of AAV5/2-CaMKII::GCaMP3 (UNC) at the coordinates (in mm) 1.7 anterior-posterior (AP), 0.3 mediolateral (ML), and -2.2 dorsoventral (DV). 350 micron thick coronal slices were prepared from these animals 15-27 days after injection. Slices preparation followed our previously described protocol 18. The one notable deviation from the published protocol was

that immediately after brain slices were prepared, they were transferred to an N-Methyl-D-Glucamine (NMDG)-based recovery solution for 10 min before being transferred to ACSF for the remainder of their recovery. The NMDG-based solution was maintained at 32° C, and consisted of the following (in mM): 93 N-Methyl-D-Glucamine (NMDG), 93 HCl, 2.5 KCl, 1.2 NaH₂PO₄, 30 NaHCO₃, 25 glucose, 20 HEPES, 5 Na-ascorbate, 5 Na-pyruvate, 2 thiourea, 10 magnesium sulfate, 0.5 calcium chloride. This NMDG preparation method was used to improve the overall health of adult slices to ensure sufficient amounts of activity for analysis. ACSF contained the following (in mM): 126 NaCl, 26 NaHCO₃, 2.5 KCl, 1.25 NaH₂PO₄, 1 MgCl₂, 2 CaCl, and 10 glucose. All recordings were at 32.5 ± 1° C. All drugs were dissolved in water to make 10mM and 20mM stock aliquots for the (-) quinpirole and carbachol respectively, which were thawed and diluted into ACSF before each experiment. Application of quinpirole began 10-30 min after the start of an experiment and all analysis was performed on data acquired after quinpirole had been applied for at least 10 minutes.

Imaging

GCaMP3 imaging was performed on a Olympus BX51 upright microscope outfitted with a 20x 1.0NA water immersion lens with 0.5x reducer (Olympus) and ORCA-ER CCD Camera (Hamamatsu Photonics). Illumination was delivered using a Lambda DG4 arc lamp (Sutter Instruments). Light path was delivered through a 472/30 excitation filter, 495nm single band dichroic, and a 496nm long pass emission filter (Semrock). GCaMP3 movies consisted of 36000 frames acquired at 10Hz (1 hr) with 4x4 sensor binning yielding a final resolution of

256 x 312 pixels. Light power during imaging was 100–500 $\mu\text{W}/\text{mm}^2$. The open source Micro Manager software suite (v1.4, NIH) was used to control all camera parameters and acquire movies. Any movies that had significant drift, movement, or lacked significant amounts of activity were excluded from further analysis. Specifically, when movies that exhibited drift greater than 0.25 soma diameters, the experiment was terminated prematurely, and any data collected was excluded from analysis. Significant movement could also be detected during ICA by the appearance of elliptical rather than circular segments.

Signal extraction

All analyses and signal extraction was performed using MATLAB (Mathworks). Locations of cells were automatically identified using a modified version of the published CellSort 1.1 toolbox [13]. Signals were extracted from movies and the baseline fluorescence function, F_0 , was calculated for every trace using the mode of the kernel density estimate over a 100s rolling window, implemented via the MATLAB function `ksdensity` following the procedure outlined in (Mukamel et al., 2009b). All signal traces shown represent normalized versions of the $\frac{F-F_0}{F_0}$ trace.

Threshold based event detection was performed on the traces by detecting increases in $\frac{F-F_0}{F_0}$ exceeding 2.5σ over one second, and then further thresholding these events by keeping only those events which exceeded a 4σ increase over two seconds. σ is the standard deviation of $\frac{F-F_0}{F_0}$, calculated over the entire movie. Thus all detected events have a deviation of at least 4σ from baseline. After identifying these events in the calcium signal from a cell, the

cell was considered active during the entire period from the beginning to the peak of the event. The beginning of the event was defined as the first point for which $\frac{F-F_0}{F_0}$ increases by 2.5σ within 1 second and by 4σ within 2 seconds. The peak of the event was defined as the local maximum of the entire event, from the beginning of the event until $\frac{F-F_0}{F_0}$ returns to the same baseline value. We then created a matrix in which each row corresponds to a neuron, and each column corresponds to a frame. Entries in this matrix were 1 if a given neuron was active during a given frame, and 0 otherwise. All subsequent analyses were performed on this two-dimensional representation of network activity over time (c.f. Fig. 2.1C). Correlations between cells were calculated between the event trains corresponding to those two cells after subtracting the mean level of activity from each event train. Thus the correlation between two event trains and is where is a rolling mean of the event train calculated by convolving with a Gaussian kernel with $\sigma = 50$ sec.

Distance dependent surrogates

For figure S4.3, the distance dependent surrogates were constructed by first measuring the correlations as a function of distance across all experiments as shown in figure 4.2C. These correlations were then binned into 50 micron bins and then a random draw was made for each pair of neurons from a the distribution of correlations corresponding to that pair of neurons physical distance. This resulted in construction of surrogate correlation matrices where correlation was purely distance dependent.

Creating and analyzing functional networks

Edges or functional connections between neurons were determined on the basis of significant correlation. First all datasets were shuffled 100 times and from this a null distribution of correlations was estimated and the threshold for an edge, Σ , was set as any correlation was above the 99% threshold of the null distribution. While the correlation distributions were mean subtracted and thus two tailed, rarely were there any significant negative correlations and thus only positive correlations were considered. In the case of fig. S4.2 where the threshold was done on a per cell basis, networks were shuffled 500 times and the 99% threshold was used for each individual interaction to determine significance. The clustering coefficient for each node was calculated by considering the subgraph of the k nodes connected to a given node then calculating where is the total number of edges between the k nodes connected to the principal node divided by the total possible number of edges between all k nodes which is $\frac{k(k-1)}{2}$. So a clustering coefficient of 1 would mean that all possible edges between k neighbors exist whereas a clustering coefficient of 0 would indicate that none of the k neighbors share an edge. Networks used as benchmarks for random networks were generated using an erdos-renyi random model. In this model, all possible edges are equally likely and networks were generated to match the number of nodes and edge probability of a given dataset. Random path length/clustering coefficients denoted in Fig 4.2D/E represent the average of 100 random networks.

Creating surrogate datasets

As described in more detail in the Results, we created surrogate datasets using the times and durations of active states observed in the original dataset, but reassigned the cell identities associated with each active state. The cell identities were reassigned based on an iterative process and a target correlation matrix, C . To create a surrogate dataset based on the original pattern of correlations, we simply used the correlation matrix obtained from the original dataset, $C_{i,orig}$. To obtain $C_{i,rand}$ we randomly re-arranged the values in $C_{i,orig}$ (while maintaining the symmetry of the correlation matrix). Based on the target correlation matrices, defined above, we re-assigned the cell identities associated with each active state. Specifically, we randomly selected an active state, i . We then found all the active states that overlapped with this active state. Next, we selected the subset of these active states for which new cell identities had already been assigned. Call this set X . Let r_j represent the number of timepoints over which active state $j \in X$ overlaps with active state i , and let n_j represent the identity of the cell assigned to active state $j \in X$. L_i and L_j are the lengths of active states i and j , respectively. Then we constructed a vector,

$$\vec{P}_i = \sum_{j \in X}^0 \frac{r_j}{\sqrt{L_i L_j}} \frac{\vec{C}_{n_j}}{\vec{C}_{n_j}^i}$$

where \vec{C}_{n_j} represents row j of the target correlation matrix, i.e. the target correlations between neuron n_j and the other neurons, and $\vec{C}_{n_j}^i$ contains the current values of the correlations between neuron n_j and the other neurons based on the re-assigned portion of the surrogate dataset. We set elements of \vec{P}_i to zero if the corresponding neuron had already

been assigned to an active state that overlaps with active state i , i.e. element n_j of \vec{P}_i was set to zero $\forall j \in X$. Finally, we assigned active state i to the neuron corresponding to the maximum value of \vec{P}_i . If all the elements of \vec{P}_i were zero, e.g. because there are no overlapping active states that have had new cell identities assigned, then we simply used the cell identity of this active state from the original dataset. For surrogate datasets that were fit to correlation matrices from other experiments, if there was a mismatch in the number of cells in the two experiments, the experiment with more cells was downsampled by randomly choosing a subset of those neurons. For example if activity from dataset A with 80 neurons was going to be fit to a matrix from dataset B with 70 neurons, only 70 randomly chosen neurons from dataset A would be fit to dataset B.

Detecting and quantifying temporal sequences of activity

We implemented a template-matching algorithm to identify stereotyped patterns of sequential activity [14,15]. Briefly, we began with a reference event in a given cell i and then identified all of the other cells that became active in a 5 sec (50 frame) window following the reference event. This was stored as a "template vector" of cell IDs and activation times relative to the reference event (i.e. offset times). This template was then shifted to each subsequent event of cell i . Any events that align with the template constitute a pattern. One frame of jitter was allowed when matching event times. We identified all possible patterns of sequential activation that begin with the reference event. For example, if a reference event in cell 1 was followed sequentially by events in cells 5, 17, and 37, then we would identify

both a four-cell sequence ($1 \Rightarrow 5 \Rightarrow 17 \Rightarrow 37$) and 3 three-cell sequences ($1 \Rightarrow 5 \Rightarrow 17$, $1 \Rightarrow 5 \Rightarrow 37$, $1 \Rightarrow 17 \Rightarrow 37$) as illustrated in figure 4.3. A pattern vector containing the cell IDs and offset times of each matched event was stored for each identified sequence. If this pattern vector matched an existing pattern vector - again allowing one frame of jitter - then it was counted as an additional incidence of that pattern; otherwise, it was stored as a new pattern. For the purpose of defining unique patterns, patterns had to repeat at least three times in data to be counted. This process was repeated iteratively, and every active state in every cell was used as a reference event. The algorithm was not parallelized and required 4 h per dataset running on a 2.0 GHz dual-core processor.

Statistical analysis

Unless otherwise noted, we used the Mann Whitney U-test to compare pairs of groups, repeated measures ANOVA to compare multiple groups, and the two-tailed Kolmogorov-Smirnov (KS) test to compare pairs of distributions. Error bars where shown indicate standard error unless otherwise noted. Cell identification, signal extraction and normalization, event detection, and all other data analysis was done using fully automated routines that were independent of the investigator and thus blinded.

4.5 Figures

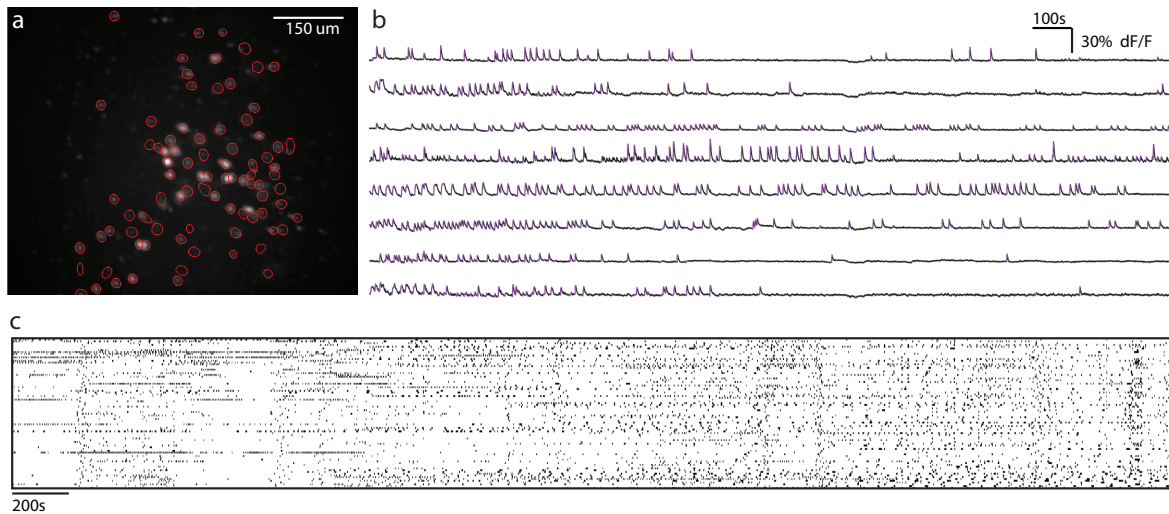


Figure 4.1: **Single photon imaging of GCaMP signals resolves simultaneous activity from many neurons in prefrontal slices**

a, Regions of Interest (ROIs) obtained by an automated algorithm, superimposed on an image showing the maximum standard deviation projection of each pixel during a movie of GCaMP fluorescence. **b**, Sample GCaMP signals from 8 neurons. The overlaid purple lines indicated times where we detected that each neuron was active. **c**, Example raster of spontaneous network activity for a single 60 min experiment with 82 neurons.

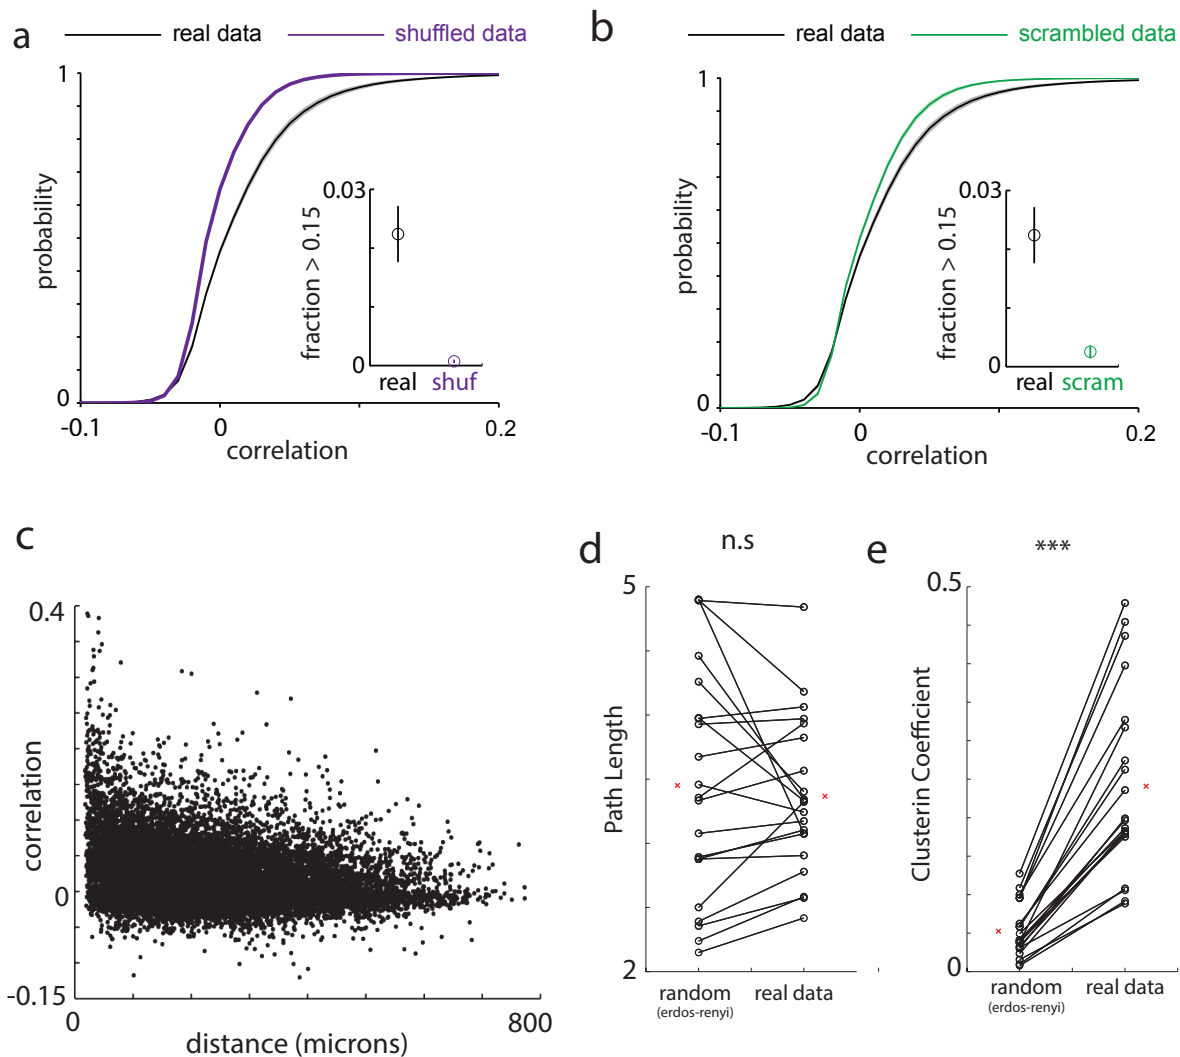


Figure 4.2: **Spontaneous prefrontal network activity is enriched in positive correlations and events in which multiple neurons are co-active.**

a, Cumulative probability distributions for correlations in real datasets (black line; $n = 29$ experiments) vs. those observed in data shuffled by shifting each neurons event train by a different random amount (shuffled, purple line) **b**, Cumulative probability distributions for correlations in real datasets (black line; $n = 29$ experiments) vs. randomly reassigning the neuronal identity associated with each event (scrambled, green line). **c**, Correlations between neurons, plotted as a function of distance between those neurons. The correlation is weakly related to the distance between two neurons, but high correlations occur even at large distances. **d**, The average path length (right) and average path length of a random erdos-renyi connected network of the same number of nodes and edge probability (left). **e**, The average clustering coefficient (right) and average clustering coefficient of a random erdos-renyi connected network of the same number of nodes and edge probability (left). *** $p < 0.001$ by Wilcoxon rank sum test.

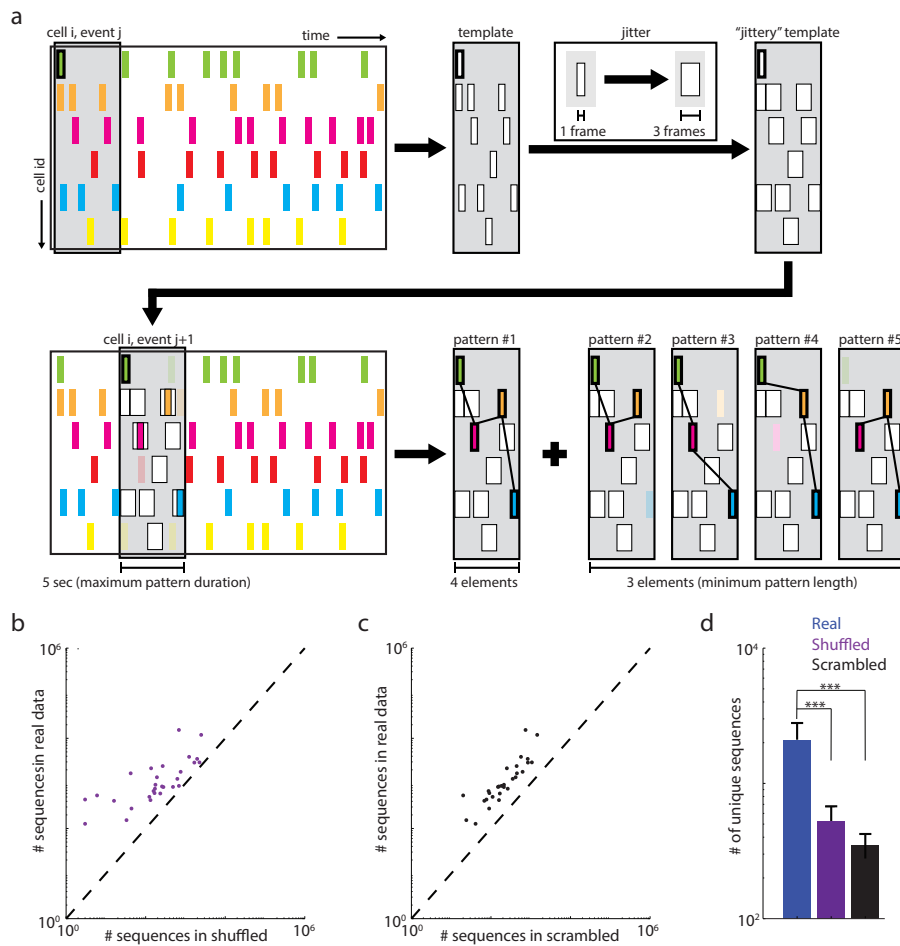


Figure 4.3: **Prefrontal microcircuits spontaneously generate many more stereotyped sequences of activity than expected by chance.**

a, Schematic outlining the detection of temporal patterns. Briefly, a reference event was chosen, and from that event a template was constructed, based on all the active events that occurred during the following 5 seconds. This template was counted as an existing pattern if it matched a previously generated template (allowing for up to 1 frame of jitter), or a new pattern otherwise. Each subpattern contained within the original template was similarly compared. This process was repeated iteratively using all events in all cells as the reference event. **b**, For each dataset ($n = 29$), the number of unique temporal patterns is plotted against the number of patterns in a shuffled version of the same dataset. Each dataset was shuffled by shifting the activity raster for each neuron over in time by a random amount (different for each neuron). All 29 points lie above the unity line, indicating that actual datasets always contain more unique temporal patterns than expected by chance. **c**, For each dataset ($n = 29$), the number of unique temporal patterns is plotted against the number of patterns in a scrambled version of the same dataset. Each scrambled dataset was generated from the original dataset by randomly re-assigning the identity of the cell associated with each epoch of activity. Again, all 29 points lie above the unity line, indicating that actual datasets always contain more unique temporal patterns than expected by chance. **d**, Bar plot summarizing the data plotted in b and c. *** $p < 0.001$ by Wilcoxon rank sum test

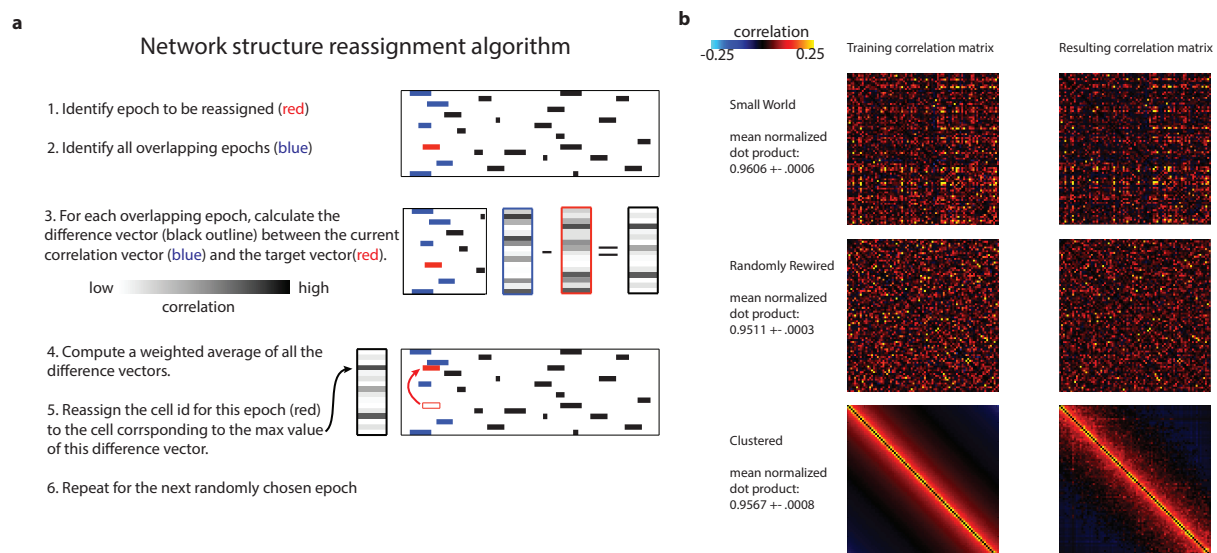


Figure 4.4: **Generation of surrogate datasets with various functional organizations.**

a, Overview of our method for generating each surrogate dataset. The method re-assigns the cell identities associated with each epoch of activity in the original dataset based on a target correlation matrix. **b**, Example of the correlation matrices used for generating surrogate datasets (left) and the correlation matrices for the resulting surrogate datasets (right). Top: Training correlation matrix based on the original pattern of correlations. Middle: Training correlation matrix based on a randomly rearranged correlation matrix. Bottom: Training correlation matrix based on a maximally clustered correlation matrix.

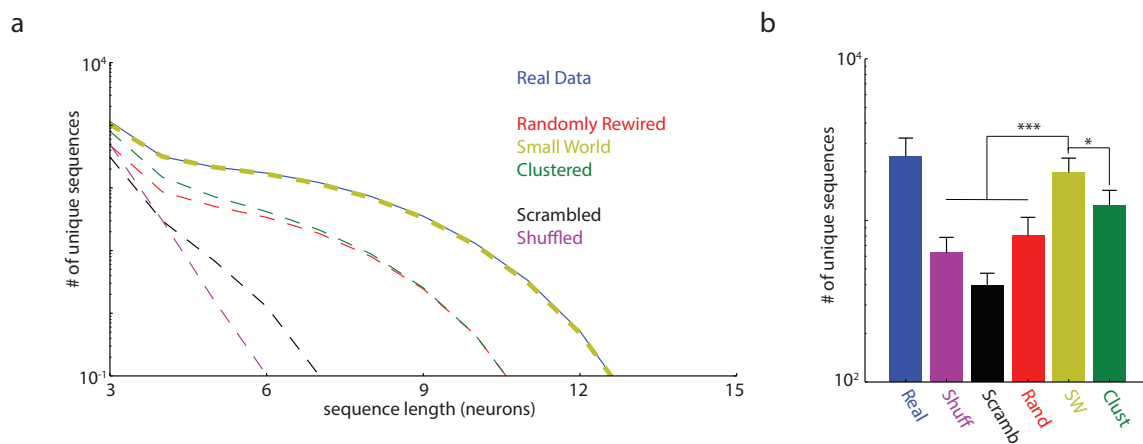



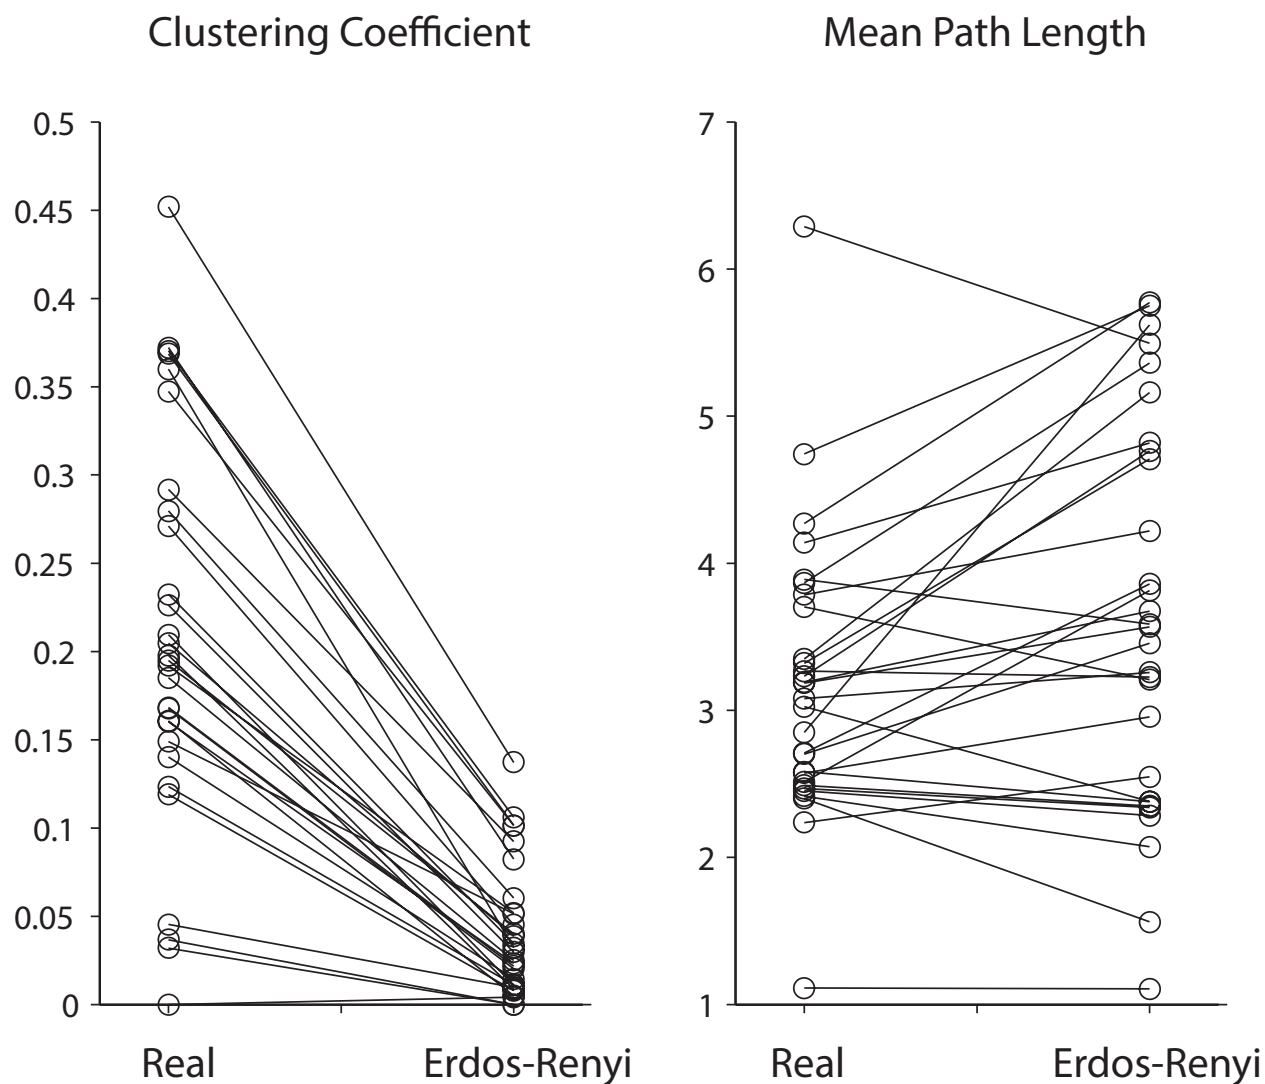
Figure 4.5: **Small world networks outperform clustered and random networks, and reproduce the levels of sequences observed in actual data.**

a, The average number of unique temporal sequences is plotted as a function of the sequence length (in neurons) for surrogate datasets representing various types of networks, as well as our experimentally recorded datasets. For each experimental dataset ($n = 29$), we generated one surrogate dataset of each type, such that all of these surrogate datasets conformed to the temporal pattern of activity in the experimental dataset. **b**, Bar graph summarizing the total number of unique temporal sequences within various types of datasets. *** $p < 0.001$, * $p < 0.05$ by Wilcoxon rank sum test



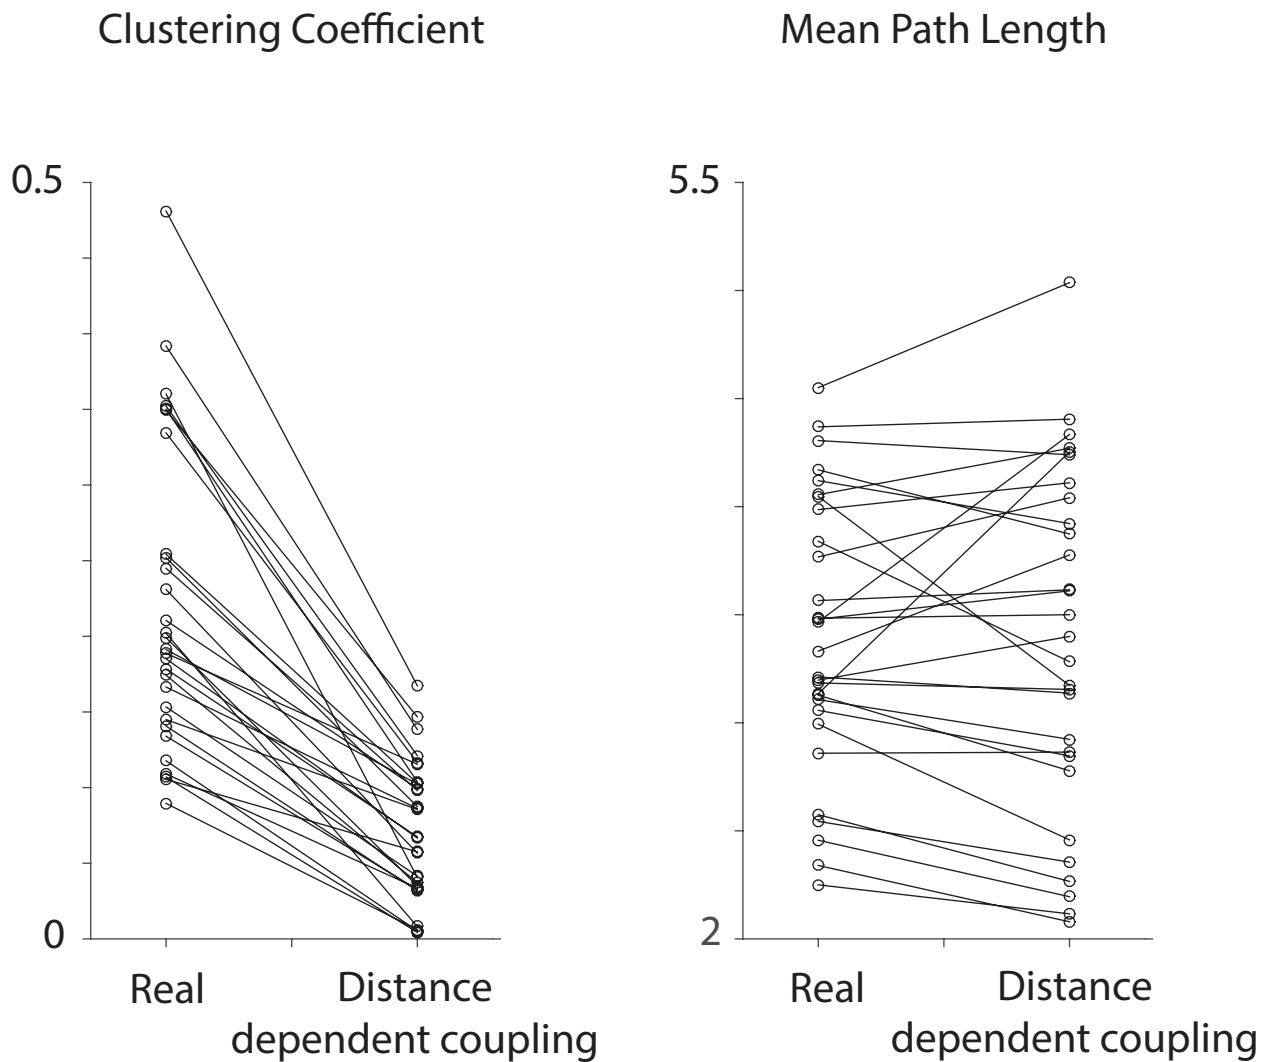
Supplementary Figure 4.1: **The small-world organization of correlations is robust to changes in the event detection threshold.**

We re-analyzed our original datasets using thresholds of either 3, 4, 5, or 6 standard deviations for event detection, then constructed networks based on the pairwise correlation matrix in each case. In every case, the significant correlations form a small world network. Each dot represents the ratio of either the Clustering Coefficient, C , or the Path Length, L , relative to the Clustering Coefficient or Path Length for an erdos-renyi network (C_{er} or L_{er} , respectively), matched to the number of nodes and edge probability. For every threshold, the corresponding network shows non-random clustering and random path-lengths, characteristic of a small world topology.



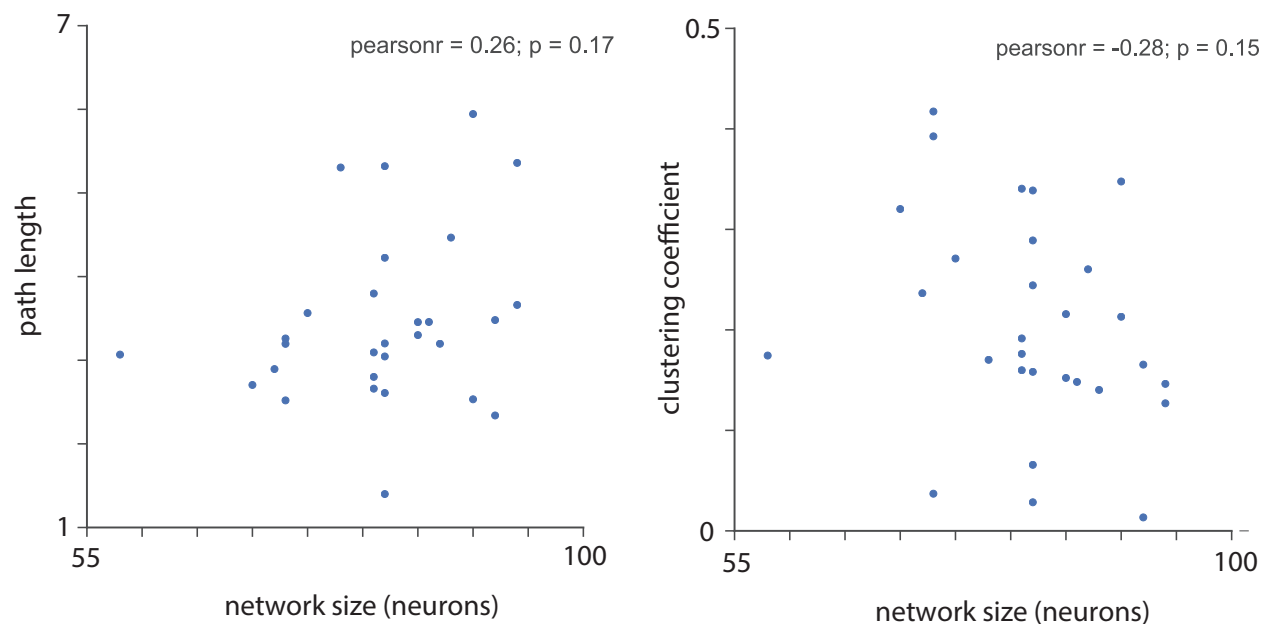
Supplementary Figure 4.2: **The small world organization is maintained even after determining a unique threshold for strong correlations for each cell.**

We generated networks based on the pattern of strong correlations between neurons, but determined a unique threshold for strong correlations for each cell (rather than using a single threshold for all cells in one experiment as was done in our original analysis). Specifically, we defined the threshold for strong correlations as the 99.5th percentile of the shuffled correlation distribution for each individual cell. Each dot represents the ratio of either the Clustering Coefficient, C , or the Path Length, L , relative to the Clustering Coefficient or Path Length for an erdos-renyi network (C_{er} or L_{er} , respectively), matched to the number of nodes and edge probability. Networks generated in this way display non-random clustering and random path-lengths, characteristic of a small world topology.



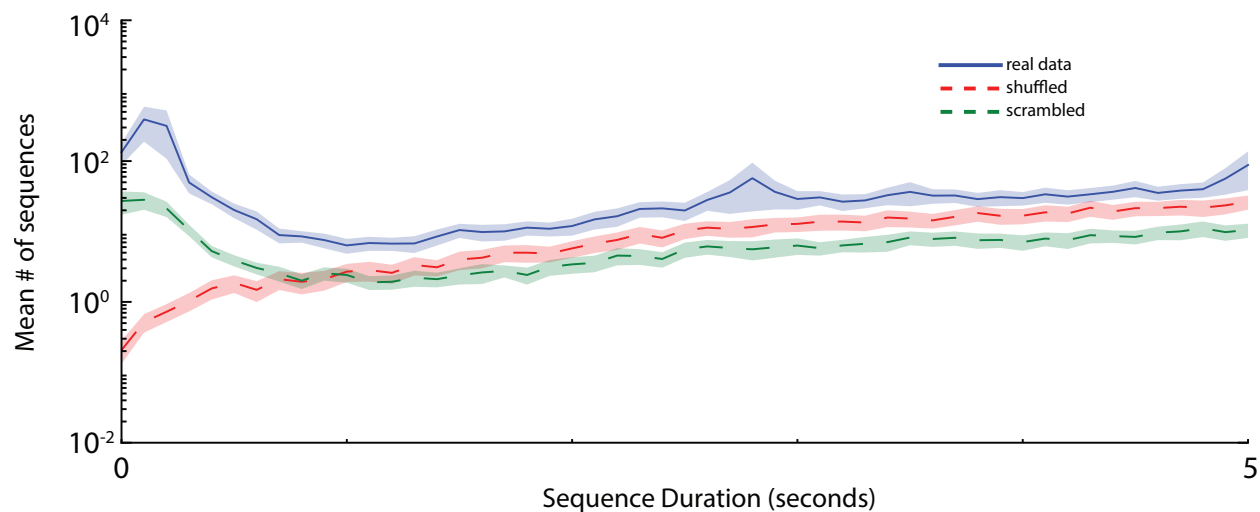
Supplementary Figure 4.3: **Prefrontal microcircuits are more clustered than expected based on a random, distance-dependent coupling rule**

We generated networks by randomly assigning edges between nodes based on the distance-dependent distribution shown in figure 2C. Networks generated in this manner (Distance) exhibit much less clustering than do networks based on the actual pattern of correlations we observed (Real).



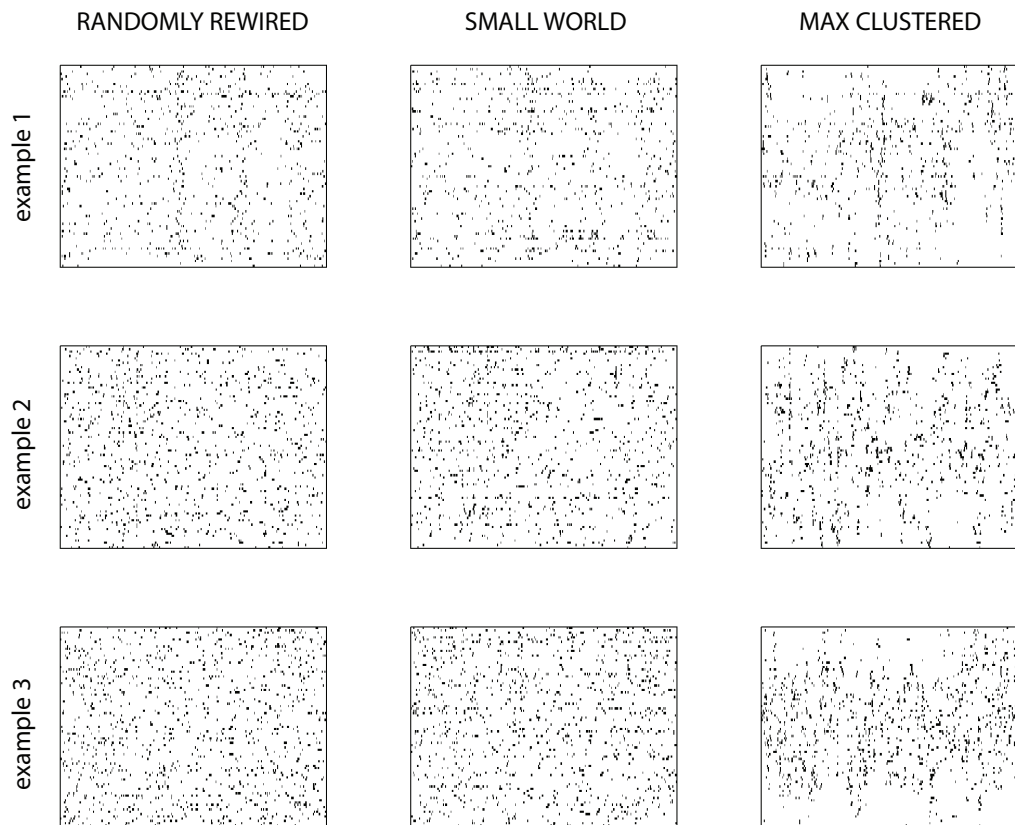
Supplementary Figure 4.4: **Path lengths and clustering coefficients are both independent of network size**

Neither path length nor clustering coefficient correlate with network size (bootstrapped p-values of 0.17 and 0.15, respectively).



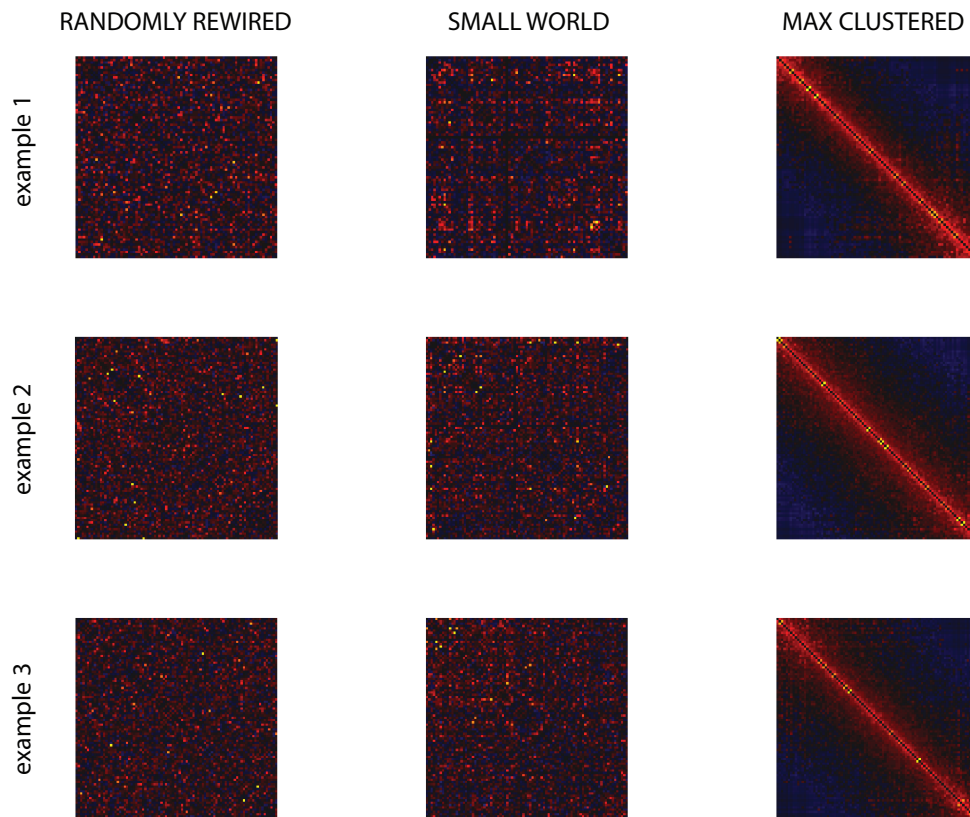
Supplementary Figure 4.5: **Durations of identified multineuron sequences**

Plot showing the distribution of durations for repetitive multineuron sequences identified within actual datasets (blue line), shuffled data (dashed red line), and scrambled data (dashed green line). Each plot has been averaged over all experiments. Most sequences observed in real data have durations <0.5 seconds.



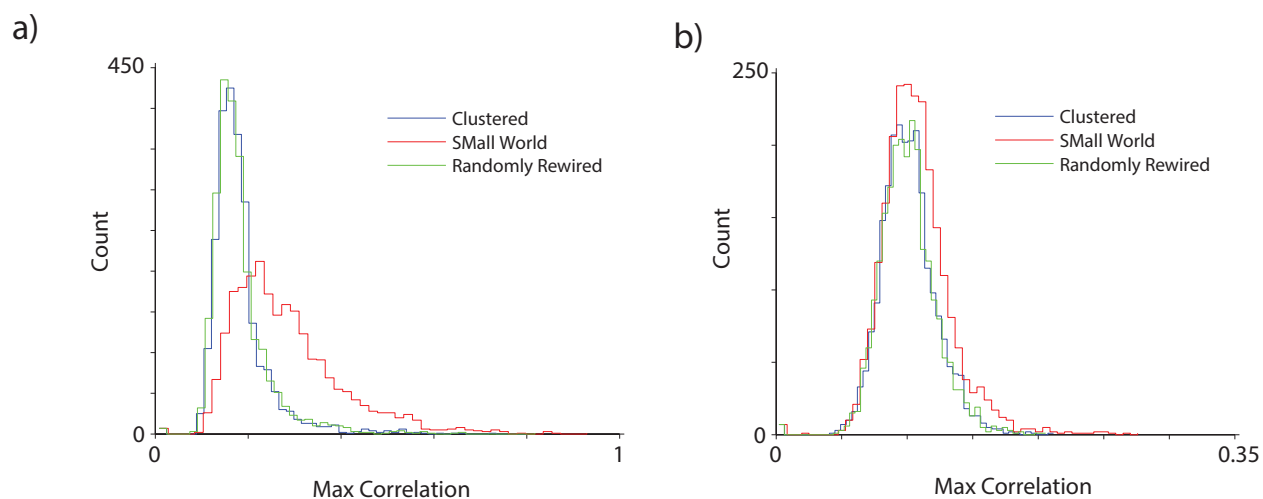
Supplementary Figure 4.6: **Sample activity rasters from surrogate datasets.**

Examples of the three different types of surrogate datasets, each generated for three different experiments. For each experiment, we generated surrogate datasets based on the original pattern of activity, and either a randomly rewired correlation matrix (left), the original small world correlation matrix (middle), or a maximally clustered correlation matrix (right). Thus all three of the datasets in each row, which correspond to different network topologies, nevertheless have the same temporal level of total network activity. I.e. at each point in time, the same number of neurons will be active in all three networks. Each raster represents 600 seconds of data.



Supplementary Figure 4.7: **Example correlation matrices from surrogate datasets.**

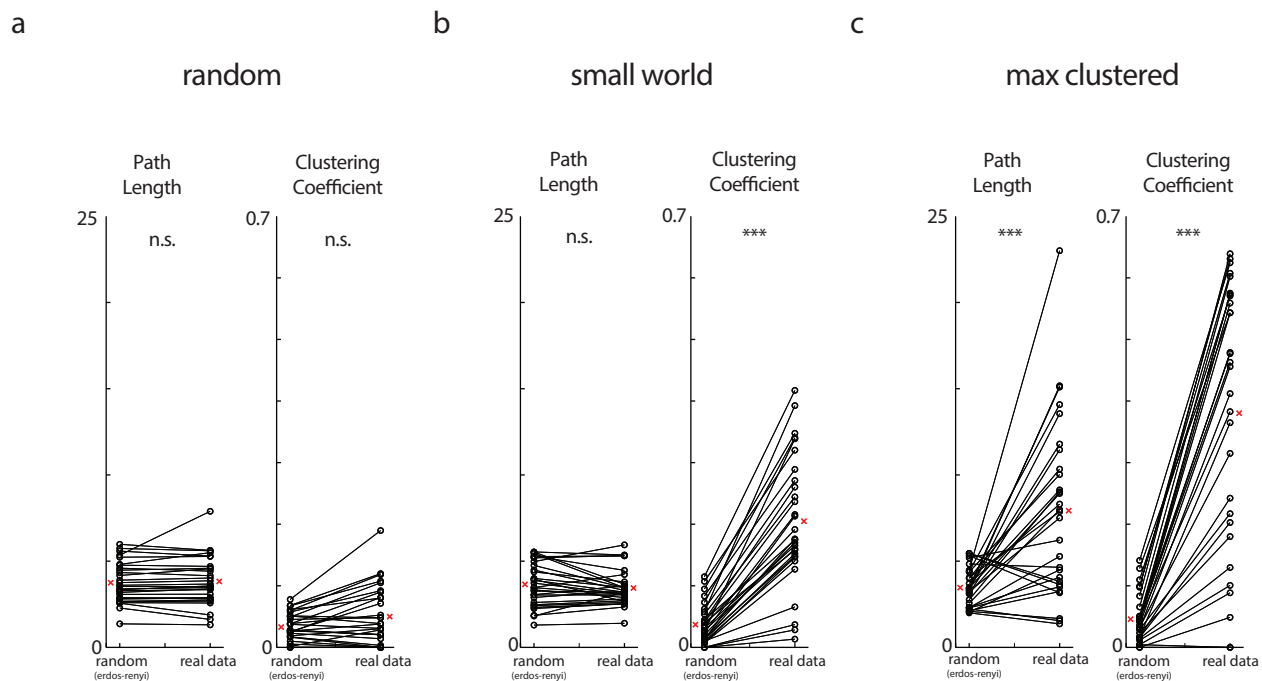
Correlation matrices corresponding to each of the activity rasters presented in Figure S4.6.



Supplementary Figure 4.8: **Distribution of correlations between neurons in surrogate datasets and the most similar neurons in actual datasets.**


a, For each neuron in a surrogate dataset, we found the maximum correlation between its event train and the event train of a neuron in the original (actual) dataset, i.e. the correlation with the most similar actual neuron. The distribution of these correlations is plotted for each type of surrogate dataset. In all cases, these correlations tend to be substantially <1 , indicating neurons in the surrogate datasets are not just renamed versions of the original neurons. Note however that neurons in small-world surrogate datasets (red) tend to have slightly higher correlations with actual neurons than neurons in either maximally clustered (blue) or randomly rewired (green) datasets.

b, Same as in a), except that surrogate datasets have now been generated using the correlation matrix from one experiment, and the activity raster from a different experiment. This procedure eliminates the slightly higher correlations between simulated neurons and real neurons observed in the small world surrogate datasets shown in a).



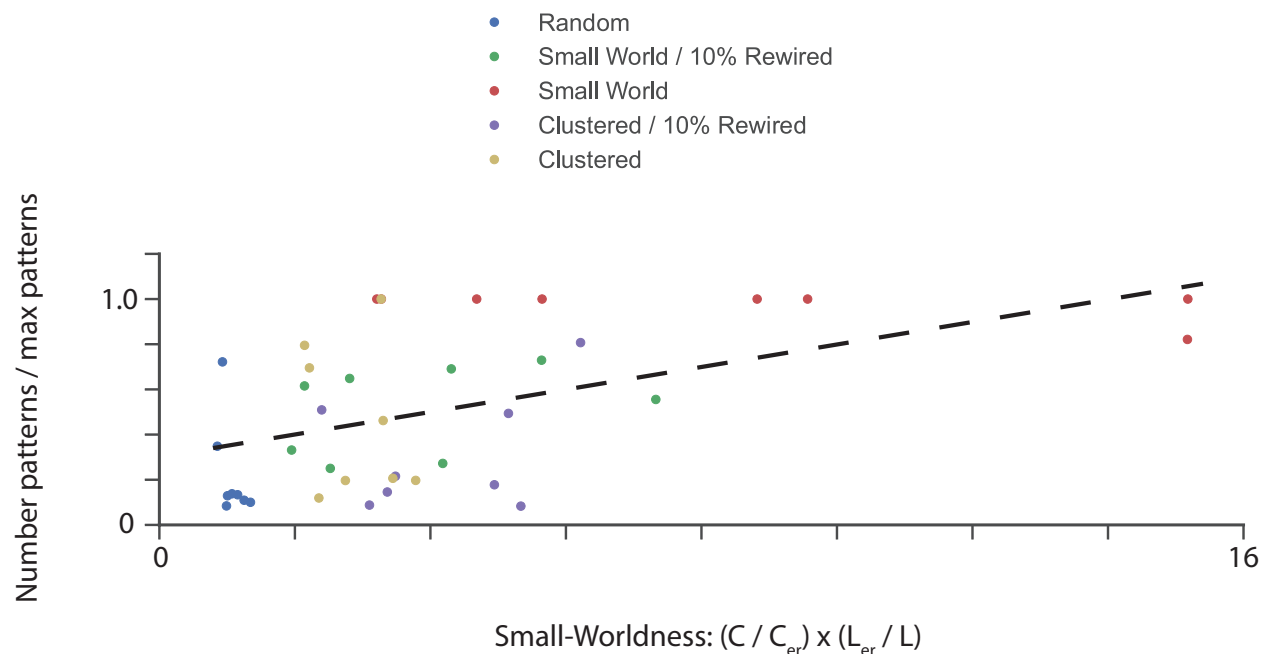
Supplementary Figure 4.9: **Path lengths and clustering coefficients for surrogate datasets**

a, Path length (left) and clustering coefficient (right) for surrogate datasets fit to randomly rewired correlation matrices as compared to erdos-renyi random networks. Randomly rewired networks match the path lengths and clustering coefficients of erdos-renyi networks. **b**, Path length (left) and clustering coefficient (right) for surrogate datasets fit to maximally clustered correlation matrices as compared to erdos-renyi random networks. Clustered networks have greater path lengths and clustering than erdos-renyi random networks. **c**, Path length (left) and clustering coefficient (right) for surrogate datasets fit to small world correlation matrices as compared to erdos-renyi random networks. Small world networks have similar path lengths and greater clustering compared to erdos-renyi random networks.



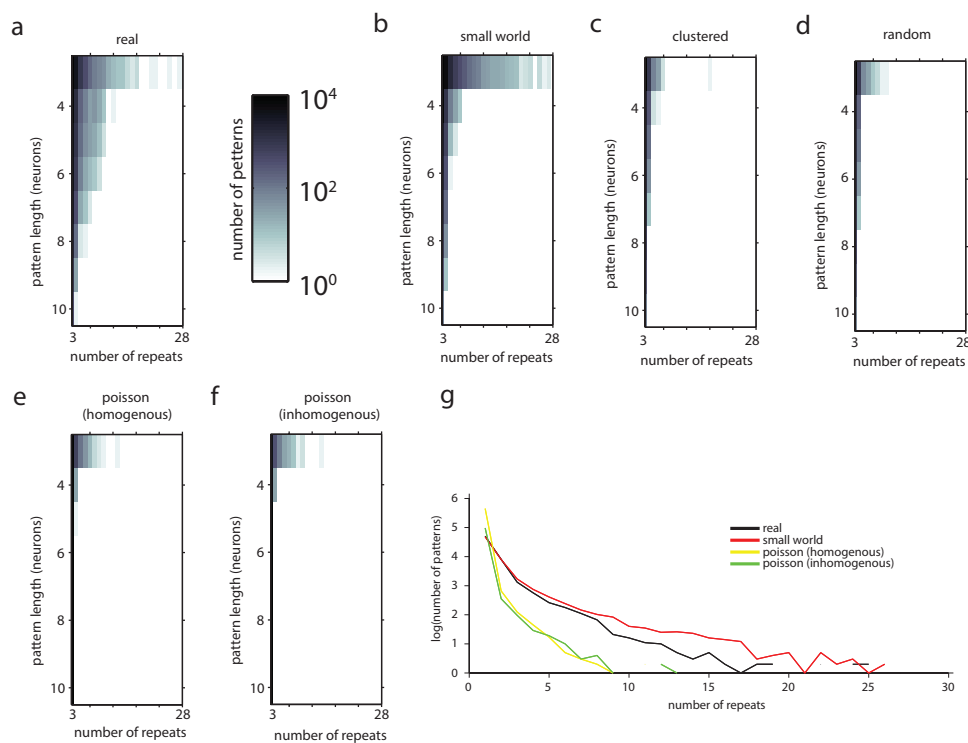
Supplementary Figure 4.10: **Small world surrogate datasets based on correlation matrices from one experiment and activity rasters from another outperform clustered and random networks.**

Surrogate datasets generated using activity rasters from one experiment and correlation matrices from different experiments exhibit maximal numbers of sequential patterns of activity when the correlation matrix is small world than when the correlation matrix is maximally clustered or randomly rewired. For this analysis we generated 3 surrogate datasets of each type (each one based on a correlation matrix from a different experiment) for each activity raster. ** $p < 0.01$, * $p < 0.05$ by Wilcoxon rank sum test



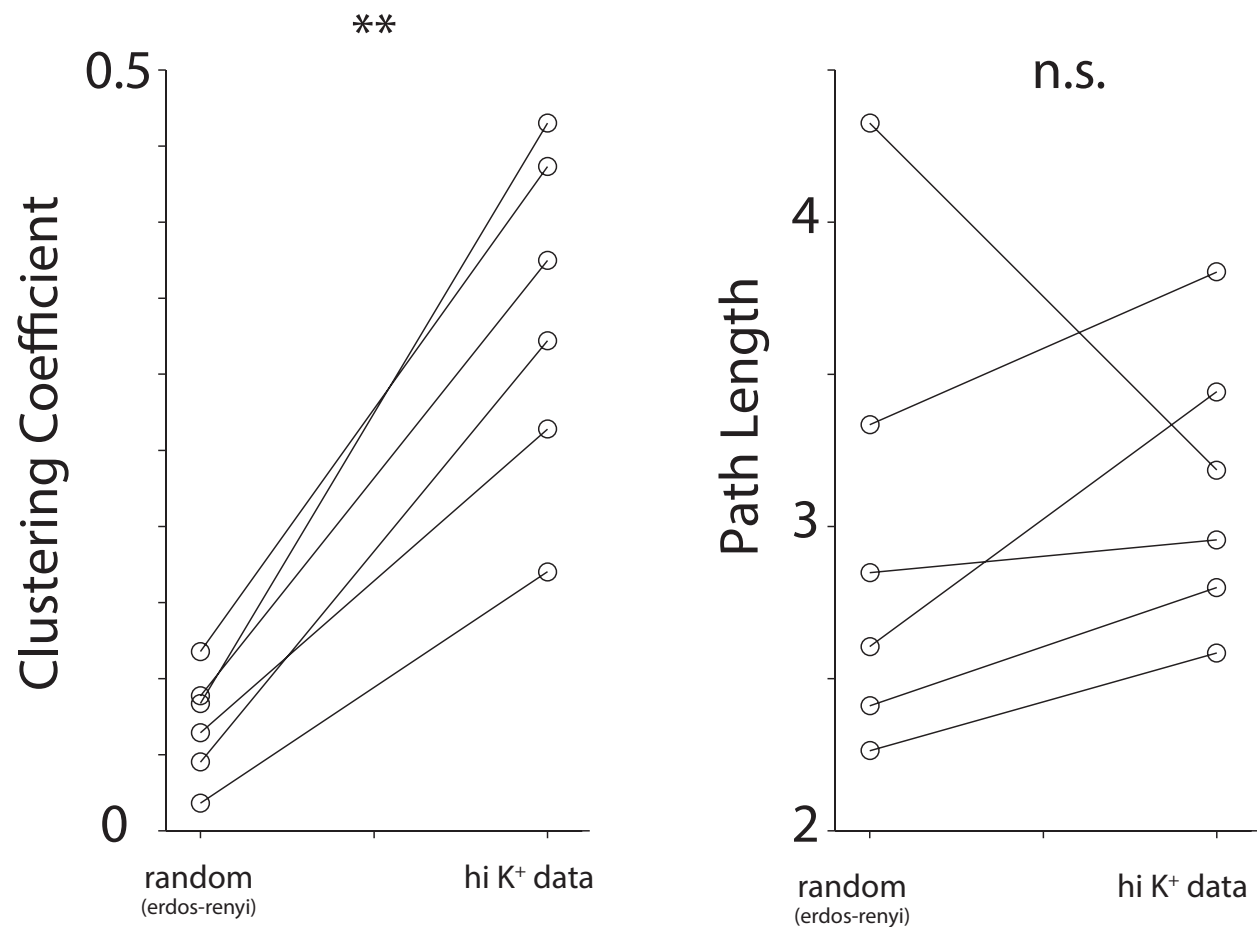
Supplementary Figure 4.11: **Number of patterns increases as a function of small-worldness.**

For each of 7 original datasets, 5 surrogate datasets were generated: small-world, clustered, and random datasets, as well as 2 intermediate datasets generated by randomly reshuffling 20% of the connections from either the small world or clustered datasets. The small-worldness of each datasets was defined as $sw = \frac{C}{C_{er}} \frac{L_{er}}{L}$, where C and L are the clustering coefficient and path length of the network, and C_{er} and L_{er} are the average clustering coefficient and path length for erdos-renyi random networks thus yielding a continuous metric of small-worldness and a continuum of networks with varying degrees of small-worldness. The Y axis shows the number of patterns observed in each surrogate dataset, normalized by the maximal number of patterns observed across any of the 5 surrogate datasets generated from the same original dataset. The number of patterns is plotted as a function of the small-worldness of each dataset (on the X axis). The dotted line was obtained from these points by linear regression and has an R value of 0.44, and a bootstrapped p-value of 0.0041.



Supplementary Figure 4.12: **Distribution of number of sequence length vs. number of repeats across all experiments.**

a, Heatmap representing the number of sequences that were of a certain length and certain number of repetitions across all experiments. The colormap is represented on a logarithmic scale. ($n=29$) **b**, Heatmap representing the number of sequences that were of a certain length and certain number of repetitions across in small world surrogates. ($n=29$) **c**, Heatmap representing the number of sequences that were of a certain length and certain number of repetitions across in clustered surrogates. ($n=29$) **d**, Heatmap representing the number of sequences that were of a certain length and certain number of repetitions in rasters generated from rate-matched homogenous independent poisson neurons. ($n=29$) **e**, Heatmap representing the number of sequences that were of a certain length and certain number of repetitions in rasters generated from rate-matched inhomogenous independent poisson neurons. ($n=29$) **g**, Plot of the number of sequences that repeat a certain amount of time across all datasets in experimental data (black line), small-world surrogates (red-line), rate matched homogenous poisson neurons (yellow line), and rate matched inhomogenous poisson neurons (green line). Poisson neurons fail to capture the patterns that repeat >5 times.



Supplementary Figure 4.13: **Datasets exposed to active ACSF instead of carbachol also exhibit a small-world organization.**

Experiments on slices driven using a high potassium/low calcium active acsf instead of carbachol also display random path lengths and non-random clustering. ** $p < 0.01$ by Wilcoxon rank sum test

Chapter 5

Extracting and reducing network wide structure in chronic intracranial electrocorticographic (ECoG) recordings

5.1 Introduction

Information processing at the level of networks in our brains endows us with the capacity to perform specific yet fundamental tasks in our everyday lives. One major challenge to studying information processing at this emergent level has been that the ability to record from multiple neurons or regions has been historically limited. Recent advances in functional

magnetic resonance imaging (fMRI) have enabled neuroscientists to probe communication across multiple brain regions simultaneously albeit with limited temporal resolution (Harrison et al., 2015; Smith et al., 2013). fMRI studies have begun to identify multiple distinct resting state networks such as the default mode network and dorsal/ventral attentional networks (RSNs) (Fox et al., 2006). These networks have begun to shed light on how information is processed in the human brain and hold promise in the study of neuropsychiatric disease (Greicius, 2008; Zhou et al., 2010).

However, identifying such networks using fMRI has the noted caveats that only slow-varying networks can be identified due to the slow nature of the BOLD signal being recorded. Absent from these datasets is the ability to observe high frequency brain rhythms that have been posited to be important for multiple social and cognitive processes, as well as represent an important substrate for information processing between brain areas (Fries, 2009; von Stein and Sarnthein, 2000). Chronic intracranial electrocorticographic (ECoG) recordings can record local electrical signals with much higher spatial and temporal frequency than fMRI. This is of vital importance especially given recent evidence that rhythms at different frequencies may subserve distinct roles in how and where information is routed (Watrous et al., 2013). Thus identifying frequency specific subnetworks represents an important step towards understanding how information is flexibly routed in the brain and how dysfunction in this information processing might underlie neuropsychiatric disease conditions.

Recent developments in analytical as well as experimental methods have started to enable the study of time varying brain wide dynamics as well as start to identify putative network-

level biomarkers of neuropsychiatric disease (Calhoun et al., 2014). However these studies have been limited to the slow dynamics of fMRI and thus have not been able to explore networks as defined by higher frequency rhythms in the brain. Here we present a novel analysis of chronic intracranial ECoG recordings across various areas in the mesolimbic circuit. Using pairwise signal coherence and dimensionality reduction techniques, we show that indeed significant interactions across distant brain areas exist and that these dynamics are surprisingly low dimensional. We also show that these identified subnetworks contain a large amount of redundancy that allows these networks to be summarized by one or a few highly predictive key interactions thus further reducing the dimensionality. We also show that these identified subnetworks are frequency specific and that similar results can be obtained using either principal components analysis (PCA) or independent components analysis (ICA) suggesting that brain wide coherences can be broadly approximated by the first two cumulants of their distributions and lack significant higher order cumulants. The approach outlined here could serve as a template for future studies of brain-wide network function and highlights the low-dimensional nature of seemingly complex brain-wide interactions.

5.2 Results

We analyzed data acquired from medically refractory epilepsy patients chronically implanted with a combination of surface and depth electrodes over multiple areas in the mesolimbic circuit. Cortical coverage was generally comprised of some subset of cingulate, insula, orbitofrontal, fronto-parietal, and temporal coverage. Subcortical coverage generally comprised

of amygdala and/or hippocampal coverage. Despite diversity of electrode locations between patients, most patients had ≥ 110 electrodes covering various parts of mesolimbic areas. Due to this variability in electrode placement across patients, we focused largely on global metrics of organization but will also present data from individual patients. Coverage across all 8 patients summarized in this study is included in 5.1.

ECoG signals were collected from these patients continuously over the course of 10-21 days and patients underwent neuropsychiatric testing including a Beck Depression Index (BDI) and Beck Anxiety Index (BAI) presented in table 5.2. In addition, some of these patients were asked to perform immediate mood scalar (IMS) tests to measure their within and across day variability in mood state. This mood data will be the focus of a later study in which we attempt to identify a neural biomarker for variations in mood state. Instead in this study we focus on more global statistics of the data such as the dimensionality of such datasets and how to identify key interactions between brain regions to target for future closed-loop control of neuropsychiatric disease conditions.

Extracting relevant network structure

In order to identify network-wide interactions we started by computing pairwise magnitude squared coherences across electrodes. Signal coherence in this context is defined as the cross spectral density over the product of the auto-spectral densities of the two ECoG signals. This measure can also be thought of as a linear correlation between two signals in the frequency domain. Thus coherence between two signals varies as a function of frequency and

we analyzed coherences separately for several frequency bands of interest namely delta (δ , 1-4 Hz), theta (θ , 4-7 Hz), alpha (α , 7-15 Hz), beta (β , 15-30 Hz), low gamma (γ , 30-70 Hz), and high gamma (Γ , 70-120 Hz).

Given that electrodes from the same grid or strip tend to have much higher correlation in their coherences due largely to spatial proximity, we further reduced the dimensionality of the coherences to one value between two regions by taking the average of all the electrodes for a given region. For example the single value representing the coherence between amygdala and hippocampus represents the mean coherence across the 12 different combinations of coherences between the 4 amygdala electrodes and the 4 hippocampal electrodes. Since the coherence measured between two regions lacks directionality, the matrix representing the network-wide interactions at a given point in time is mirror symmetric about the diagonal and values along the diagonal correspond to the average coherence across all electrodes within a given region (Fig. 5.1A). For future reference, we will call this matrix the network interaction matrix.

Computing this network interaction matrix in non-overlapping 10 second bins over time lets us compute a time series with momentary estimates of network-wide interactions over time. In order to identify network-wide interactions, we applied Principal Components Analysis (PCA) to the data to identify the principal component (PC) vectors that accounted for the largest amount of the variance in our data subject to relative orthogonality. Examples of such principal components are shown in supplementary figures 5.4, 5.5, 5.6. The first PC accounts generally accounted for 35% of the total variance in the data depending on the

frequency band. (Fig. 5.1B) These PCs representing covarying interactions within and between regions that can be thought of as representing distinct subnetworks of activity or different modes of activity across subsets of regions.

One way to test whether the structure present in these PCs is a trivial consequence of the statistics of the coherences or whether it represents real network-wide interactions is to compare the eigenvalue distribution of real data with that of shuffled surrogates. We generated shuffled surrogates by splitting the entirety of the data into 5 segments and then randomly shuffling those segments relative to each other as diagrammed in supplementary figure 5.1. This method preserves the overall mean, variance, and autocorrelation of each coherence signal but should disrupt any cross-correlation between areas. We performed PCA on these shuffled surrogates and find that in general you only observe 1-2 non-zero loadings per PC in contrast to the real data where we observe many more non-zero loadings.

To estimate the number of relevant PCs, we compared the eigenvalue distribution of the real data compared to that obtained from shuffled data and asked how many PCs exceed the eigenvalue of the corresponding shuffled surrogate (Fig. 5.1C; Supplementary figure 5.2). At lower frequency bands there were 4 significant PCs on average for delta band PCs, decreasing with increasing frequency to 3 significant PCs in the high gamma range (Fig. 5.1C). The significant PCs in each dataset accounted for 45-60% of the total variance of the datasets depending on frequency (Fig. 5.1D). Thus, PCA can extract PCs that represent network-wide interactions and a large amount of the variance can be accounted for with only the first few PCs.

Further reduction: identifying key interactions

The identification of network-wide interactions in PCs that account for large amounts of variance in our data opens the possibility of exploring the behavior within these different PCs. Figure 5.2 illustrates the first alpha band PC of a given patient that shows strong interactions between subtemporal, amygdala, insula, and cingulate electrodes. In the case of the example patient presented in figure 5.2, PC1 represents a subtemporal, insula, cingulate, amygdala subnetwork, PC2 represents a cingulate, insula, amygdala subnetwork, and PC3 would represent an amygdala and hippocampus subnetwork.

If we now consider the subtemporal, amygdala, insula and cingulate subnetwork defined by PC1 we wondered whether we could assay how active this or any other subnetwork was at any given point in time. To this end we computed the linear projection of the original coherence data onto the PC vectors identified producing a single value that provides a scalar readout for how well the state of the network of coherences at a given point in time is aligned with a given subnetwork (Fig 5.2A/B). This linear projection can be thought of as the amount of activation of a subnetwork defined by a given PC.

While this single scalar is a proxy for the amount of activation within a given subnetwork, it is a value that requires us to observe activity across the full set of regions to calculate. In future experiments, we wondered whether the effective activation of a subnetwork could be inferred from one or a few of the corresponding interactions so as to limit surgical intervention in future patients. In an attempt to build this reduced representation of the state of the subnetwork, we started by asking how correlated each interaction was with the

projection onto that subnetwork. A sample distribution of correlations between interactions and the projection onto PC1 is illustrated in figure 5.2B. What you can see is that while most interactions carry very little information about the state of the subnetwork, a few key interactions are highly informative of the state of the network with some of them such as the INSaST interaction having an R^2 value of 0.7 (Fig. 5.2B). These highly informative key interactions also display a high amount of redundancy as evidenced by the similarity between the aST-aST, iCIN-AMY, and aST-pINS, highlighting that one does not need to know the coherences in all of these regions to estimate the state of that subnetwork at any given point in time (Fig. 5.2B). Thus for the example patient in the alpha band illustrated in figure 5.2, all 3 of the subnetwork projections can be approximated well using a linear combination of the top 3 key interactions as determined by those with the highest predictability with respect to the projection of that subnetwork (Fig. 5.2C). Further example of this approach of identifying subnetworks and then reducing these subnetworks to a few key interactions are shown in supplementary figures 5.4, 5.5, 5.6.

Importantly, there exists a basal amount of correlation that will be introduced by the loading of a given interaction within that PC. To control for this, we computed the correlation of a given coherence over time with the projection onto a PC that excluded that interaction, thus removing any effect of loading value on the ultimate projection. Doing so, still allows certain key interactions to account for 70% of the variance observed in the EC71 alpha PC1 subnetwork over time (Fig. 5.3A). Importantly, in shuffled surrogates, the predictability of coherence signals against projections excluding that interaction falls to zero (Fig. 5.3B).

While thus far we have illustrated this result for the alpha band for patient EC71, this same result holds in the alpha band across all patients (Fig. 5.3C, D, E). Lastly, this result holds across frequency bands suggesting that the ability to reduce the activity of these subnetworks to key interactions represents something fundamental about the organization of neural activity in meso-scale circuits (Supplementary figures. 5.7, 5.8). Thus, meso-scale networks are organized in a low-dimensional manner in which there exists significant redundancy between areas that allows the observation of the state of a given subnetwork by focusing on just one or a few key interactions. What sort of underlying statistical structure can produce such patterns will be an important question for future studies.

Frequency specific subnetworks

Given that the advantage of ECoG signals as compared to fMRI BOLD signals lies in the ability to probe higher frequency information, we explored whether the subnetworks we observed represented broad band networks present at multiple frequencies or whether they were isolated to specific frequencies. To answer this question we computed pairwise coherences and performed PCA to extract the top 3 subnetworks for each of the 6 frequency bands mentioned previously: delta (δ , 1-4 Hz), theta (θ , 4-7 Hz), alpha (α , 7-15 Hz), beta (β , 15-30 Hz), low gamma (γ , 30-70 Hz), and high gamma (Γ , 70-120 Hz). We asked whether there was any similarity between the subnetworks identified in the different frequencies by computing the R^2 between the top 3 PCs in one frequency band as compared to the top 3 PCs in another frequency band (Fig. ??A). We then matched PCs across frequency band

by assigning PC from a given frequency to its matched PC from another frequency band by simply assigning it to the PC with which it had the highest correlation. The distribution of matched correlations and unmatched correlations were largely non-overlapping, suggesting that there existed a fair amount of consistency in PCs across frequency (Fig. ??B). We next computed the mean correlation of the matched PCs between all of the different frequency bands as a measure of how similar or dissimilar the subnetworks in the different bands were with respect to each other. We found that subnetworks in closer frequency bands e.g. alpha and beta, were more similar than those from distant frequency bands e.g. delta and gamma (Fig. ??C). Quantifying more systematically this frequency dependent bias, PCs were most similar for neighboring frequency bands (1 band difference) at a mean of 0.7 and dropped systematically as a function distance in frequency space to 0.4 for bands separated by 5 other bands e.g. delta and high gamma (Fig. ??D).

Lastly, we wondered how well this approach would generalize to other dimensionality reduction techniques. In using PCA, the PCs extracted represent the largest sources of variance subject to the constraint of linear orthogonality to other PCs. Another commonly used method for identifying individual sources of variance is independent components analysis (ICA) which yields independent components (ICs) representing independent sources of variance. We repeated our dimensionality reduction of the pairwise coherences as discussed previously except using the ICA algorithm to identify subnetworks instead of PCA and found remarkably similar results (Fig. 5.5). Dimensionality reduction of the coherences using ICA instead of PCA yields ICs that are qualitatively very similar to the previously computed

PCs (Fig. 5.5A/B). In addition, the ICs can be used interchangeably with the PCs, with key interactions identified using PCA having very high predictability against subnetworks defined using ICA (Fig. 5.5C/D). This analysis argues that the specific dimensionality reduction technique used is not as important and reflects that global brain dynamics e.g. pairwise coherences can be largely described using only the first two cumulants e.g. mean and variance.

5.3 Discussion

In this study, we sought to extract and characterize network-wide structure in chronic intracranial ECoG recordings across the mesolimbic circuit. Using time-varying spectral coherence and PCA, we were able to define functional subnetworks, representing groups of covarying interactions within and between brain regions. We also show that within these subnetworks there exist highly informative key interactions, which can be used to infer the current state of that subnetwork with great precision. Importantly, these key interactions were present in multiple patients and across multiple frequency bands suggesting that this low-dimensional representation represents a global principle of networks in the mesolimbic circuit.

Of course, one limitation of our findings is that these networks are defined functionally and not anatomically, however studies have shown the two to be intimately linked (Stafford et al., 2014; Wang et al., 2013). It will be important to verify any novel identified subnetworks, but many of the recurring interactions extracted in our analyses such as the amygdala

and hippocampal interactions, have been previously shown to exist (Pitkanen et al., 2000). Functional studies can thus be an important tool to explore patterns of neural information processing as well as guide future anatomical studies.

Dimensionality of functionally defined ECoG networks

Most sensory and cognitive functions are carried out in concert by large numbers of neurons. With the advent of multi-neuron recording techniques that grant experimenters access to the simultaneous activity of ensembles of neurons, statistical techniques for understanding emergent properties have become critically important. Here we took advantage of one of the more common dimensionality reduction techniques, PCA, and show that functionally defined ECoG networks are low-dimensional in nature. While a similar result can be obtained using ICA, it should be noted that many other assumptions can be made about the nature of the data and that other dimensionality techniques could be equally applicable (Cunningham and Yu, 2014). It is possible that rather than assuming a linear manifold of underlying states as is the assumption, a non-linear technique such as linear discriminant analysis might yield better identification of subnetworks (Scholkopf and Mullert, 1999). Exploration of multiple dimensionality reduction techniques in extraction of network topology from coherences will be an important aspect of future studies.

Implications for frequency-specific subnetworks

Meso-scale neural signals such as the BOLD signal recorded in fMRI or ECoG signals have enabled simultaneous measurement of multiple human brain areas simultaneously with the noted caveat that such signals represent the collective activity of hundreds to thousands of neurons. The ECoG signals studied here likely represent a heterogenous mix of neurons and signals, and thus a key question becomes how to extract these individual signals? These signals could be existing in distinct frequency bands, be coupled to distinct areas, or some combination thereof. The fact that we are able to extract frequency specific subnetworks suggests that it is indeed possible to separate out distinct signals present in the global ECoG signal by examining specific frequency bands and across region interactions. As more studies seek to identify neuropsychiatric disease biomarkers, it will be important to develop analytical techniques, such as the approach presented here, for demixing the many heterogenous signals present in the global neural ECoG signal.

Towards closed-loop control

Neuropsychiatric disease remains one of the toughest human conditions to adequately treat. This is likely due to the fact these diseases rather than reflecting gross anatomical changes or aberrations are often belied by subtle changes to the networks and patterns involved in information processing in the brain. With the advent of new molecular, genetic, as well as electrical tools we are now able to both study and treat some of these diseases with greater precision. One such promising approach is the development of closed loop stimulation

paradigms for neural circuits in the brain. However, if networks in the brain are inherently high dimensional this could prove difficult as it would require multiple distributed inputs and readouts for the system. The identification here of key interactions, which allow for the readout of an instantaneous network state across multiple distributed regions opens up the possibility for sensing and affecting distributed subnetworks, given that you can identify these key interactions a-priori. Further analysis could help disentangle whether these key interactions are drivers or simply readouts of network activity but this could also be assayed directly through stimulation experiments. Regardless of the result, the low dimensionality of the networks in the mesolimbic circuit coupled with identifying key interactions opens up the possibility of developing closed loop control of network state in a circuit level treatment of neuropsychiatric disease.

As systems neuroscience moves in a direction of recording with higher density, increased spatial resolution, increased temporal resolution, and simultaneously across multiple areas we will face constraints on the data analysis front. One challenge for the field will be to find ways to make sense of such large volumes of data and glean reasonable biological insights. The approach outlined here illustrates that under certain assumptions, networks in the brain can be quite low-dimensional. In addition, this low dimensional nature can be characterized quite effectively using coherence and dimensionality reduction to identify subnetworks of interactions.

5.4 Materials and Methods

Data collection

Data were collected from eight subjects with intractable epilepsy who were implanted with chronic subdural grid electrodes as part of a pre-operative procedure to localize the epileptogenic focus. The surgeons determined electrode placement and treatment based solely on the clinical needs of each patient. Data were recorded at the University of California, San Francisco (UCSF) hospital. All subjects gave written informed consent to participate in the study in accordance with the University of California, San Francisco Institutional Review Board. ECoG data were acquired using a omega system and sampled at either 256 or 512 Hz. ECoG data were individually referenced to the average potential across all electrodes in a given grid, strip, or electrode. Electrode locations and general coverage is listed in table 5.1.

Electrophysiology pre-processing

All electrophysiological data were analyzed in PYTHON using custom scripts. ECoG data were individually referenced to the average potential across all electrodes in a given grid, strip, or electrode. Signals were also bandpass filtered at 60 and 120 Hz with a bandwidth of 4Hz. Movement and signal artifacts were removed using an independent components analysis (ICA) approach. Signals were decomposed into X independent components then components corresponding to noise were identified using a support vector machine (SVM)

trained on the power spectra of the signals. The SVM training set was based on training data that was labeled by a human observer and this technique yielded >95% accuracy in training data. The signals were then recomposed from the signal components excluding the noise components resulting in a largely artifact free signal.

Coherence calculations

We calculated pairwise coherences between all electrode pairs in 10-second non-overlapping intervals. The coherence between two time varying signals x and y was defined as

$$C_{xy} = \frac{P_{xy}^2}{P_{xx}P_{yy}}$$

where P_{xy} is the cross spectral density and P_{xx} / P_{yy} are the auto-spectral densities. From this measurement we subtracted an estimated noise floor by computing the coherence of a phase-randomized signal generated by computing the discrete fourier transform or decomposing the signal into the form of

$$X(f) = \sum_{n=0}^{N-1} x(t_n) e^{2\pi i f n \Delta t}$$

This form can also be rewritten as

$$X(f) = A(f) e^{i\phi(f)}$$

where $A(f)$ is the amplitude and $\phi(f)$ is the phase component of each frequency component f . A 'phase-randomized' version of this signal would simply be permuting all of the $\phi(f)$ for each frequency and assigning them randomly across frequency, then computing the inverse fourier transform, yielding a signal with the same distribution of $A(f)$ across frequencies but

a new random $\phi(f)$ and thus controlling for any residual coherence expected from signals of the corresponding power spectra.

Dimensionality reduction/predictability analysis

Coherences computed across all electrode pairs were then averaged within each region to reduce the dimensionality from $n\text{Electrodes} \times n\text{Electrodes}$, to $n\text{Regions} \times n\text{Regions}$, where the entry for each region to region interaction represents the mean coherence across all electrode pairs from the two regions. This is also diagrammed in figure 5.1. We then subdivided these coherence measurements into several frequency bands of interest namely delta (δ , 1-4 Hz), theta (θ , 4-7 Hz), alpha (α , 7-15 Hz), beta (β , 15-30 Hz), low gamma (γ , 30-70 Hz), and high gamma (Γ , 70-120 Hz). This resulted in time series that were $n\text{Regions} \times n\text{Regions} \times \text{Time}$ and PCA was performed on this resulting representation for each frequency band separately. Coherences were mean-subtracted prior to performing PCA. Principal components analysis was carried out using the singular value decomposition (SVD) built into the sklearn python package. Shuffled surrogates were generated by subdividing the coherence measurements into 5 separate chunks of random length and then permuting the chunks relative to each other as diagrammed in figure 5.1.

We defined the projection onto each principal component as the linear projection of the coherence matrix onto the axis defined by each principal component, this represents the sum total activity along that axis at a given point in time. Predictability of each coherence against the total projection was computed as the square of the pearson correlation coefficient

between the Z scored versions of the projection and the coherence at a given point in time.

The pearson correlation value is defined as

$$\rho = \frac{A \cdot B}{\|A\| \|B\|}$$

where $A \cdot B$ is the dot product and $\|A\|$ and $\|B\|$ are the magnitudes of those two signals.

5.5 Figures

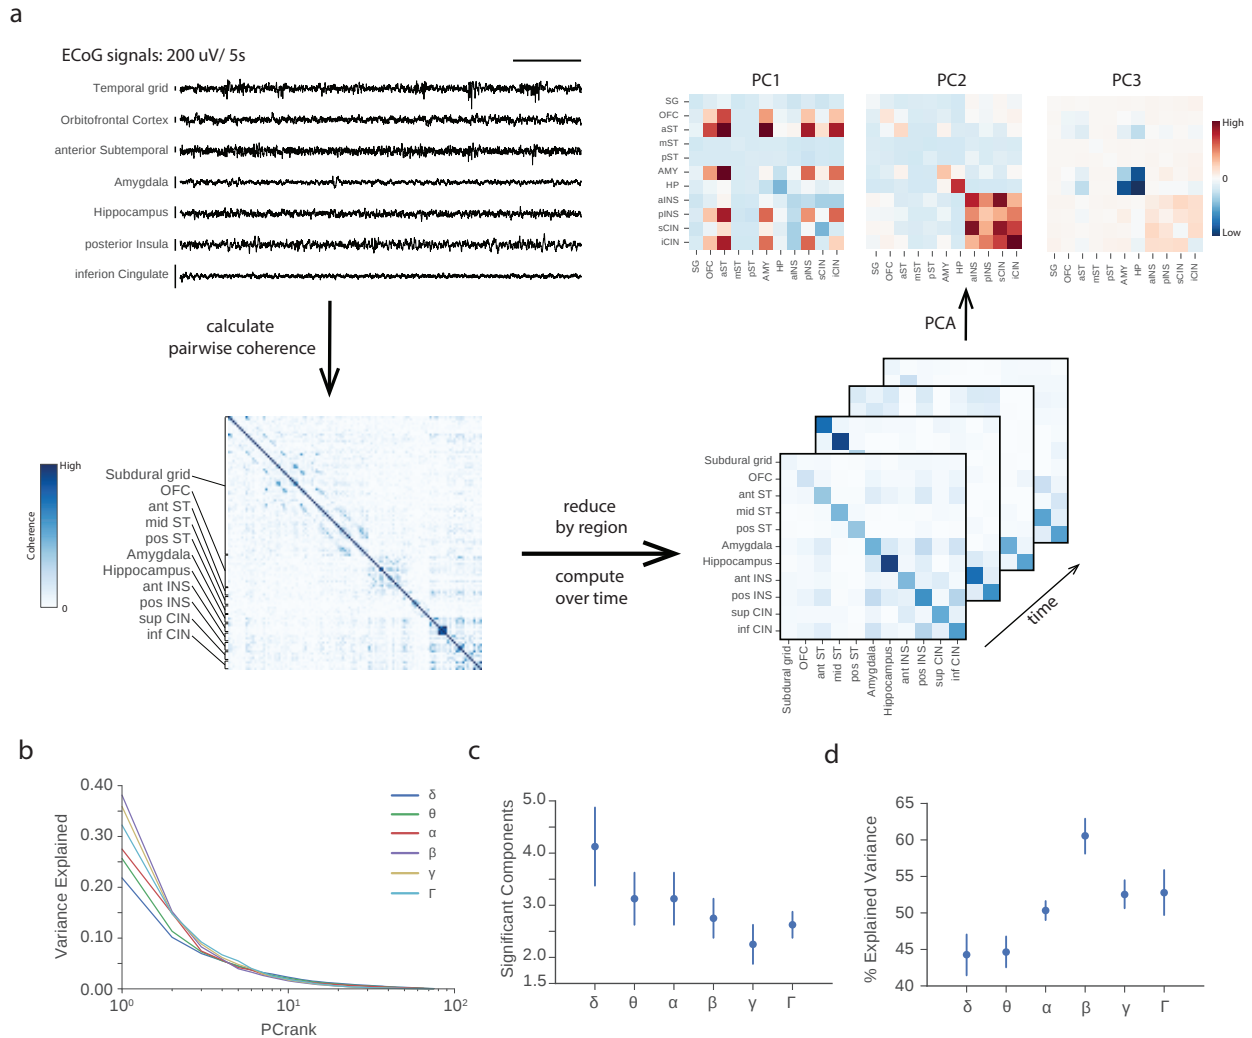


Figure 5.1: **Extracting network wide structure in ECoG data use pairwise coherence and dimensionality reduction.**

a, Schematic diagramming the workflow for processing of neural ecog data into principal components. Pairwise coherences are calculated from filtered and artifact-removed ECoG data across all possible electrode pairs in 10 second intervals (bottom left). This large representation is then reduced by grouping electrodes from the same grid or depth electrode as representing the same point in space (bottom right). PCA is performed on that resulting reduced dataset in different frequency bands resulting in example PCs like the ones shown (top right). **b**, Percentage of variance explained for each PC as function of rank of the PC. Each line corresponds to the mean across 8 patients and colored according to frequency band as indicated in the legend. **c**, Number of significant components, defined as the number of PCs with eigenvalues above that expected from shuffled data as a function of frequency band. On average, data exhibits about 3-4 significant PCs. **d**, Percentage of variance explained by significant components as a function of frequency bands. On average about 50

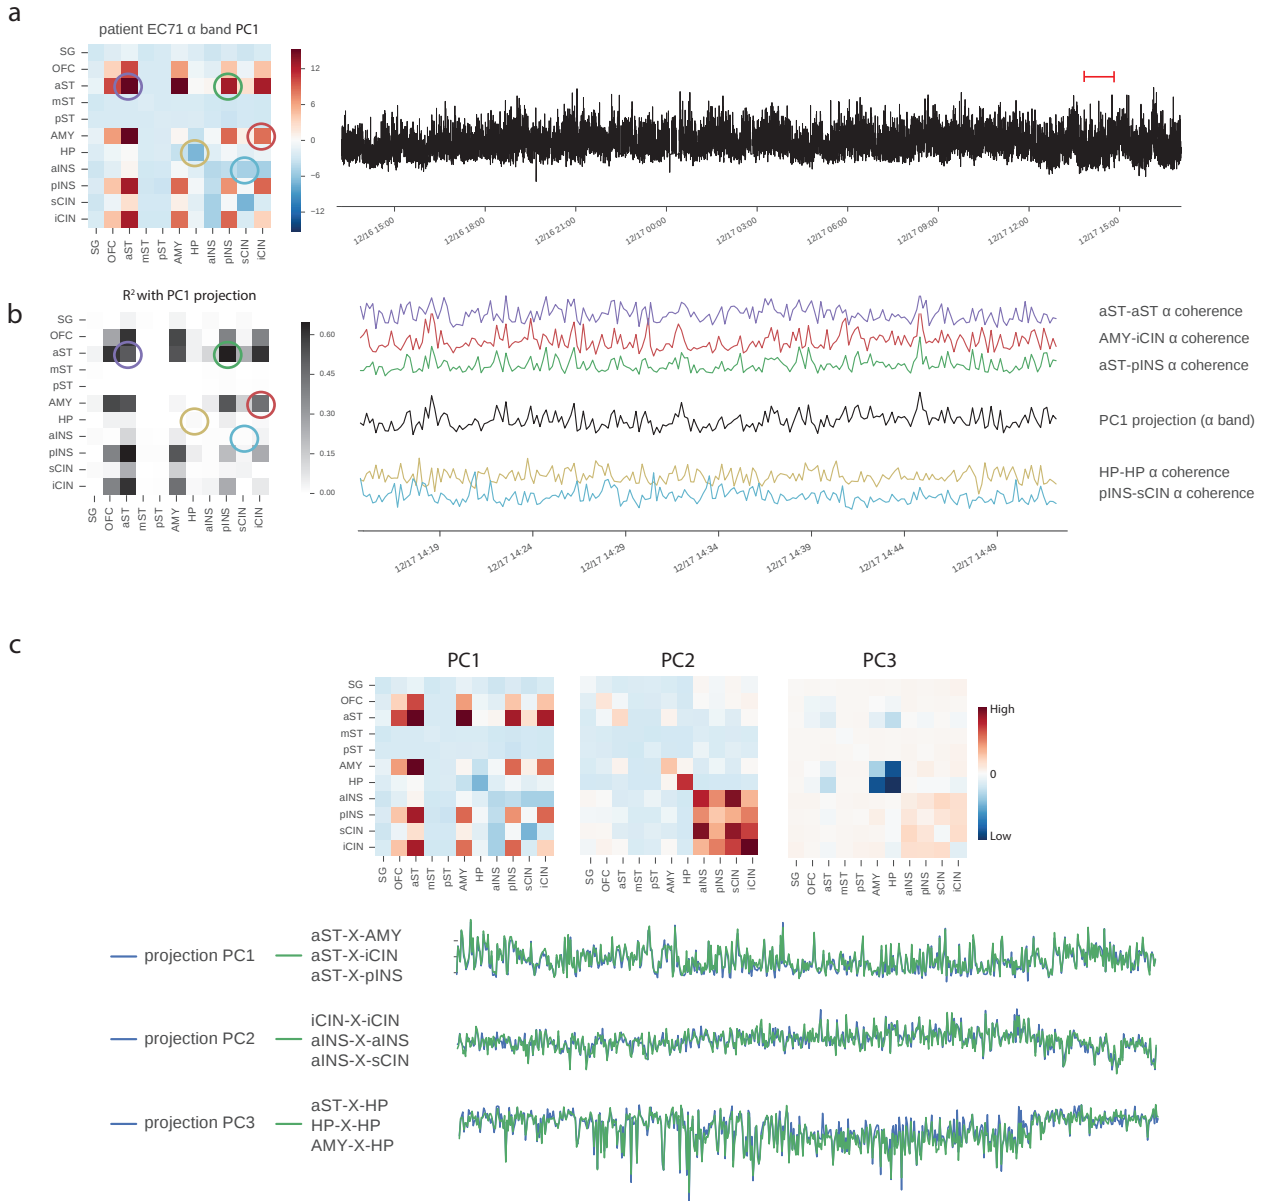


Figure 5.2: Analyzing principal component projections and extracting lower dimensional representations.

a, PC1 from alpha band of EC71 with accompanying linear projection of coherences onto that PC, i.e. a representation of what is happening in those groups of interactions at any given point in time. Red demarcation denotes time period blown up in part B. **b**, R^2 value of the coherence value of each given interaction with the PC1 projection. Example traces showing 3 highly correlated traces (top; aST-aST, AMY-iCIN, aST-pINS) and 2 uncorrelated traces (bottom; HP-HP, pINS-sCIN) with the PC1 projection (black) **c**, Top: Principal components (from left to right) PC1, PC2, and PC3 of alpha band coherences. Bottom: Reconstruction of projections onto (from top to bottom) PC1, PC2, and PC3 using just the top 3 most informative key interactions for that given PC

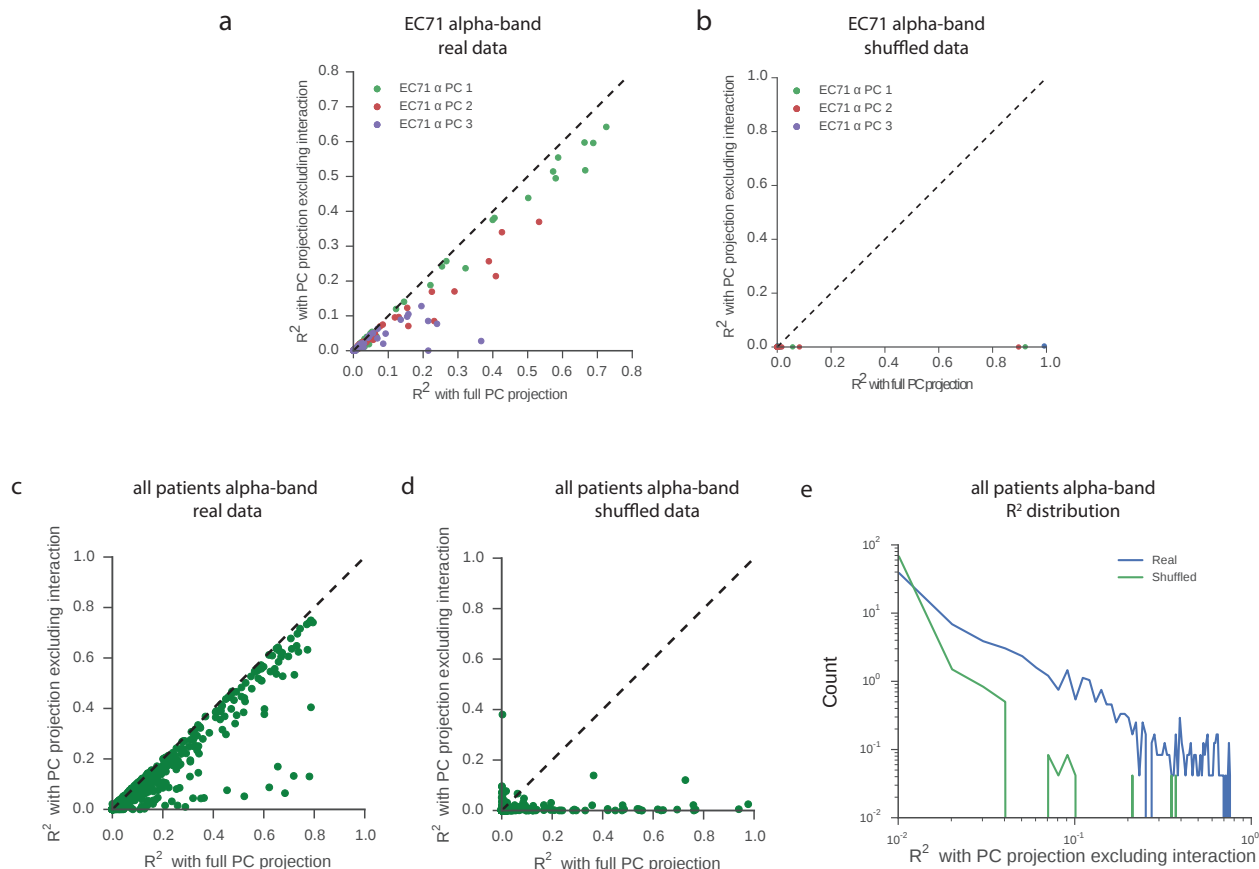


Figure 5.3: Key interactions present across patients.

a, R^2 value of each coherence interaction against the projection onto PC1 (green), PC2 (red), and PC3 (purple) for either the full projection on the Y axis, or a version of the projection excluding that interaction on the X axis. Most points that have high Y values initially, maintain high X values. **b**, Same plot as in A except for shuffled surrogates. Interactions with high R^2 against the original projection, have no predictability against projections excluding that interaction. Highlighting that shuffled data loses the redundancy observed in real data. **c**, Same plot as in A but for the first 3 PCs across all patients in the alpha band. **d**, Same plot as in B but for the first 3 PCs across all patients in the alpha band. **e**, Distribution of X axis values from C (Blue), and the X axis values from D (green), highlighting that the redundancy and this high predictability against the projections excluding that interaction are only present in real data.

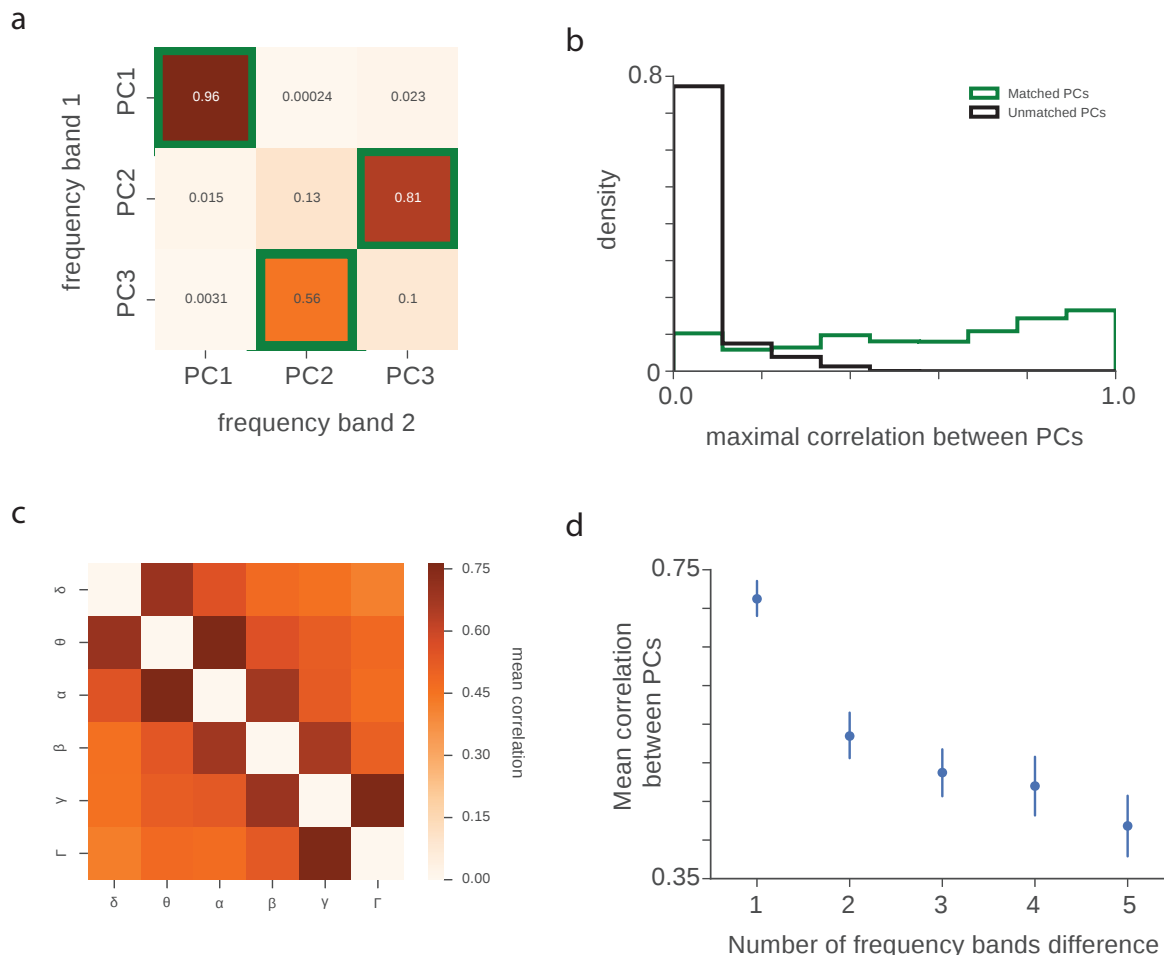


Figure 5.4: **Subnetworks identified via PCs are frequency specific.**

a, Schematic illustrating how a top 3 PC from one frequency band can be matched to its equivalent in another frequency band by computing the correlation between the two components. PCs that are determined to be the same and considered matched are outlined in green and represent the largest correlation between a seed PC and the top 3 PCs from another frequency band. **b**, Distribution of R2 values for all of the matched (green) or unmatched (black) PCs. The distribution of matched vs. non-matched PCs are largely non-overlapping. **c**, Mean of the correlation between Matched PCs from each of the 6 frequency bands. Stronger correlations are viewed towards the diagonal, indicating that closer frequency bands share more similar PCs. **d**, Quantification of mean correlation between matched PCs as a function of either 1,2,3, 4, or 5 frequency bands of difference. PCs computed from neighboring frequency bands more closely resemble each other than those computed from distant frequency bands.

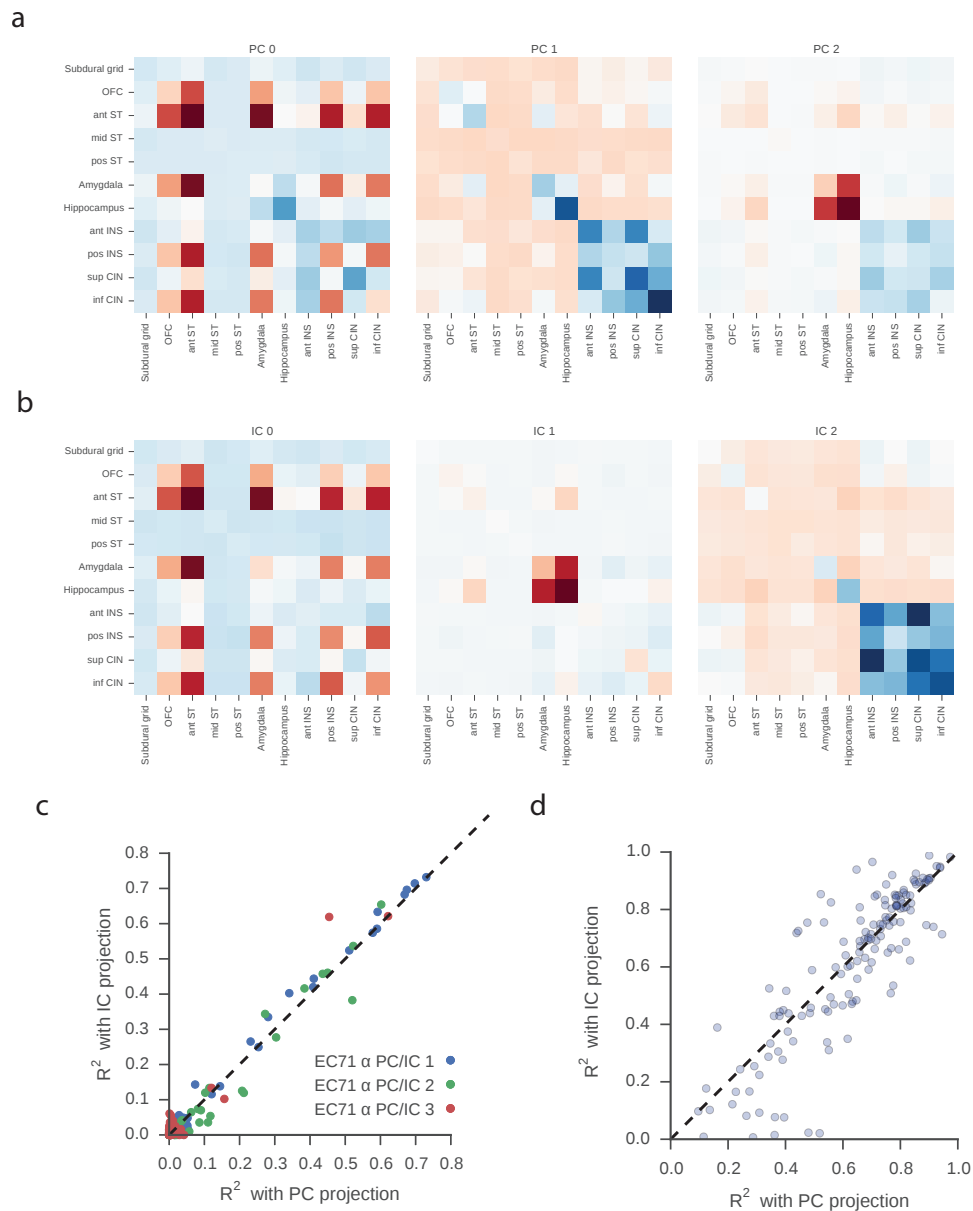



Figure 5.5: **PCA and ICA yield similar subnetworks.**

a, Example first 3 principal components computed from patient EC79 in the alpha band **b**, Example first 3 independent components computed from patient EC79 in the alpha band **c**, Mean correlation of interactions with the top 3 principal components (X axis) and the mean correlation with the projection of the closest matching independent component (Y axis) for patient EC79 in the alpha band. Due to the similarity between principal components and independent components, interactions with high correlation to one, retain a high correlation to the other. **d**, Same plot as in C, except showing only the top 3 key interactions as determined from the principal components, and their correlation with the most closely matched independent components (Y axis). Data plotted represents top 3 key interactions for each of the top 3 principal components across all frequencies and patients.


Real Coherence



Shuffled Surrogate


Supplementary Figure 5.1: **Method for generating shuffled surrogates.**

Schematic for generating shuffled surrogates. Coherences are divided into 5 epochs of random length and then randomly permuted resulting in a coherence surrogate that has the same mean, variance, and autocorrelation but any cross-correlation is broken.



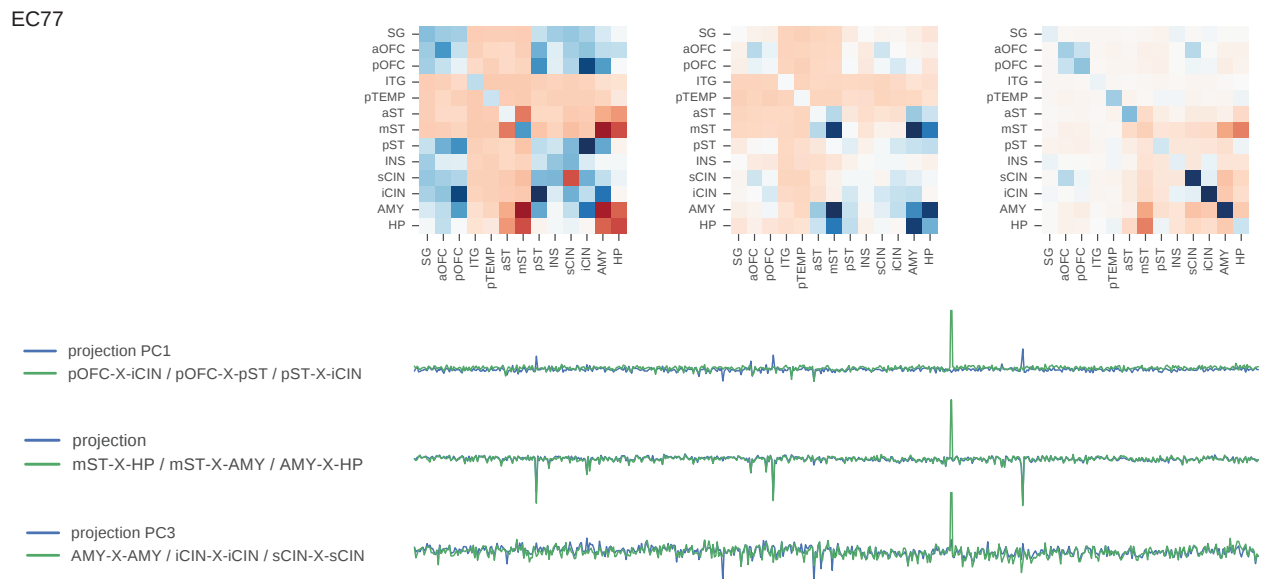
Supplementary Figure 5.2: **Eigenvalue distributions of real vs. shuffled data across patients.**

Eigenvalues of real data (black) vs. shuffled (green). All traces are normalized by the top eigenvalue from real data for each individual experiment. Shading represents SEM. Each panel corresponds to the distributions computed from a given frequency band (from left to right): delta (1-3 Hz), theta (4-7 Hz), alpha (8-15 Hz), beta (15-30 Hz), low gamma (30-70 Hz), and high gamma (70-120 Hz). In most cases 3-4 PCs have eigenvalues above that expected from shuffled surrogates.



Supplementary Figure 5.3: **Descriptive statistics for patient EC71 alpha band data.**

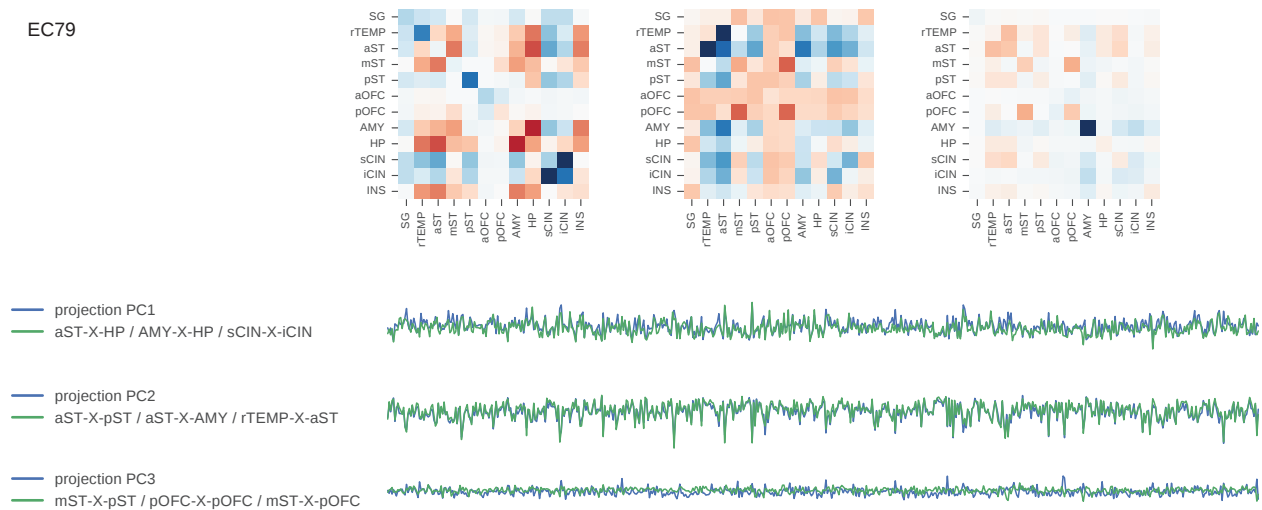
Left: Mean of coherence for each given interaction corresponding to data presented in figure 2, **Right:** Variance of coherence for each given interaction corresponding to data presented in figure 2



Supplementary Figure 5.4: **Example alpha band PC data from EC77.**

Top: Principal components (from left to right) PC1, PC2, and PC3 of alpha band coherences. **Bottom:** Reconstruction of projections onto (from top to bottom) PC1, PC2, and PC3 using just the top 3 most informative key interactions for that given PC

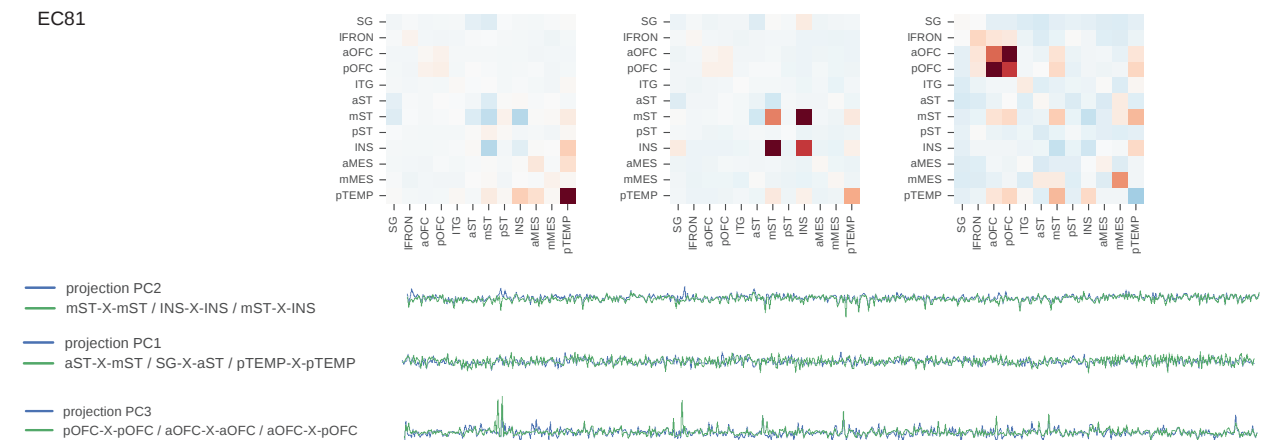
EC79



Supplementary Figure 5.5: **Example alpha band PC data from EC80.**

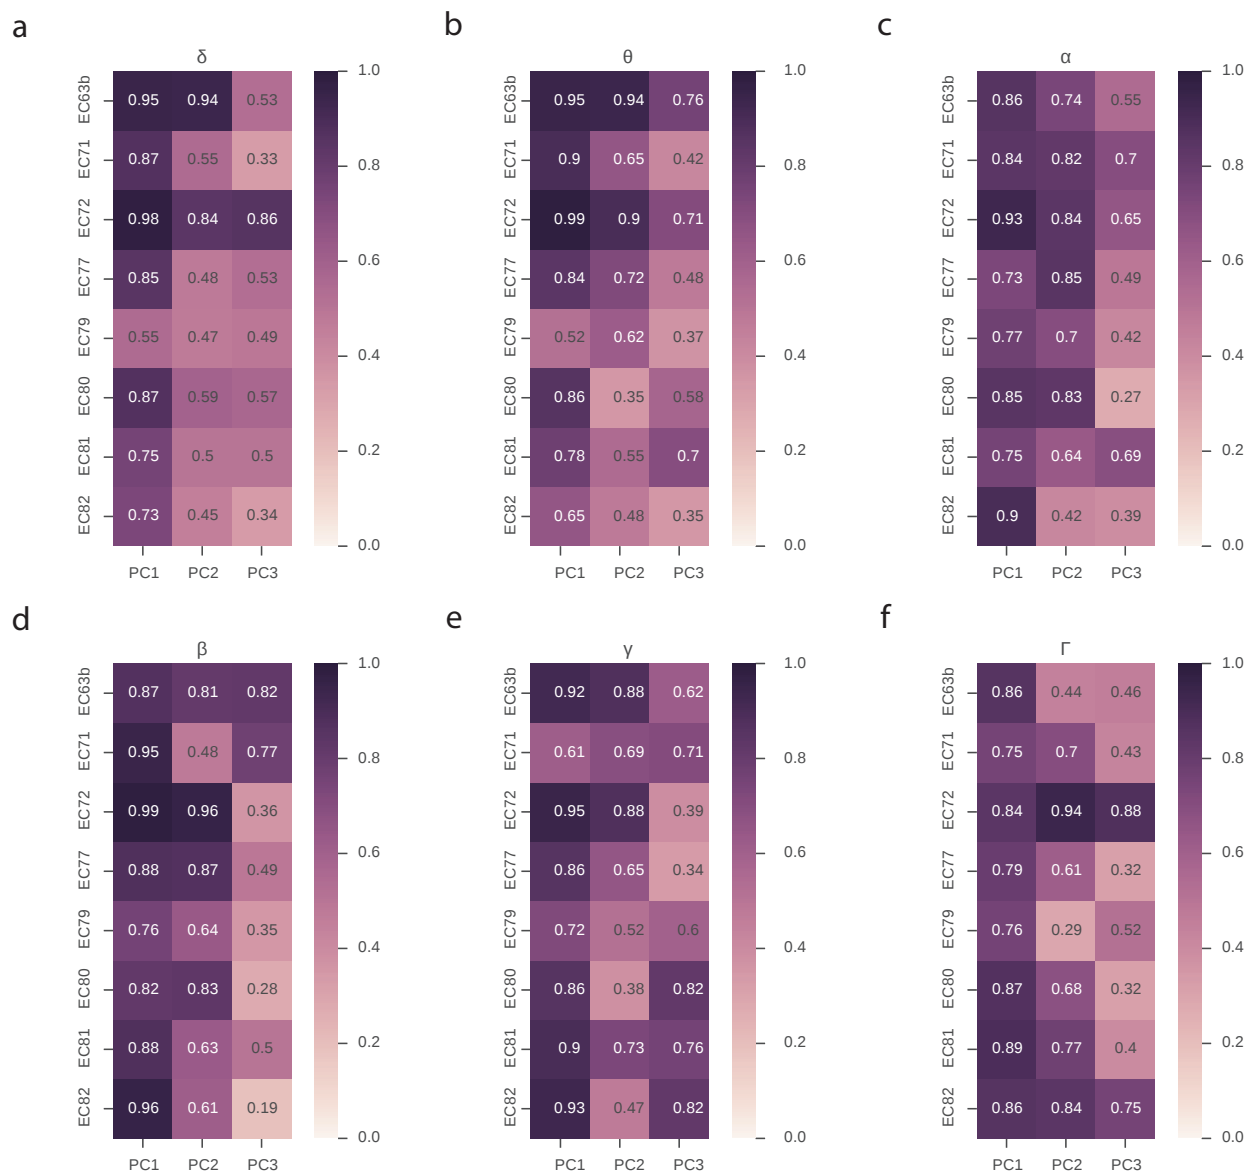
Top: Principal components (from left to right) PC1, PC2, and PC3 of alpha band coherences.
Bottom: Reconstruction of projections onto (from top to bottom) PC1, PC2, and PC3 using just the top 3 most informative key interactions for that given PC

EC81



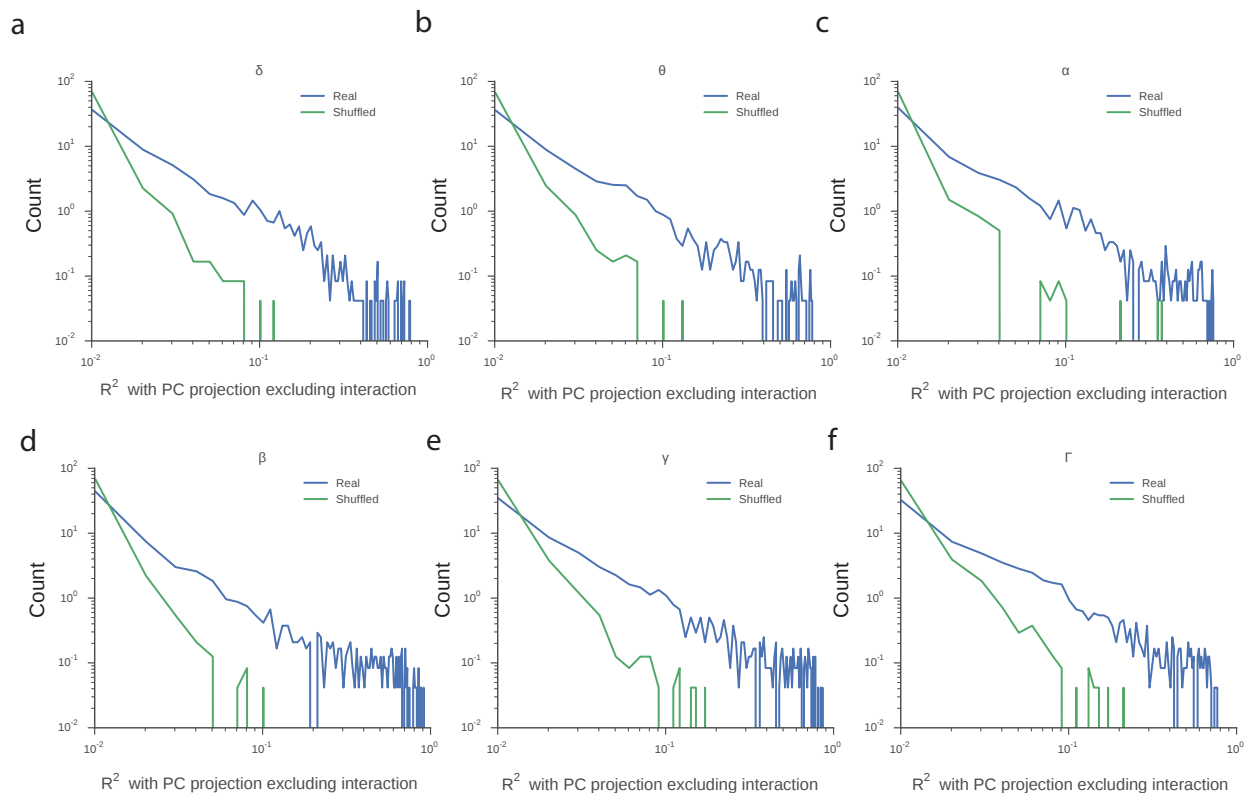
Supplementary Figure 5.6: **Example alpha band PC data from EC81.**

Top: Principal components (from left to right) PC1, PC2, and PC3 of alpha band coherences.
Bottom: Reconstruction of projections onto (from top to bottom) PC1, PC2, and PC3 using just the top 3 most informative key interactions for that given PC



Supplementary Figure 5.7: **Explained variance from top 3 key interactions in each principal component across frequencies and patients.**

a, R^2 value of the original PC projection with that generated using only the top 3 interactions for a given PC in the delta band (1-3 Hz) **b**, R^2 value of the original PC projection with that generated using only the top 3 interactions for a given PC in the theta band (4-7 Hz) **c**, R^2 value of the original PC projection with that generated using only the top 3 interactions for a given PC in the alpha band (7-15 Hz) **d**, R^2 value of the original PC projection with that generated using only the top 3 interactions for a given PC in the beta band (15-30 Hz) **e**, R^2 value of the original PC projection with that generated using only the top 3 interactions for a given PC in the low gamma band (30-70 Hz) **f**, R^2 value of the original PC projection with that generated using only the top 3 interactions for a given PC in the high gamma band (70-120 Hz)



Supplementary Figure 5.8: **Key interactions present across multiple frequencies.**

a, Distribution of R^2 values of each interaction against a projection excluding that interaction for real data (blue) and shuffled data (green) in the delta frequency band **b**, Distribution of R^2 values of each interaction against a projection excluding that interaction for real data (blue) and shuffled data (green) in the theta frequency band **c**, Distribution of R^2 values of each interaction against a projection excluding that interaction for real data (blue) and shuffled data (green) in the alpha frequency band **d**, Distribution of R^2 values of each interaction against a projection excluding that interaction for real data (blue) and shuffled data (green) in the beta frequency band **e**, Distribution of R^2 values of each interaction against a projection excluding that interaction for real data (blue) and shuffled data (green) in the low gamma frequency band **f**, Distribution of R^2 values of each interaction against a projection excluding that interaction for real data (blue) and shuffled data (green) in the high gamma frequency band

5.6 Tables

patient	Long region name	Short region name	patient	Long region name	Short region name
EC63	Subdural grid	SG	EC79	Subdural grid	SG
	inf Frontal	iFRON		right Temporal	rTEMP
	lat Temporal	ITEMP		ant ST	aST
	OFC	OFC		mid ST	mST
	ant ST	aST		pos ST	pST
	mid ST	mST		ant OFC	aOFC
	pos ST	pST		pos OFC	pOFC
	inf CIN	iCIN		Amygdala	AMY
	sup CIN	sCIN		Hippocampus	HP
	INS	INS		sup CIN	sCIN
	Amygdala	AMY		inf CIN	iCIN
	Hippocampus	HP		INS	INS
EC71	Subdural grid	SG	EC80	Amygdala	AMY
	OFC	OFC		ER-Ctx	ER-CTX
	ant ST	aST		Hipp-head	Hipp-head
	mid ST	mST		Hipp-body	Hipp-body
	pos ST	pST		Hipp-tail	Hipp-tail
	Amygdala	AMY		Cuneus	CUNEUS
	Hippocampus	HP		Calcarine	CAL
	ant INS	aINS		inf CIN	iCIN
	pos INS	pINS		INS	INS
	sup CIN	sCIN		sup CIN	sCIN
	inf CIN	iCIN		OFC	OFC
	EC72	Subdural grid		SG	EC81
ant OFC		aOFC	lat Frontal	iFRON	
pos OFC		pOFC	ant OFC	aOFC	
ST A		ST A	pos OFC	pOFC	
ST B		ST B	ITG	ITG	
ST C		ST C	ant ST	aST	
ST D		ST D	mid ST	mST	
Amygdala		AMY	pos ST	pST	
Posterior depth		PD	INS	INS	
EC77		Subdural grid	SG	EC82	
	ant OFC	aOFC	ant Frontal		aFrontal
	pos OFC	pOFC	ant HP		aHP
	ITG	ITG	pos HP		pHP
	pos Temporal	pTEMP	ant ST		aST
	ant ST	aST	mid ST		mST
	mid ST	mST	pos ST		pST
	pos ST	pST	ITG		ITG
	INS	INS	sup OFC		sOFC
	sup CIN	sCIN	inf OFC		iOFC
	inf CIN	iCIN	med Frontal		mFrontal
	Amygdala	AMY			
Hippocampus	HP				

Table 5.1: Summary of ECoG coverage in each patient

pID	BDI	BAI
EC63b	36	22
EC71	9	7
EC72	5	6
EC77	5	7
EC79	28	8
EC80	8	2
EC81	18	16
EC82	?	?

Table 5.2: Summary of the baseline depression index (BDI) and the baseline anxiety index (BAI) for each patient. Values between 0-10 are considered normal, 10-20 are considerate moderate depression/anxiety, and over 20 is considered severe anxiety/depression

patient	Delta		
	PC1	PC2	PC3
EC63b	{SG -X- SG} {AMY -X- AMY} {AMY -X- HP}	{INS -X- HP} {iCIN -X- iCIN} {iCIN -X- INS}	{iCIN -X- sCIN} {AMY -X- AMY} {HP -X- HP}
EC71	{AMY -X- iCIN} {AMY -X- pINS} {pINS -X- iCIN}	{iCIN -X- iCIN} {sCIN -X- sCIN} {aINS -X- sCIN}	{aINS -X- aINS} {aINS -X- sCIN} {sCIN -X- sCIN}
EC72	{aOFC -X- pOFC} {PD -X- PD} {pOFC -X- pOFC}	{ST C -X- PD} {PD -X- PD} {AMY -X- PD}	{SG -X- pOFC} {aOFC -X- aOFC} {aOFC -X- pOFC}
EC77	{pOFC -X- sCIN} {sCIN -X- iCIN} {pOFC -X- iCIN}	{pOFC -X- pST} {pOFC -X- AMY} {sCIN -X- sCIN}	{aST -X- AMY} {mST -X- AMY} {AMY -X- HP}
EC79	{sCIN -X- iCIN} {pST -X- sCIN} {aST -X- sCIN}	{aOFC -X- iCIN} {iCIN -X- iCIN} {sCIN -X- iCIN}	{mST -X- mST} {pOFC -X- pOFC} {mST -X- pOFC}
EC80	{CAL -X- sCIN} {CAL -X- OFC} {iCIN -X- sCIN}	{OFC -X- LIN-GYR} {INS -X- LIN-GYR} {CAL -X- CAL}	{AMY -X- HP-head} {HP-head -X- HP-head} {AMY -X- AMY}
EC81	{pTEMP -X- pTEMP} {INS -X- pTEMP} {aMES -X- pTEMP}	{INS -X- mMES} {aST -X- mMES} {mST -X- mMES}	{aOFC -X- aOFC} {IFRON -X- pOFC} {aOFC -X- pOFC}
EC82	{pHP -X- pHP} {aHP -X- aHP} {aHP -X- pHP}	{aFrontal -X- sOFC} {aFrontal -X- ITG} {aFrontal -X- mFrontal}	{iOFC -X- iOFC} {aFrontal -X- mST} {iOFC -X- mFrontal}

Table 5.3: Top 3 key interactions, ordered from top to bottom for each patient for each of the top 3 principal components in the Delta (1-3Hz) band

patient	Theta		
	PC1	PC2	PC3
EC63b	{SG -X- SG} {AMY -X- AMY} {AMY -X- HP}	{INS -X- HP} {HP -X- HP} {iCIN -X- INS}	{INS -X- INS} {iCIN -X- INS} {HP -X- HP}
EC71	{aST -X- iCIN} {aST -X- pINS} {pINS -X- iCIN}	{aINS -X- aINS} {aINS -X- sCIN} {sCIN -X- sCIN}	{HP -X- HP} {HP -X- iCIN} {AMY -X- HP}
EC72	{aOFC -X- aOFC} {aOFC -X- pOFC} {pOFC -X- pOFC}	{ST C -X- PD} {PD -X- PD} {AMY -X- PD}	{ST B -X- AMY} {aOFC -X- pOFC} {aOFC -X- aOFC}
EC77	{SG -X- sCIN} {pST -X- iCIN} {pOFC -X- iCIN}	{mST -X- HP} {AMY -X- HP} {mST -X- AMY}	{sCIN -X- sCIN} {sCIN -X- AMY} {iCIN -X- iCIN}
EC79	{aST -X- sCIN} {sCIN -X- iCIN} {pST -X- sCIN}	{rTEMP -X- AMY} {aST -X- AMY} {rTEMP -X- aST}	{iCIN -X- INS} {iCIN -X- iCIN} {sCIN -X- iCIN}
EC80	{sCIN -X- OFC} {iCIN -X- OFC} {iCIN -X- sCIN}	{AMY -X- Hipp-head} {CAL -X- iCIN} {CAL -X- CAL}	{CUNEUS -X- CUNEUS} {AMY -X- AMY} {Hipp-head -X- Hipp-head}
EC81	{INS -X- pTEMP} {aMES -X- pTEMP} {pTEMP -X- pTEMP}	{mST -X- mST} {INS -X- INS} {mST -X- INS}	{pOFC -X- pOFC} {aOFC -X- aOFC} {aOFC -X- pOFC}
EC82	{pHP -X- aST} {aHP -X- pHP} {pHP -X- mST}	{aHP -X- pHP} {aFrontal -X- iOFC} {iOFC -X- mFrontal}	{aST -X- mST} {iOFC -X- mFrontal} {aFrontal -X- iOFC}

Table 5.4: Top 3 key interactions, ordered from top to bottom for each patient for each of the top 3 principal components in the Theta (4-7Hz) band

Alpha			
patient	PC1	PC2	PC3
EC63b	{HP -X- HP}	{HP -X- HP}	{INS -X- INS}
	{AMY -X- AMY}	{iCIN -X- iCIN}	{AMY -X- AMY}
	{AMY -X- HP}	{iCIN -X- INS}	{HP -X- HP}
EC71	{aST -X- AMY}	{iCIN -X- iCIN}	{aST -X- HP}
	{aST -X- iCIN}	{aINS -X- aINS}	{HP -X- HP}
	{aST -X- pINS}	{aINS -X- sCIN}	{AMY -X- HP}
EC72	{ST C -X- AMY}	{ST C -X- AMY}	{ST A -X- ST C}
	{aOFC -X- pOFC}	{ST C -X- PD}	{aOFC -X- aOFC}
	{pOFC -X- pOFC}	{AMY -X- PD}	{aOFC -X- pOFC}
EC77	{pOFC -X- iCIN}	{mST -X- HP}	{AMY -X- AMY}
	{pOFC -X- pST}	{mST -X- AMY}	{iCIN -X- iCIN}
	{pST -X- iCIN}	{AMY -X- HP}	{sCIN -X- sCIN}
EC79	{aST -X- HP}	{aST -X- pST}	{mST -X- pST}
	{AMY -X- HP}	{aST -X- AMY}	{pOFC -X- pOFC}
	{sCIN -X- iCIN}	{rTEMP -X- aST}	{mST -X- pOFC}
EC80	{sCIN -X- OFC}	{HP-head -X- HP-head}	{CAL -X- sCIN}
	{iCIN -X- OFC}	{AMY -X- AMY}	{CAL -X- CAL}
	{iCIN -X- sCIN}	{AMY -X- HP-head}	{CAL -X- iCIN}
EC81	{mST -X- mST}	{aST -X- mST}	{pOFC -X- pOFC}
	{INS -X- INS}	{SG -X- aST}	{aOFC -X- aOFC}
	{mST -X- INS}	{pTEMP -X- pTEMP}	{aOFC -X- pOFC}
EC82	{aHP -X- mST}	{iOFC -X- iOFC}	{iOFC -X- mFrontal}
	{pHP -X- mST}	{aFrontal -X- iOFC}	{aST -X- aST}
	{aHP -X- pHP}	{iOFC -X- mFrontal}	{mST -X- mST}

Table 5.5: Top 3 key interactions, ordered from top to bottom for each patient for each of the top 3 principal components in the Alpha (8-15Hz) band

patient	Beta		
	PC1	PC2	PC3
EC63b	{AMY -X- AMY} {iCIN -X- iCIN} {iCIN -X- INS}	{iCIN -X- INS} {AMY -X- AMY} {AMY -X- HP}	{AMY -X- HP} {AMY -X- AMY} {HP -X- HP}
EC71	{HP -X- HP} {AMY -X- AMY} {AMY -X- HP}	{aST -X- AMY} {OFC -X- aST} {aST -X- aST}	{aINS -X- aINS} {aST -X- aST} {iCIN -X- iCIN}
EC72	{ST C -X- PD} {ST C -X- AMY} {AMY -X- PD}	{PD -X- PD} {aOFC -X- aOFC} {pOFC -X- pOFC}	{ST B -X- ST C} {AMY -X- AMY} {ST B -X- PD}
EC77	{mST -X- HP} {mST -X- AMY} {AMY -X- HP}	{SG -X- sCIN} {pST -X- iCIN} {pOFC -X- iCIN}	{iCIN -X- iCIN} {AMY -X- AMY} {sCIN -X- sCIN}
EC79	{AMY -X- HP} {aST -X- HP} {pST -X- HP}	{pST -X- iCIN} {aST -X- pST} {pST -X- sCIN}	{rTEMP -X- pOFC} {pOFC -X- pOFC} {mST -X- pOFC}
EC80	{iCIN -X- OFC} {iCIN -X- INS} {iCIN -X- sCIN}	{AMY -X- AMY} {Hipp-head -X- Hipp-head} {AMY -X- Hipp-head}	{CAL -X- sCIN} {CAL -X- CAL} {CAL -X- iCIN}
EC81	{mST -X- mST} {INS -X- INS} {mST -X- INS}	{mST -X- mMES} {mMES -X- mMES} {mMES -X- pTEMP}	{aOFC -X- pOFC} {pOFC -X- pOFC} {pTEMP -X- pTEMP}
EC82	{aHP -X- pHP} {aHP -X- mST} {pHP -X- mST}	{iOFC -X- iOFC} {pST -X- pST} {iOFC -X- mFrontal}	{pHP -X- aST} {aST -X- mST} {aST -X- aST}

Table 5.6: Top 3 key interactions, ordered from top to bottom for each patient for each of the top 3 principal components in the Beta (15-30Hz) band

patient	Gamma		
	PC1	PC2	PC3
EC63b	{iCIN -X- iCIN} {AMY -X- AMY} {iCIN -X- INS}	{iCIN -X- iCIN} {iCIN -X- INS} {AMY -X- AMY}	{OFC -X- AMY} {INS -X- AMY} {iCIN -X- AMY}
EC71	{pINS -X- iCIN} {sCIN -X- iCIN} {iCIN -X- iCIN}	{AMY -X- AMY} {HP -X- HP} {AMY -X- HP}	{mST -X- pST} {mST -X- pINS} {pST -X- pINS}
EC72	{ST C -X- PD} {ST C -X- AMY} {AMY -X- PD}	{AMY -X- AMY} {PD -X- PD} {pOFC -X- pOFC}	{AMY -X- AMY} {SG -X- SG} {PD -X- PD}
EC77	{SG -X- sCIN} {mST -X- mST} {pOFC -X- iCIN}	{AMY -X- HP} {AMY -X- AMY} {HP -X- HP}	{sCIN -X- sCIN} {mST -X- iCIN} {mST -X- INS}
EC79	{pST -X- iCIN} {pST -X- sCIN} {sCIN -X- iCIN}	{aST -X- HP} {HP -X- HP} {pST -X- HP}	{rTEMP -X- pOFC} {pOFC -X- pOFC} {mST -X- pOFC}
EC80	{iCIN -X- sCIN} {CAL -X- OFC} {iCIN -X- OFC}	{INS -X- LIN-GYR} {CAL -X- CAL} {CAL -X- iCIN}	{AMY -X- AMY} {HP-head -X- HP-head} {AMY -X- Hipp-head}
EC81	{mST -X- mST} {INS -X- INS} {mST -X- INS}	{INS -X- pTEMP} {pTEMP -X- pTEMP} {mST -X- pTEMP}	{SG -X- pTEMP} {SG -X- mST} {mMES -X- mMES}
EC82	{aHP -X- pHP} {aHP -X- mST} {pHP -X- mST}	{iOFC -X- iOFC} {iOFC -X- mFrontal} {mFrontal -X- mFrontal}	{sOFC -X- sOFC} {pST -X- pST} {ITG -X- ITG}

Table 5.7: Top 3 key interactions, ordered from top to bottom for each patient for each of the top 3 principal components in the low Gamma (30-70Hz) band

patient	GammaHi		
	PC1	PC2	PC3
EC63b	{SG -X- AMY} {iCIN -X- INS} {AMY -X- AMY}	{iCIN -X- iCIN} {INS -X- INS} {iCIN -X- INS}	{INS -X- AMY} {OFC -X- iCIN} {iCIN -X- AMY}
EC71	{mST -X- pST} {mST -X- pINS} {pST -X- pINS}	{sCIN -X- iCIN} {sCIN -X- sCIN} {iCIN -X- iCIN}	{HP -X- HP} {SG -X- sCIN} {AMY -X- sCIN}
EC72	{AMY -X- AMY} {ST C -X- AMY} {AMY -X- PD}	{ST C -X- PD} {AMY -X- PD} {pOFC -X- pOFC}	{AMY -X- PD} {PD -X- PD} {SG -X- SG}
EC77	{SG -X- sCIN} {SG -X- iCIN} {pOFC -X- iCIN}	{sCIN -X- sCIN} {HP -X- HP} {AMY -X- AMY}	{mST -X- pST} {mST -X- iCIN} {mST -X- INS}
EC79	{pST -X- sCIN} {sCIN -X- sCIN} {sCIN -X- iCIN}	{rTEMP -X- aOFC} {pOFC -X- AMY} {aOFC -X- AMY}	{mST -X- mST} {rTEMP -X- mST} {mST -X- pOFC}
EC80	{CAL -X- OFC} {iCIN -X- sCIN} {iCIN -X- OFC}	{OFC -X- LIN-GYR} {INS -X- LIN-GYR} {CAL -X- iCIN}	{sCIN -X- LIN-GYR} {CAL -X- sCIN} {sCIN -X- sCIN}
EC81	{SG -X- INS} {INS -X- INS} {mST -X- INS}	{INS -X- pTEMP} {pTEMP -X- pTEMP} {mST -X- pTEMP}	{SG -X- pTEMP} {SG -X- mST} {mST -X- mST}
EC82	{aHP -X- aHP} {aHP -X- mST} {aHP -X- pHP}	{aFrontal -X- ITG} {aFrontal -X- aFrontal} {ITG -X- ITG}	{mFrontal -X- mFrontal} {ITG -X- ITG} {aFrontal -X- aFrontal}

Table 5.8: Top 3 key interactions, ordered from top to bottom for each patient for each of the top 3 principal components in the high Gamma (70-120Hz) band

Chapter 6

Concluding remarks

Population calcium imaging studies

The studies outlined here reveal some novel ways in which neural circuits can represent and process information. In the first 3 studies we were able to exploit the high spatial and temporal resolution afforded by calcium imaging to study statistical patterns at the level of individual neurons. These studies inform us about the emergent representations generated by neuromodulators such as dopamine and acetylcholine. Our preparation using acute prefrontal slices is advantageous in that we can isolate specific microcircuits and assay various neuromodulatory conditions in isolation. The major downside of this preparation is that we do not know for sure how well these patterns observed might be mimicked in the behaving animal. Indeed temporally stereotyped sequences of activation similar to those observed have been shown in-vivo (Fujisawa et al., 2008; Harvey et al., 2012). In addition, a cholinergic dependent decorrelation can be elicited in visual cortex in the awake mouse

(Goard and Dan, 2009). Thus future work will hopefully leverage large-scale recordings in the behaving animal to observe these neuromodulatory effect in-vivo and ascertain how these emergent properties might contribute to behavior.

Despite this limitation, an in-vitro assay as well as a recapitulation of the patterns and statistics observed in-vivo opens up the exploration of the specific mechanisms these computations. In much the same way that slice electrophysiology has allowed for the careful dissection of cell-intrinsic biophysics, similar studies could be used to elucidate the cell-autonomous and network interactions underlying the cholinergic dependent decorrelation. Many modeling studies have posited a role for inhibition in the decorrelation of activity, this is a theory that could in fact be directly tested using our preparation (Bernacchia and Wang, 2013; Tetzlaff et al., 2012; Wiechert et al., 2010). In addition, the top-down effects in visual circuits have been shown to be mediated via VIP positive interneurons suggesting that perhaps the decorrelation acts via similar mechanisms (Fu et al., 2014). Lastly, abnormalities in acetylcholine receptors have been noted in post-mortem autistic human tissue (Deutsch et al., 2010; Lee et al., 2002). While these deficiencies were noted in a nicotinic receptor while we used a muscarinic agonist, it seems likely from other data in our lab that carbachol likely has nicotonic effects as well at the concentrations that we used in our experiments. Ultimately, neuropsychiatric disease conditions likely have multiple biological underpinnings that converge at the level of microcircuits. Finding circuit level biomarkers (such as the deficit in cholinergically dependent decorrelation) for such diseases will likely not only help shed light on disease mechanisms, but hopefully also aid in drug development as

such biomarkers could be used in the high throughput screening of novel drugs or compounds for treatment.

Chroic ECoG network analysis conclusions

The final study presented here examined how one could reduce network-wide interactions across multiple mesolimbic regions and summarize them as the linear combination of one or a few key-interactions. This finding highlights that mesolimbic networks have the potential to be quite low-dimensional in nature. This in turn is advantageous in the context of attempting to treat dysfunction that results from abnormalities in network-wide information processing. These key interactions that we identified in multiple patients represent plausible foci for either stimulating or reading out the activity within a given brain subnetwork that may have dysfunction in neuropsychiatric disease conditions such as depression and anxiety.

Of interest now, will be to perform stimulation experiments in relevant regions and assay the effect on coherence at these key-interactions. Can stimulation at one region drive coordinated and dispersed changes across the whole entirety of the network? What are the optimal stimulation parameters for affecting many distributed regions within that network? Lastly, it is important to recall that the original purpose of our coherence study was an attempt to identify a neural signature for depression and anxiety conditions for use in a closed-loop environment. The interactions identified here represent excellent candidates for identifying a neural signature but more data from patients must be collected to see whether the interactions identified do indeed contain a neurally based signature of a depressed or

anxious state. Identification of such a signature is a crucial next step in the development of any closed-loop therapy of neuropsychiatric disease.

As technologies enable experimentalists to record from more and more neurons we will be able to further probe the computations that take place in the circuits of the brain. However this deluge of data and exploding of experimental capacity will require clever analytical methods to be able to glean simple and actionable insight from these experiments. The work outlined in this thesis is one such attempt to try to think about these problems and develop some analytical tools to deal with these kinds of datasets and hopefully inform future experiments and studies on the mechanisms of information processing at the level of neuronal networks and how such mechanisms might become aberrant in a disease state.

Bibliography

- Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of neurophysiology*, 70(4):1629–1638.
- Arnsten, A. F., Wang, M. J., and Paspalas, C. D. (2012). Neuromodulation of thought: flexibilities and vulnerabilities in prefrontal cortical network synapses. *Neuron*, 76(1):223–239.
- Ashok, L. and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.*, 15(11):1498–505.
- Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nature reviews. Neuroscience*, 7(5):358–366.
- Aviel, Y., Mehring, C., Abeles, M., and Horn, D. (2003). On embedding synfire chains in a balanced network. *Neural computation*, 15(6):1321–1340.
- Baeg, E., Kim, Y., Huh, K., I, M., Kim, H., and Jung, M. (2003). Dynamics of population code for working memory in the prefrontal cortex. *Neuron*, 40(1):177–188.

- Bassett, D. S. and Bullmore, E. (2006). Small-world brain networks. *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry*, 12(6):512–523.
- Bernacchia, A. and Wang, X. J. (2013). Decorrelation by recurrent inhibition in heterogeneous neural circuits. *Neural computation*, 25(7):1732–1767.
- Borghuis, B. G., Tian, L., Xu, Y., Nikonov, S. S., Vardi, N., Zemelman, B. V., and Looger, L. L. (2011). Imaging light responses of targeted neuron populations in the rodent retina. *J. Neurosci.*, 31(8):2855–67.
- Boyle, L. and Kaufmann, W. E. (2010). The behavioral phenotype of FMR1 mutations. *American journal of medical genetics. Part C, Seminars in medical genetics*, 154C(4):469–476.
- Brody, C. (1999). Correlations without synchrony. *Neural computation*, 11(7):1537–1551.
- Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature neuroscience*, 7(5):446–451.
- Calhoun, V. D., Miller, R., Pearlson, G., and Adal, T. (2014). The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron*, 84(2):262–274.
- Casten, K. S., Gray, A. C., and Burwell, R. D. (2011). Discrimination learning and attentional set formation in a mouse model of fragile x. *Behavioral neuroscience*, 125(3):473–479.

- Chang, J., Gilman, S. R., Chiang, A. H., Sanders, S. J., and Vitkup, D. (2015). Genotype to phenotype relationships in autism spectrum disorders. *Nature neuroscience*, 18(2):191–198.
- Chen, T. W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., Schreiter, E. R., Kerr, R. A., Orger, M. B., Jayaraman, V., Looger, L. L., Svoboda, K., and Kim, D. S. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458):295–300.
- Clapcote, S. J., Lipina, T. V., Millar, J., Mackie, S., Christie, S., Ogawa, F., Lerch, J. P., Trimble, K., Uchiyama, M., Sakuraba, Y., Kaneda, H., Shiroishi, T., Houslay, M. D., Henkelman, R., Sled, J. G., Gondo, Y., Porteous, D. J., and Roder, J. C. (2007). Behavioral phenotypes of *discl* missense mutations in mice. *Neuron*, 54(3):387–402.
- Cohen, M. R. and Maunsell, J. H. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature neuroscience*, 12(12):1594–1600.
- Cornish, K., Sudhalter, V., and Turk, J. (2004). Attention and language in fragile x. *Mental retardation and developmental disabilities research reviews*, 10(1):11–16.
- Cunningham, J. P. and Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature neuroscience*, 17(11):1500–1509.
- Curtis, C. E. and Lee, D. (2010). Beyond working memory: the role of persistent activity in decision making. *Trends Cogn. Sci. (Regul. Ed.)*, 14(5):216–22.

- Dalley, J. W., Cardinal, R. N., and Robbins, T. W. (2004). Prefrontal executive and cognitive functions in rodents: neural and neurochemical substrates. *Neuroscience and biobehavioral reviews*, 28(7):771–784.
- Dembrow, N. C., Chitwood, R. A., and Johnston, D. (2010). Projection-specific neuromodulation of medial prefrontal cortex neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(50):16922–16937.
- Deutsch, S. I., Urbano, M. R., Neumann, S. A., Burket, J. A., and Katz, E. (2010). Cholinergic abnormalities in autism: is there a rationale for selective nicotinic agonist interventions? *Clinical neuropharmacology*, 33(3):114–120.
- Dickson, P. E., Corkill, B., Eric, M., Miller, M. M., Calton, M. A., Goldowitz, D., Blaha, C. D., and Mittleman, G. (2013). Effects of stimulus salience on touchscreen serial reversal learning in a mouse model of fragile x syndrome. *Behavioural brain research*, 252:126–135.
- Diesmann, M., Gewaltig, M., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533.
- Downes, J. H., Hammond, M. W., Xydias, D., Spencer, M. C., Becerra, V. M., Warwick, K., Whalley, B. J., and Nasuto, S. J. (2012). Emergence of a small-world functional network in cultured neurons. *PLoS computational biology*, 8(5):e1002522.
- Durstewitz, D. and Seamans, J. K. (2008). The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-o-methyltransferase genotypes and schizophrenia. *Biol. Psychiatry*, 64(9):739–49.

- Durstewitz, D., Vittoz, N. M., Floresco, S. B., and Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, 66(3):438–48.
- Fellous, J. and Sejnowski, T. (2000). Cholinergic induction of oscillations in the hippocampal slice in the slow (0.5-2 hz), theta (5-12 hz), and gamma (35-70 hz) bands. *Hippocampus*, 10(2):187–197.
- Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., and Raichle, M. E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences of the United States of America*, 103(26):10046–10051.
- Freeman, J. (2015). Open source tools for large-scale neuroscience. *Current opinion in neurobiology*, 32:156–163.
- Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual review of neuroscience*, 32:209–224.
- Fu, Y., Tucciarone, J. M., Espinosa, J., Sheng, N., Darcy, D. P., Nicoll, R. A., Huang, Z., and Stryker, M. P. (2014). A cortical circuit for gain control by behavioral state. *Cell*, 156(6):1139–1152.
- Fujisawa, S., Amarasingham, A., Harrison, M. T., and Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature neuroscience*, 11(7):823–833.

- Fuster, J. and Alexander, G. (1971). Neuron activity related to short-term memory. *Science*, 173(3997):652–4.
- Gadow, K. D., Devincent, C. J., Pomeroy, J., and Azizian, A. (2005). Comparison of DSM-IV symptoms in elementary school-age children with PDD versus clinic and community samples. *Autism : the international journal of research and practice*, 9(4):392–415.
- Gadow, K. D., J, D. C., Pomeroy, J., and Azizian, A. (2004). Psychiatric symptoms in preschool children with PDD and clinic and comparison samples. *Journal of autism and developmental disorders*, 34(4):379–393.
- Ganguli, S., Huh, D., and Sompolinsky, H. (2008). Memory traces in dynamical systems. *Proceedings of the National Academy of Sciences of the United States of America*, 105(48):18970–18975.
- Gerhard, F., Pipa, G., Lima, B., Neuenschwander, S., and Gerstner, W. (2011). Extraction of network topology from Multi-Electrode recordings: Is there a Small-World effect? *Frontiers in computational neuroscience*, 5:4.
- Gibson, J. R., Bartley, A. F., Hays, S. A., and Huber, K. M. (2008). Imbalance of neocortical excitation and inhibition and altered UP states reflect network hyperexcitability in the mouse model of fragile x syndrome. *Journal of neurophysiology*, 100(5):2615–2626.
- Gil, Z., Connors, B., and Amitai, Y. (1997). Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron*, 19(3):679–686.

- Goard, M. and Dan, Y. (2009). Basal forebrain activation enhances cortical coding of natural scenes. *Nature neuroscience*, 12(11):1444–1449.
- Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron*, 61(4):621–634.
- Goncalves, J., Anstey, J. E., Golshani, P., and Carlos, P. (2013). Circuit level defects in the developing neocortex of fragile x mice. *Nature neuroscience*, 16(7):903–909.
- Greicius, M. (2008). Resting-state functional connectivity in neuropsychiatric disorders. *Current opinion in neurology*, 21(4):424–430.
- Grienberger, C. and Konnerth, A. (2012). Imaging calcium in neurons. *Neuron*, 73(5):862–885.
- Harrison, S. J., Woolrich, M. W., Robinson, E. C., Glasser, M. F., Beckmann, C. F., Jenkinson, M., and Smith, S. M. (2015). Large-scale probabilistic functional modes from resting state fMRI. *NeuroImage*, 109:217–231.
- Harvey, C. D., Coen, P., and Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, 484(7392):62–68.
- Hasselmo, M. (1995). Neuromodulation and cortical function: modeling the physiological basis of behavior. *Behavioural brain research*, 67(1):1–27.

- Hennequin, G., Vogels, T. P., and Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*, 82(6):1394–1406.
- Hikida, T., Hanna, J., Seshadri, S., Oishi, K., Hookway, C., Kong, S., Wu, D., Xue, R., Andradé, M., Tankou, S., Mori, S., Gallagher, M., Ishizuka, K., Pletnikov, M., Kida, S., and Sawa, A. (2007). Dominant-negative DISC1 transgenic mice display schizophrenia-associated phenotypes detected by measures translatable to humans. *Proceedings of the National Academy of Sciences of the United States of America*, 104(36):14501–14506.
- Humphries, M. D. and Gurney, K. (2008). Network ‘small-world-ness’: a quantitative method for determining canonical network equivalence. *PloS one*, 3(4):e0002051.
- Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., and Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 304(5670):559–64.
- Kawaguchi, Y. (1997). Selective cholinergic modulation of cortical GABAergic cell subtypes. *Journal of neurophysiology*, 78(3):1743–1747.
- Kodama, M., Fujioka, T., and Duman, R. S. (2004). Chronic olanzapine or fluoxetine administration increases cell proliferation in hippocampus and prefrontal cortex of adult rat. *Biological psychiatry*, 56(8):570–580.
- Laje, R. and Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature neuroscience*, 16(7):925–933.

- Lee, M., C, M., Graham, A., Court, J., Jaros, E., Perry, R., Iversen, P., Bauman, M., and Perry, E. (2002). Nicotinic receptor abnormalities in the cerebellar cortex in autism. *Brain : a journal of neurology*, 125(Pt 7):1483–1495.
- Luczak, A., Barthó, P., and Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–25.
- Luis, C., Miller, J. K. E., Hamm, J. P., Jackson, J., and Yuste, R. (2015). Endogenous sequential cortical activity evoked by visual stimuli. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 35(23):8813–8828.
- Luongo, F. J., Horn, M. E., and Sohal, V. S. (2015). Putative Microcircuit-Level substrates for attention are disrupted in mouse models of autism. *Biological psychiatry*.
- M, M. S., Neal, S. J., Lin, Q., Hughes, Z., and Smith, D. G. (2013). The BTBR mouse model of autism spectrum disorders has learning and attentional impairments and alterations in acetylcholine and kynurenic acid in prefrontal cortex. *PloS one*, 8(4):e62189.
- Marder, E. (2015). Understanding brains: details, intuition, and big data. *PLoS biology*, 13(5):e1002147.
- Markram, H., Maria, T., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature reviews. Neuroscience*, 5(10):793–807.
- Markram, H., Rinaldi, T., and Markram, K. (2007). The intense world syndrome—an alternative hypothesis for autism. *Frontiers in neuroscience*, 1(1):77–96.

- McConnell, S. (1991). The generation of neuronal diversity in the central nervous system. *Annual review of neuroscience*, 14:269–300.
- Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. *Neuron*, 63(6):879–888.
- Moon, J., Beaudin, A., Verosky, S., Driscoll, L., Weiskopf, M., Levitsky, D., Crnic, L., and Strupp, B. (2006). Attentional dysfunction, impulsivity, and resistance to change in a mouse model of fragile x syndrome. *Behavioral neuroscience*, 120(6):1367–1379.
- Moore, S., Turnpenny, P., Quinn, A., Glover, S., Lloyd, D., Montgomery, T., and Dean, J. (2000). A clinical study of 57 children with fetal anticonvulsant syndromes. *Journal of medical genetics*, 37(7):489–497.
- Mukamel, E. A., Nimmerjahn, A., and Schnitzer, M. J. (2009a). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, 63(6):747–60.
- Mukamel, E. A., Nimmerjahn, A., and Schnitzer, M. J. (2009b). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, 63(6):747–760.
- O’Connor, D., Peron, S. P., Huber, D., and Svoboda, K. (2010). Neural activity in barrel cortex underlying vibrissa-based object localization in mice. *Neuron*, 67(6):1048–1061.
- Ohiorhenuan, I. E., Mechler, F., Purpura, K. P., Schmid, A. M., Hu, Q., and Victor, J. D. (2010). Sparse coding and high-order correlations in fine-scale cortical networks. *Nature*, 466(7306):617–621.

- Pafundo, D. E., Miyamae, T., Lewis, D. A., and Guillermo, G. (2013). Cholinergic modulation of neuronal excitability and recurrent excitation-inhibition in prefrontal cortex circuits: implications for gamma oscillations. *The Journal of physiology*, 591(Pt 19):4725–4748.
- Parikh, V. and Sarter, M. (2008). Cholinergic mediation of attention: contributions of phasic and tonic increases in prefrontal cholinergic activity. *Annals of the New York Academy of Sciences*, 1129:225–235.
- Pitkanen, A., Pikkarainen, M., Nurminen, N., and Ylinen, A. (2000). Reciprocal connections between the amygdala and the hippocampal formation, perirhinal cortex, and postrhinal cortex in rat. a review. *Annals of the New York Academy of Sciences*, 911:369–391.
- Pletnikov, M., Ayhan, Y., Nikolskaia, O., Xu, Y., Ovanesov, M., Huang, H., Mori, S., Moran, T., and Ross, C. (2008). Inducible expression of mutant human DISC1 in mice is associated with brain and behavioral abnormalities reminiscent of schizophrenia. *Molecular psychiatry*, 13(2):173–86, 115.
- Rasalam, A., Hailey, H., Williams, J., Moore, S., Turnpenny, P., Lloyd, D., and Dean, J. (2005). Characteristics of fetal anticonvulsant syndrome associated autistic disorder. *Developmental medicine and child neurology*, 47(8):551–555.
- Roxin, A., Hakim, V., and Brunel, N. (2008). The statistics of repeating patterns of cortical activity can be reproduced by a model network of stochastic binary neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(42):10734–10745.

- Runfeldt, M. J., Sadovsky, A. J., and N, M. J. (2014). Acetylcholine functionally reorganizes neocortical microcircuits. *Journal of neurophysiology*, 112(5):1205–1216.
- Sadovsky, A. J. and MacLean, J. N. (2013). Scaling of topologically similar functional modules defines mouse primary auditory and somatosensory microcircuitry. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(35):14048–60, 14060a.
- Sakata, S. and Harris, K. D. (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron*, 64(3):404–18.
- Sanchez-Vives, M. V. and McCormick, D. A. (2000). Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat. Neurosci.*, 3(10):1027–34.
- Schneider, T. and Przewocki, R. (2005). Behavioral alterations in rats prenatally exposed to valproic acid: animal model of autism. *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology*, 30(1):80–89.
- Schneider, T., Roman, A., Agnieszka, B., Kubera, M., Budziszewska, B., Schneider, K., and Przewocki, R. (2008). Gender-specific behavioral and immunological alterations in an animal model of autism induced by prenatal exposure to valproic acid. *Psychoneuroendocrinology*, 33(6):728–740.
- Schneidman, E., Berry, M. J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012.

- Scholkopf, B. and Mullert, K. (1999). Fisher discriminant analysis with kernels. *Neural networks for signal processing IX*, 1:1.
- Seidemann, E., Meilijson, I., Abeles, M., Bergman, H., and Vaadia, E. (1996). Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 16(2):752–768.
- Shu, Y., Hasenstaub, A., and A, M. D. (2003). Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–93.
- Smith, S. M., Vidaurre, D., Beckmann, C. F., Glasser, M. F., Jenkinson, M., Miller, K. L., Nichols, T. E., Robinson, E. C., Gholamreza, S., Woolrich, M. W., Barch, D. M., Uurbil, K., and Van Essen, D. C. (2013). Functional connectomics from resting-state fMRI. *Trends in cognitive sciences*, 17(12):666–682.
- Sporns, O. and Zwi, J. D. (2004). The small world of the cerebral cortex. *Neuroinformatics*, 2(2):145–162.
- Stafford, J. M., Jarrett, B. R., Oscar, M., Mills, B. D., Cain, N., Mihalas, S., Lahvis, G. P., Lattal, K., Mitchell, S. H., David, S. V., Fryer, J. D., Nigg, J. T., and Fair, D. A. (2014). Large-scale topology and the default mode network in the mouse connectome. *Proceedings of the National Academy of Sciences of the United States of America*, 111(52):18745–18750.
- Stevenson, I. H., London, B. M., Oby, E. R., Sachs, N. A., Reimer, J., Englitz, B., David, S. V., Shamma, S. A., Blanche, T. J., Mizuseki, K., Zandvakili, A., Hatsopoulos, N. G.,

- Miller, L. E., and Kording, K. P. (2012). Functional connectivity and tuning curves in populations of simultaneously recorded neurons. *PLoS computational biology*, 8(11):e1002775.
- Stewart, C. V. and Plenz, D. (2006). Inverted-U profile of dopamine-NMDA-mediated spontaneous avalanche recurrence in superficial layers of rat prefrontal cortex. *J. Neurosci.*, 26(31):8148–59.
- Stirman, J. N., Smith, I. T., Kudenov, M. W., and Smith, S. L. (2014). Wide field-of-view, twin-region two-photon imaging across extended cortical networks. *bioRxiv*.
- Sussillo, D. and Abbott, L. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557.
- Takeuchi, D., Hirabayashi, T., Tamura, K., and Miyashita, Y. (2011). Reversal of interlaminar signal between sensory and memory processing in monkey temporal cortex. *Science*, 331(6023):1443–7.
- Tetzlaff, T., Helias, M., Einevoll, G. T., and Diesmann, M. (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLoS computational biology*, 8(8):e1002596.
- Tian, L., Hires, S., Mao, T., Huber, D., Chiappe, M., Chalasani, S. H., Petreanu, L., Akerboom, J., A. M. S., Schreiter, E. R., Bargmann, C. I., Jayaraman, V., Svoboda, K., and Looger, L. L. (2009a). Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat. Methods*, 6(12):875–81.

- Tian, L., Hires, S., Mao, T., Huber, D., Chiappe, M., Chalasani, S. H., Petreanu, L., Akerboom, J., A. M. S., Schreiter, E. R., Bargmann, C. I., Jayaraman, V., Svoboda, K., and Looger, L. L. (2009b). Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nature methods*, 6(12):875–881.
- von Stein, A. and Sarnthein, J. (2000). Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *International journal of psychophysiology : official journal of the International Organization of Psychophysiology*, 38(3):301–313.
- Wang, Z., Chen, L. M., Négyessy, L., Friedman, R. M., Mishra, A., Gore, J. C., and Roe, A. W. (2013). The relationship of anatomical and functional connectivity to resting-state connectivity in primate somatosensory cortex. *Neuron*, 78(6):1116–1126.
- Watrous, A. J., Tandon, N., Conner, C. R., Pieters, T., and Ekstrom, A. D. (2013). Frequency-specific network connectivity increases underlie accurate spatiotemporal memory retrieval. *Nature neuroscience*, 16(3):349–356.
- Watts, D. and Strogatz, S. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.
- Welch, J. M., Lu, J., Rodriguiz, R. M., Trotta, N. C., Peca, J., Ding, J. D., Feliciano, C., Chen, M., Adams, J., Luo, J., Dudek, S. M., Weinberg, R. J., Calakos, N., Wetsel, W. C., and Feng, G. (2007). Cortico-striatal synaptic defects and OCD-like behaviours in sapap3-mutant mice. *Nature*, 448(7156):894–900.

- Wiechert, M. T., Judkewitz, B., Riecke, H., and Friedrich, R. W. (2010). Mechanisms of pattern decorrelation by recurrent neuronal circuits. *Nature neuroscience*, 13(8):1003–1010.
- Willsey, A., Sanders, S. J., Li, M., Dong, S., Tebbenkamp, A. T., Muhle, R. A., Reilly, S. K., Lin, L., Fertuzinhos, S., Miller, J. A., Murtha, M. T., Bichsel, C., Niu, W., Cotney, J., AG, E., Gockley, J., Gupta, A. R., Han, W., He, X., Hoffman, E. J., Klei, L., Lei, J., Liu, W., Liu, L., Lu, C., Xu, X., Zhu, Y., Mane, S. M., Lein, E. S., Wei, L., Noonan, J. P., Roeder, K., Devlin, B., Sestan, N., and State, M. W. (2013). Coexpression networks implicate human midfetal deep cortical projection neurons in the pathogenesis of autism. *Cell*, 155(5):997–991007.
- Winterer, G. and Weinberger, D. R. (2004). Genes, dopamine and cortical signal-to-noise ratio in schizophrenia. *Trends Neurosci.*, 27(11):683–90.
- Xiang, Z., Huguenard, J., and Prince, D. (1998). Cholinergic switching within neocortical inhibitory networks. *Science (New York, N.Y.)*, 281(5379):985–988.
- Yamada, Y. and Mikoshiba, K. (2012). Quantitative comparison of novel GCaMP-type genetically encoded Ca^{2+} indicators in mammalian neurons. *Front Cell Neurosci*, 6:41.
- Yu, S., Huang, D., Singer, W., and Nikolic, D. (2008). A small world of neuronal synchrony. *Cerebral cortex (New York, N.Y. : 1991)*, 18(12):2891–2901.
- Zhao, S., Ting, J. T., Atallah, H. E., Qiu, L., Tan, J., Gloss, B., Augustine, G. J., Deisseroth, K., Luo, M., Graybiel, A. M., and Feng, G. (2011). Cell typespecific channelrhodopsin-2

transgenic mice for optogenetic dissection of neural circuitry function. *Nature methods*, 8(9):745–752.

Zhou, J., Greicius, M. D., Gennatas, E. D., Growdon, M. E., Jang, J. Y., Rabinovici, G. D., Kramer, J. H., Weiner, M., Miller, B. L., and Seeley, W. W. (2010). Divergent network connectivity changes in behavioural variant frontotemporal dementia and alzheimer’s disease. *Brain : a journal of neurology*, 133(Pt 5):1352–1367.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

Please sign the following statement:

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.



Author Signature

11/23/15
Date