Harnessing the Variability of Neuronal Activity: from Single Neurons to Networks

Eric S. Kuebler, PhD Candidate

School of Psychology, Faculty of Social Sciences

University of Ottawa, Ottawa, Ontario, Canada

Submitted in partial requirement of PhD in Experimental Psychology.

This page was intentionally left blank.

Acknowledgements

I would primarily like to thank my advisor Dr. Jean-Philippe Thivierge for his patience, support, and guidance throughout the course of my graduate studies.

I would especially like to thank Dr. André Longtin and the Neural Dynamics Group at uOttawa for contributing to my knowledge of experimental and computational neuroscience. Dr. Denis Cousineau for his leadership with the QUIBB group and for invaluable feedback during my graduate studies. Drs. Joseph S. Tauskela and Marzia Martina for hosting me at the National Research Council (Ottawa, Ontario, Canada), as well as teaching me about the chemistry, hands-on biology of *in vitro* neuron cultures, and allowing me to make multi-electrode recordings with their equipment and preparations. Dr. Trudy Bergere for her inspiration and motivation, as well as her support as my mentor throughout both undergraduate and graduate studies.

I would like to thank Drs. Jean-Claude Béïque, Richard Naud, Maurice Chacron, and John Lewis for taking the time to read and review this work. As well as, Drs. Andreas T. Schaefer, Francis K. Skinner, Julio C. Martinez-Trujillo, Nicholas Swindale, and Alfonso Renart for invaluable feedback during this work. I would also like to thank all the support staff and administration at uOttawa for their patience and aid throughout the course of my graduate studies.

I would also like to thank my graduate colleagues for many insightful discussions and feedback: Dr. Idu Azugo, Laurence Morrisette, Taylor Hatchard, Philippe Vincent-Lamarre, Nareg Berberian, Matthew Ross, Matias Calderini, Teresa Allan, and Melissa Howard.

I would like to thank my loving partner Erica Traccitto, my parents, Paul and Patricia Kuebler, my brother, Andrew Kuebler, my nephews, Benjamin and Nathan Kuebler, many of my relatives, and all my friends. Without all their patience and support this would not have been possible.

Table of contents

Statement on published work

This thesis document is based on the following published work:

Kuebler, E. S., Tauskela, J. S., Aylsworth, A., Zhao, X., & Thivierge, J. P. (2015). Burst predicting neurons survive an in vitro glutamate injury model of cerebral ischemia. *Sci Rep, 5*, 17718. doi:10.1038/srep17718

Kuebler, E. S., & Thivierge, J. P. (2014). Spiking variability: Theory, measures and implementation in Matlab. *The quantitative methods for psychology, 10*(2), 131-142.

Kuebler, E. S. K., Bonnema, E., McCorriston, J., & Thivierge, J. P. (2013). *Stimulus discrimination in networks of spiking neurons.* Paper presented at the International Joint Conference on Neural Networks, Dallas, Texas.

As well as the following submitted work:

Kuebler, E.S., Longtin, A., Bent, N., Vincent-Lamarre, P., & Thivierge, J. P. (submitted: BICY-D- 17-00035R1). *Accumulation of Spike Time Variance Imposes Strict Constraints on Coding with Oscillations*. Biological Cybernetics.

Kuebler, E.S., Calderini, M., Lambert, P., & Thivierge, J. P. *Optimal readout of neural activity near criticality*. Book chapter submitted for “The Functional Role of Critical Dynamics in Neural Systems”. Springer.

Statement of contribution by the candidate

In Chapter One, we use a combination of *in vitro* experimental data collected by another lab (Schaefer, Angelo, Spors, & Margrie, 2006), paired with our own numerical simulations and analyses. The *in vitro* data was not recorded by the candidate; however,

our use of it was approved by the first author of the publication. The candidate took part in the modeling effort of the first chapter, dissemination of results at conferences and the preparation of a manuscript for submission to peer-reviewed journal. In Chapter Two and Three, we use *in vitro* data acquired by working with collaborators (National Research Council, NRC, Ottawa), that was paired with our analyses, as well as computational models. The data from Chapter One and Two was recorded prior to the candidate's arrival in the lab; however, the candidate did learn the methodology of recording with multi-electrode arrays (MEA) and performed many recordings for ongoing projects during the time of graduate studies. In Chapter Two and Three, the candidate performed all analyses of MEA recordings, contributed to the modeling effort, disseminated results at conferences and participated in the preparation of the manuscript for publication to a peer-reviewed as well as a book chapter.

Helpful tips for reading the document

i. There are three chapters, written to be read sequentially, not in article format.

ii. To differentiate between text and figure captions, text is justified, and figure captions are left aligned.

iii. Figure and equation counts contain a chapter number followed by hyphen or period, respectively.

   a. All figures are denoted by hyphens, i.e.: “1-1”, “1-2”.

   b. All equations are denoted by periods, i.e.: “1.1”, “1.2”.

iv. There is only one citation count, and all references for citations are at the end of the document.

v. Print out the five tables for easy reference (i.e., contents, acronyms, symbols, equations, and figures).

vi. Citations and references are in APA format.

Table of acronyms

(Sorted alphabetically)

µm: micrometers

2D: 2-dimensional

AMPA: α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid

ANOVA: analysis of variance

AP: action potential

APV/DNQX: (2*R*)-amino-5-phosphonovaleric acid / dinitroquinoxaline

ATP: adenosine triphosphate

BP: burst predictor

CA1: cornu ammonis 1

Ca^{2+}: calcium

Cl^{-}: cloride

DG: dentate gyrus

DIV: days *in vitro*

DMSO: dimethyl siloxane

DNA: nucleic acid

EOI: electrode of interest

EPSP: excitatory post-synaptic potential

GABA: gamma-Aminobutyric acid

Hz: hertz

IPSP: inhibitory post-synaptic potential

K^{+}: potassium

kHz: kilohertz

LDA: linear discriminant analysis

LIF: leaky integrate and fire

µA: microamps

MEA: multi-electrode array

ms: milliseconds

MUA: multi-unit activity

mV: millivolt

Na^+: sodium

NDMA: N-methyl-D-aspartate

OB: olfactory bulb

pA: Pico amp

PC: population coupling

PCA: principal components analysis

PI: propidium iodide

pS: Pico siemens

PSTH: peri-stimulus time histogram

PTX: picrotoxin

RC: resistor capacitance

ROC: receiver operating characteristic

S.D.: standard deviation

SEM: standard error of mean

SIM: simulation

SOC: self-organized criticality

TBOA: DL-threo-β-benzyloxyaspartic acid

vs.: versus

Summary

Neurons and networks of the brain may use various strategies of computation to provide the neural substrate for sensation, perception, or cognition. To simplify the scenario, two of the most commonly cited neural codes are firing rate and temporal coding, whereby firing rates are typically measured over a longer duration of time (i.e., seconds or minutes), and temporal codes use shorter time windows (i.e., 1 to 100 ms). However, it is possible that neurons may use other strategies. Here, we highlight three methods of computation that neurons, or networks, of the brain may use to encode and/or decode incoming activity. First, we explain how single neurons of the brain can utilize a neuronal oscillation, specifically by employing a 'spike-phase' code wherein responses to stimuli have greater reliability, in turn increasing the ability to discriminate between stimuli. Our focus was to explore the limitations of spike-phase coding, including the assumptions of low firing rates and precise timing of action potentials. Second, we examined the ability of single neurons to track the onset of network bursting activity, namely 'burst predictors'. In addition, we show that burst predictors were less susceptible to an *in vitro* model of neuronal stroke (i.e., excitotoxicity). Third, we discuss the possibility of distributed processing with neuronal networks of the brain. Specifically, we show experimental and computational evidence supporting the possibility that the population activity of cortical networks may be useful to downstream classification. Furthermore, we show that when network activity is highly variable across time, there is an increase in the ability to linearly separate the spiking activity of various networks. Overall, we use the results of both experimental and computational methods to highlight three strategies of computation that neurons and networks of the brain may employ.

General Introduction

Since the late 1920s we have been able to record from nerve fibers (Adrian & Bronk, 1928), and these recordings have improved our understanding of how various brain processes work. In addition, we have begun to understand how activity in the brain gives rise to sensations, perceptions, and cognitions. One observation as a function of neuronal recordings is that neurons tend to 'spike' (i.e., reach action potential) in a variable fashion across time (Softky & Koch, 1993; Tolhurst, Movshon, & Dean, 1983; Werner & Mountcastle, 1963), even when presented with the same stimulus (Mainen & Sejnowski, 1995). Theoretical research shows that when neurons are interconnected in a balanced fashion, the firing activity of a single neuron can be characterized as chaotic (van Vreeswijk & Sompolinsky, 1996, 1998). From a computational perspective, the variation in neuronal activity can be useful, for example when the variability of two neurons varies in a similar fashion, termed 'noise correlations' (Abbott & Dayan, 1999; Cohen & Kohn, 2011; Moreno-Bote et al., 2014; Pillow et al., 2008; Shadlen & Newsome, 1998). However, some portion of the neural activity is believed to be 'noise' that comes from a wide-range of different sources, especially in the context of making experimental recordings (London, Roth, Beeren, Hausser, & Latham, 2010). One of the foremost burning questions in neuroscience is how single neurons and networks overcome variability and provide the substrate for a consistent and reliable stream of sensations based on stimuli in the environment. Here, we employ experimental and theoretical methods to investigate three plausible strategies for successful computation with both single neurons and neuronal networks.

In Chapter One, we discuss neuronal oscillations, in the context of both single neurons and networks. Here, we use a combination of *in vitro* experimental data collected by another lab (Schaefer et al., 2006), paired with our own numerical simulations and analyses. In Chapter Two, we examined the ability of single neurons to predict spontaneous bursts of network activity. Finally, in Chapter Three, we studied whether population bursting activity, that is characterized as 'at or near a critical point', is useful for linear decoding. In Chapter Two and Three, we use *in vitro* data acquired by working with collaborators (National Research Council, NRC, Ottawa), that was paired with our analyses and computational models. Overall, we use two different methods of *in vitro* recordings, paired with computational modeling to investigate the strategies that neurons and networks of the brain use.

Neuronal oscillations

Oscillations are a common feature of brain recordings across many regions, preparations, and species (Buzsaki & Draguhn, 2004). Oscillations are a hallmark of membrane potential recordings with single neurons, and are characterized by subthreshold fluctuations that drive the periodic firing of individual neurons (Buzsaki & Wang, 2012; Engel, Fries, & Singer, 2001; X. J. Wang, 2010). Despite the wide-spread attention of researchers, the limitations of encoding stimuli by employing oscillations have not been fully elucidated.

Oscillations can be generated by several mechanisms. Layer 5 pyramidal neurons may generate subthreshold oscillations intrinsically via sodium- and calcium-dependent conductance (Silva, Amitai, & Connors, 1991). Pacemaker neurons of the thalamus may generate oscillations intrinsically by calcium spikes that trigger a pattern of action

potentials caused by sodium and potassium (Bal & McCormick, 1993). Similar intrinsic mechanisms generate oscillations in clusters of interneurons of the medulla called the Pre-Bötzinger complex responsible for respiratory rhythms (Butera, Rinzel, & Smith, 1999a, 1999b). There are also conditions where oscillations are generated by an interplay between excitation and inhibition (Bartos et al., 2002; Nicolas Brunel & Wang, 2003; Doiron, Chacron, Maier, Longtin, & Bastian, 2003; Ferguson, Chatzikalymniou, & Skinner, 2017; Gan, Weng, Pernía-Andrade, Csicsvari, & Jonas, 2017).

The results of both experimental and computational work suggest that oscillations carry functional benefits. For instance, the phase of oscillating neuronal activity is correlated with enhanced stimulus detection and discrimination (Monto, Palva, Voipio, & Palva, 2008; Y. Wang, Iliescu, Ma, Josic, & Dragoi, 2011). Computational work captures these findings (Masuda & Doiron, 2007) and highlights the functional benefits of synchronization in enhancing the reliability of responses to stimuli (H. P. Wang, Spencer, Fellous, & Sejnowski, 2010; X. J. Wang, 2010) as well as signal propagation (Akam & Kullmann, 2010, 2012, 2014; Fries, Nikolic, & Singer, 2007; Singer, 1999).

Oscillations give rise to a few possible methods of encoding information about stimuli. One possibility is that the precise spike times of neurons can be tagged with a 'phase' of oscillations, herein referred to as 'spike-phase coding' (K. D. Harris et al., 2002; Kayser, Montemurro, Logothetis, & Panzeri, 2009; Lisman, 2005; Montemurro, Rasch, Murayama, Logothetis, & Panzeri, 2008). One hypothesis is that oscillations contribute to spike timing reliability, whereby the rising phase of an oscillation provides the substrate for precise spike times, in turn increasing the ability to discriminate between stimuli

(Schaefer et al., 2006). However, to our knowledge, a full exploration of the limitations of spike-phase coding has not been conducted, and this is the topic of Chapter One.

Neuronal population activity

Brain function may be best understood by examining the structure of aggregate activity of many neurons in a network, namely 'population coding'. How single neurons contribute to population activity, and whether downstream neurons can decode this activity, remain burning questions in neuroscience.

Here, we address two main questions related to neural population activity. In Chapter Two we investigate how single neurons contribute to population activity in the form of network bursts. In Chapter Three, we examine whether population bursting activity is useful for linear decoding by downstream units.

Network bursts

Spontaneous neuronal activity *in vivo* and *in vitro* is often characterized by network bursts, whereby a substantial proportion of neurons are active in close temporal contiguity (Bellay, Klaus, Seshadri, & Plenz, 2015; Petermann et al., 2009; Yu et al., 2017). Network bursts are a common feature of several *in vitro* preparations of the central nervous system (Beggs & Plenz, 2003; Bonifazi et al., 2009; Chiappalone, Bove, Vato, Tedesco, & Martinoia, 2006; Langlois, Cousineau, & Thivierge, 2014; Shaukat & Thivierge, 2016; Vincent, Tauskela, & Thivierge, 2012; Wagenaar, Pine, & Potter, 2006). Bursts are typically characterized by short transients of rapid spiking activity flanked by silent periods. Cultures of cortical neurons begin to emit network bursts around day *in vitro* (DIV) six and gradually develop a rich repertoire of activity over the course of maturation (Wagenaar et al., 2006).

Results of experimental work suggest that network bursts are associated with the propagation of neuronal information (Kepecs & Lisman, 2003; Kepecs, Wang, & Lisman, 2002) and the formation of central pattern generators (Butera et al., 1999a). Furthermore, the statistics of spontaneous network bursts follow a power law distribution, with scaling exponents that are suggestive of a critical state that promotes information processing (Beggs, 2008) (with some controversy surrounding the issue (Touboul & Destexhe, 2010)). Overall, there may be functional benefits to bursting activity generated by populations of neurons.

There are many ways one can imagine that networks of neurons can encode information. We begin by examining the possibility that neurons in a network each behave differently prior to a burst. Experimental work has identified subsets of neurons that reliably predict the occurrence of network bursts, namely 'burst predicting neurons' (Eckmann, Jacobi, Marom, Moses, & Zbinden, 2008; Eytan & Marom, 2006)[1]. Both experimental and theoretical work suggests that these burst predicting cells have stable long-term spiking dynamics that can be modified by electrical stimulation (Zbinden, 2010, 2011). Furthermore, these 'early-to-fire' neurons have dense axonal arborizations and may be GABAergic (Bonifazi et al., 2009). Neurons that predict bursts of activity may therefore be critical to network function, especially following a neuronal insult.

Because burst-predicting neurons are suggested to play a causal role in the initiation of network activity (Bonifazi et al., 2009), their survival following a neural injury may be crucial to the maintenance of network bursts. In the hippocampus, young GABAergic interneurons are known to have long-range (or short-range, but dense) axonal arborisation, and this may play a role in their ability to fire prior to other neurons in the

local network (Bonifazi et al., 2009). Experimental work reports that hippocampal GABAergic interneurons can be more resistant than principal (pyramidal) cells to an ischemic insult (Schwartz-Bloom & Sah, 2001). Taken together, if GABAergic interneurons are burst predictors and less susceptible to neurotoxic insult, these neurons may be important to recovery of network function after insult. This link, however, to our knowledge has never been established directly. In Chapter Two, we aim to reveal some of the characteristics of neurons that fire just prior to the onset of network bursting and examine their resilience to a simulated ischemic insult.

The 'critical' state

In contrast to global oscillations, there is experimental and theoretical evidence that the brain is a critical system (Beggs, 2008; Beggs & Timme, 2012; Plenz, 2013; Shew & Plenz, 2013; Tagliazucchi, Balenzuela, Fraiman, & Chialvo, 2012). A critical system is commonly defined as one that operates near a critical point between order and disorder. During an ordered state, the system is completely predictable (locked, static, etc.). One example of a critical system is iron at a specific temperature, where the direction of momentum of particles can generate magnetism. When a piece of iron is at a low temperature, it is magnetic because all the spins occur in the same direction due to pairwise interactions among the units of the lattice (Beggs & Timme, 2012). On the other hand, during the disordered state, the system is characterized by randomness. Following the same example, when the iron is heated to a higher temperature, pairwise interactions are dominated by the impact of temperature, and all spins occur in different directions, thus the iron loses magnetism. However, when iron is heated to certain temperature (1333 ˚F), the critical point, pairwise interactions are balanced with the effect of

temperature, thus some local groups of spins occur in the same direction, others do not, and these groups constantly change over time (Beggs & Timme, 2012). Near the critical point, the iron would have local domains of magnetism based on where the pairwise interactions occur. Moreover, the sizes of groups where spins occur in the same direction vary widely, there are many small groups and just a few large groups. Most interestingly, staying near the critical point allows a system to remain free from being trapped in either of the two extreme states (i.e., order and disorder).

In neuronal networks, order is observed during the synchronous state (i.e., supercritical), and disorder during the asynchronous state (i.e., subcritical). Both synchrony and asynchrony are commonly observed in neuronal recordings, and the brain shows the ability to switch between up and down states by either rapidly amplifying or dampening activity (Beggs & Timme, 2012). Experimental work shows that the brain at rest, often referred to as the 'resting state,' indeed has many characteristics of a critical system including: power-law scaling of neuronal avalanches, long-range spatial and temporal correlations, as well as scaling of fluctuations in activity with network size (Plenz, 2013). Thus, the critical point keeps a system like the brain both reliable (i.e., presenting the same stimulus and recording a similar response) and flexible (i.e., adaptive to environmental change).

Critical systems may be beneficial from a computational standpoint. First, during the critical state, information about stimuli transmitted throughout neuronal networks in the brain is maximized (Shew, Yang, Yu, Roy, & Plenz, 2011). Second, during the critical state, neuronal activity has the largest range of responses to stimuli, in turn increasing the potential number of stimuli that can be encoded (Gautam, Hoang, McClanahan,

Grady, & Shew, 2015). For a full review of the benefits of criticality see (Shew & Plenz, 2013).

Investigations on the role of neural criticality have primarily focused on neural encoding, with little attention paid to how criticality may impact neural decoding by downstream linear units. Further, it is not known if the critical state may be optimal for decoding by downstream structures has received little attention. Specifically, under what conditions is the population activity of neuronal networks useful for decoding by downstream units? In Chapter Three we examine the ability of linear units to decode the activity of neuronal networks.

Overall, this work will highlight three biologically realistic strategies that neurons or networks of the brain may use to perform computations in a noisy environment.

Research design justification

To answer the question of how to harness the variability in neuronal activity, we have employed both computational and experimental methodologies. This is generally accepted as a powerful approach as this synergy drives both observation and explanation (Skinner, 2017).

For each of the chapters we have used a computational approach to unpack the results we have found experimentally and derive some mechanistic explanation. The beginnings of computational neuroscience can be traced back to scientists such as Louis Lapicque (1866-1952), Alan L. Hodgkin* (1914-1998), Andrew F. Huxley* (1917-2012), David H. Hubel* (1926-2013), Torsten Wiesel*, and David Marr (1945-1980) – four of which were awarded Nobel prizes for Physiology or Medicine (denoted by *). Relevant to Chapter

One, Louis Lapicque originated the leaky integrate and fire (LIF) model, when he described how the voltage of a neuron's membrane potential could be represented across time as a derivative (Abbott, 1999; N. Brunel & van Rossum, 2007; Lapicque, 1907). More recently, the LIF model has been used for both single neurons and large networks (Litwin-Kumar & Doiron, 2012; Renart et al., 2010), and has been adapted for biologically realistic variables (Naud & Gerstner, 2012; Naud, Marcille, Clopath, & Gerstner, 2008). We chose to use the LIF model because it is important to reduce a model to as few parameters as possible, simulations are computationally efficiency. Most importantly, LIF models are known to be a good fit for the firing activity of single neurons (Gerstner & Naud, 2009; Jolivet et al., 2008; Naud, Bathellier, & Gerstner, 2014). Relevant to Chapter Two and Three, we used a branching model that was originated by Francis Galton (1822-1911) and Henry W. Watson (1827-1903) to examine fertility data for large populations, termed the 'Galton-Watson branching process' (T. E. Harris, 1989). More recently, branching models have been used by Clayton Haldeman and John M. Beggs to model the activity of neural networks, specifically the stability of numerous network states (Haldeman & Beggs, 2005). We chose to use the branching model for our *in vitro* data because of a strong correspondence between experiments and models noted previously in the literature, and because it provides a baseline where the propagation of activity in the network is random. The models we have chosen are strongly supported and have been applied to many neuroscience problems.

A benefit of computational modeling is the degree of control that is not possible in experiments, allowing researchers to explore the mechanisms behind experimental results. For example, for the results of Chapter Three it was important to control the

parameters of firing activity to test the outcome on decoding (our dependent measure). A second major benefit of modeling is that this methodology allows researchers to generate predictions for future experimental research. Overall, computational modeling offers benefits that accentuate experimental work, and strengthen our understanding of how the brain works.

In Chapter Two and Three, we have used an experimental approach based on *in vitro* multi-electrode array (MEA) recordings. Previously, researchers recorded single neurons, for instance, using the patch clamp method, developed by Erwin Neher and Bert Sakmann, whereby the electrical activity of the cell is recorded via pipette inserted into the membrane (Sakmann & Neher, 1984). More recently, technologists have developed recording platforms that acquire data from several recording electrodes embedded in a petri dish, namely MEAs (Pine, 2006). MEAs were originally designed to record from muscle fibers that were dissociated and cultured (Thomas, Springer, Loeb, Berwald-Netter, & Okun, 1972). With MEAs, electrical signals are recorded extracellularly, from within the extracellular fluid, and thus are recording multiple neurons surrounding the electrode, termed ‘multi-unit activity’ (MUA). The advantage of MEAs is that researchers can record from a large number of interconnected neurons, thus providing a readout of population level activity (Beggs & Plenz, 2003; Kuebler, Tauskela, Aylsworth, Zhao, & Thivierge, 2015; Langlois et al., 2014; Plenz, 2012, 2013; Shaukat & Thivierge, 2016; Shew et al., 2011; Tauskela et al., 2008; Thivierge, 2014; Vincent, Tauskela, Mealing, & Thivierge, 2013; Vincent et al., 2012). The distinction between single neuron activity and populations is important because for some time researchers could not record from large numbers of neurons at the same time, and it is possible that some computation occurs at

the network level. Like related research, our aim was to employ MEAs as a means of measuring population activity for the investigations of both Chapter Two and Three.

The real power of experimental and computational methodologies is realized when scientists utilize them together. Experiments can drive new discoveries, and computational work can generate predictions for future research. The synergy between experimental and computational methodologies is important to the development of neuroscience, and our understanding of how the brain works (Skinner, 2017).

Chapter 1 – A Spike-Phase Neural Code and the Limitations

Abstract

Oscillations are widely observed in recordings of many different regions and preparations of the brain. One hypothesis is that a “spike-phase” neural code exists whereby stimuli are encoded in the timing of spikes relative to oscillations in the membrane potential of single neurons. Spike-phase coding has been reported in experiments both *in vivo* and *in vitro* across several regions of the brain; however, an examination of the limitations has yet to be elucidated. Here, we show that a simplified model of a noisy integrate-and-fire neuron captures key experimental results on the role of membrane potential oscillations in both the olfactory bulb and hippocampus. We also found that an analytical model, based on the variability of inhomogeneous Poisson spikes modulated with an oscillation, captured the experimental results. The analytical model proved useful in highlighting the limitations of the spike-phase code, whereby trial-to-trial variability of spike times depends upon the precise timing relative to the phase of ongoing oscillations, as well as the firing rate of the neuron. These results show that a spike-phase neural code is beneficial only under two assumptions: [1] that spikes are precisely timed relative to the phase of oscillations; [2] that firing rates remain low. Thus, our results place fundamental constraints on a broad range of findings that report evidence for this coding strategy. In addition, we discuss two alternative coding strategies based on: [1] coupled oscillations, and [2] firing rates.

Introduction

One of the foremost problems the brain must solve is recognizing objects in the environment and telling them apart. Brains, and the neurons they are composed of, however, are biological units that are typically quite noisy. Discriminating between stimuli can be a non-trivial task because neuronal responses to the same stimulus can be variable (Mainen & Sejnowski, 1995; Shadlen & Newsome, 1998). Variability in spike time responses to sensory stimuli can stem from a plethora of sources, for example: presynaptic input, sub-threshold periodic signals from other brain regions, and noisy ionic channels (Faisal, Selen, & Wolpert, 2008). Despite numerous sources of variability, to facilitate stimulus discrimination, neurons must emit a reliable signal; yet respond uniquely to different stimuli. If the brain is a noisy biological system, what mechanism might contribute to the generation of reliable and discriminable responses to stimuli?

Oscillations

Neuronal oscillations are a prominent feature in the recordings of many brain regions and are suggested to carry functional benefits. One possibility is that oscillations provide the underpinnings for encoding information about stimulus patterns. Classic examples of oscillations in the brain are studies of place and grid cells, that show both theta and gamma oscillations in the rodent hippocampus and entorhinal cortex facilitate the animal's sense of place and navigation (Black, Nadel, & O'Keefe, 1977; O'Keefe & Burgess, 2005; O'Keefe & Dostrovsky, 1971; Zhang et al., 2014; Zhang et al., 2013). Experimental work has associated oscillations with a long list of cognitions and behavioural tasks including: attention (Clayton, Yeung, & Cohen Kadosh, 2015; Schroeder & Lakatos, 2009; Song, Meng, Chen, Zhou, & Luo, 2014), top-down processing (Engel et al., 2001), motor

processes (Mackay, 1997), working memory (Roux & Uhlhaas, 2014), episodic memory (Hanslmayr & Staudigl, 2014; Staudigl & Hanslmayr, 2013), nociception (Peng & Tang, 2016), stimulus detection (Monto et al., 2008), as well as multi-sensory integration (X. J. Wang, 2010). Further, oscillations may activate stimulus-specific cell assemblies, thus promoting the selective response of individual neurons to distinct input (Brody & Hopfield, 2003; Fries et al., 2007; Laurent, 2002; Laurent & Davidowitz, 1994; Montemurro et al., 2008; Tiesinga, Fellous, & Sejnowski, 2008), and propagating signals throughout the brain (Akam & Kullmann, 2010, 2012, 2014; Fries et al., 2007; Singer, 1999). Overall, there is strong experimental support that oscillations are correlated with a wide range of sensations, perceptions, cognitions, and behaviours.

Experimental work with single neurons highlights several examples where the presence of oscillations is correlated with an increase in performance of a task. Oscillations of the brain are often recorded in the local field potential (LFP), an aggregate signal relating to the firing of neurons in the local network. For example, theta oscillations are associated with an increase the reliability of spike timing, thus boosting performance on working memory tasks (Lee, Simpson, Logothetis, & Rainer, 2005). For the purpose of this paper, we do not concentrate on oscillatory signals in LFPs, instead we are interested in subthreshold oscillations in the membrane potential of single neurons (Domnisoru, Kinkhabwala, & Tank, 2013; Kamondi, Acsady, Wang, & Buzsaki, 1998; Margrie & Schaefer, 2003). In the olfactory bulb, oscillations enhanced the neuron's ability to respond reliably to each stimulus, in turn improving its ability to discriminate between stimuli (Schaefer et al., 2006). The authors show that this effect occurred irrespective of cell type (i.e., basket cell, CA1 pyramidal, DG granule cell, etc.), assessment method (i.e.,

distance measure *vs*. mutual information), and origin of the oscillation (i.e., current injection to soma *vs*. presynaptic spikes). In the Schaefer et al. (2006) study, correlations in responses across repetitions of the same stimulus pattern generated sharp peaks and troughs in the peri-stimulus time histograms (PSTHs), in turn enhancing stimulus discrimination. The benefit of oscillations appears to be wide-spread and shown to enhance the ability of single neurons in different regions of the brain to perform various tasks; yet the underlying mechanism has not been fully elucidated.

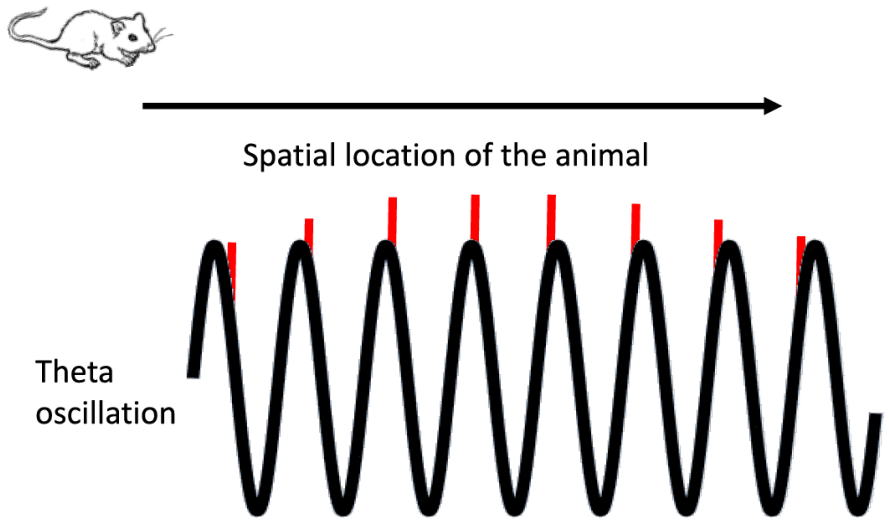

Figure 1-1. Phase precession in hippocampus depends upon precise spike timing relative to the phase of a theta oscillation. Spike times of a single neuron (red vertical lines) relate to the spatial location of the animal along one spatial dimension.

Computational studies have focused on both how oscillations can arise from the connectivity between neurons (Cressman, Ullah, Ziburkus, Schiff, & Barreto, 2009; Destexhe, McCormick, & Sejnowski, 1993; Matsuoka, 1985; Sherman & Rinzel, 1992; X.-J. Wang & Rinzel, 1992; X. J. Wang, Golomb, & Rinzel, 1995), and the resonance of oscillation frequencies (Hutcheon & Yarom, 2000). Detailed computational models of

spiking neurons have shown that oscillations enhance the coherence between proximal and distal dendritic compartments, in turn gating the propagation of signals to the soma (Remme, Lengyel, & Gutkin, 2009). The functional benefits of oscillations are widely reported, the underlying mechanisms have received less attention, and the limitations of such coding ideas are largely unknown.

Spike-phase coding hypothesis

A widely reported hypothesis is that oscillations provide the substrate for a neural code based on the timing of individual spikes relative to the phase of an oscillation, herein referred to as spike-phase coding (Engel et al., 2001; Hafting, Fyhn, Bonnevie, Moser, & Moser, 2008; K. D. Harris et al., 2002; Kayser et al., 2009; Latham & Lengyel, 2008; Lee et al., 2005; Mehta, Lee, & Wilson, 2002; O'Keefe & Burgess, 2005; Schaefer et al., 2006; Turesson, Logothetis, & Hoffman, 2012). Accordingly, oscillations may provide a 'label' for individual spikes with a phase than can ultimately be described quantitatively by degree, or qualitatively by dividing the phase into rising, peak, falling, and trough phases (Figure 1-1). One classic example of such spike-phase coding is the well-studied hippocampal phase precession, where the phase of individual action potentials relates to the animal's spatial location on a two dimensional map (K. D. Harris, 2005; Lisman, 2005). This strategy has been widely explored, for instance in the context of place cells with phase precession (Jensen & Lisman, 2000; O'Keefe & Burgess, 2005). To our knowledge, spike-phase coding has been reported in regions including the primary visual cortex (Montemurro et al., 2008), the auditory cortex (Kayser et al., 2009), the antennal lobe (Laurent, 2002), and the olfactory bulb (OB) (Fukunaga, Herb, Kollo, Boyden, & Schaefer,

2014; Margrie & Schaefer, 2003; Ranade, Hangya, & Kepecs, 2013; Schaefer et al., 2006).

In the rodent OB, stimulus-relevant information may be retrieved from the timing of spikes relative to a theta-band sniff cycle (Schaefer et al., 2006). These results were obtained *in vitro* by somatically patch clamping a neuron, injecting oscillatory currents, and delivering stimuli at various phases of an oscillation. They suggest that the ability of single neurons to discriminate between stimuli strongly depends on the phase of ongoing oscillations when stimuli are presented, an effect termed 'phase dominance' (Schaefer et al., 2006). One consequence of spike-phase coding in the OB is that pattern discrimination is limited to a specific phase of oscillation, immediately preceding the peak amplitude of each cycle. More specifically, the authors show that the ability to discriminate between neuronal responses is strongest during the trough and rising phase. Overall, these results suggest that oscillations may provide the substrate for neurons to encode information in spike times relative to the phase of ongoing local activity. Ultimately, it may be that a combination of different complementary strategies, such as firing rate and phase of oscillations, may enhance information processing beyond the capability of each strategy on its own. While this work shows that cells may encode information relative to the phase of ongoing oscillations, the limitations of spike-phase coding have not been fully elucidated.

While several models have examined the contribution of oscillations to pattern discrimination, signal propagation, long distance communication, and multisensory integration (Brody & Hopfield, 2003; Fries et al., 2007; Masuda & Doiron, 2007; Singer, 1999; H. P. Wang et al., 2010), an analysis of spike time variance accumulation is

currently lacking. Here, we examine this issue using a combination of computational and analytical methods. We employ simplified simulations of neural activity using generic leaky integrate-and-fire neurons, which capture the basic phenomenology of action potentials as well as the statistics of single neuron activity (Gerstner & Naud, 2009). To replicate experimental work (Schaefer et al., 2006), and because most neurons do not generate oscillations intrinsically, we injected a sine wave current into the membrane potential of a neuron (Steriade, 2000). The computational approach has the added advantage of a direct control over the frequency, phase, and amplitude of oscillations, while retaining fundamental aspects observed *in vivo*. As in experiments noted above, simulated cells were injected simultaneously with both an oscillatory current as well as stimuli that consisted of excitatory post-synaptic potential (EPSP-like) waveforms. Our results show that stimulus discrimination improves only under specific parameters of oscillations. Thus, our model captures several experimental findings on the functional role of oscillations and providing important constraints on a spike-phase neural code based on the phase label of action potentials. Specifically, a combination of analytical, computational, and experimental results show that such a code enhances discrimination only in a delimited portion of oscillatory phase, under assumptions of a low firing rate and precise timing of individual action potentials.

Hypotheses

i. As found in previous experiments, oscillations will improve stimulus discrimination.
ii. The phase of oscillations will play a critical role in the ability of single neurons to discriminate between stimuli.
iii. Spike timing reliability is highly dependent on the phase of oscillations.

iv. Oscillations will improve stimulus discrimination in networks of spiking neurons.

Table 1. Parameters of the leaky integrate-and-fire neuron used in simulations.

Name	Value	Units	S.D.	Description
τ	20	ms		Time constant
dt	0.05	ms		Update resolution
c_m	0.2	uF		Membrane capacitance
R	100	mΩ		Membrane resistance
V^{rest}	-65	mV	0.01	Resting membrane potential
$V^{threshold}$	-50	mV	0.01	Threshold for action potential
V^{spike}	-40	mV	0.01	Maximum potential of spike
g^{ex}	0			Excitatory conductance
g^{inh}	0			Inhibitory conductance
$\bar{g}^{ex}$	10	pS	0.04	Leak excitatory conductance
$\bar{g}^{inh}$	10	pS	0.04	Leak inhibitory conductance
E^{ex}	0	mV	0.01	Excitatory reversal potential
E^{inh}	-80	mV	0.01	Inhibitory reversal potential
I^{tonic}	30 (control)	pA	0.01	Constant input current
	0 (oscillation)	pA		
$I^{oscillation}$	4,	Hz		Oscillating input current
	60	pA		
τ^{ex}	5	Ms	0.01	Excitatory time constant for decay
τ^{inh}	10	Ms	0.01	Inhibitory time constant for decay

Methods

Leaky integrate-and-fire neuron model

The leaky integrate-and-fire (LIF) model was first proposed in a highly influential paper by Louis Lapicque (1866-1952), where he suggested that a neuron could be reduced to a parallel resistor-capacitor (RC) circuit driven by an input (Lapicque, 1907). This research was ground-breaking because it allowed computational neuroscientists to examine the effects of stimulating a neuron without experiments, and this was before we knew the mechanisms behind action potential generation (Abbott, 1999; N. Brunel & van Rossum, 2007; Johnson & Chartier, 2017, 2018).

We can decompose Lapicque's LIF model into several biological components of the membrane surrounding a cell. The lipid bilayer that partially composes the continuous membrane surrounding all cells is represented by capacitance. Here, the idea is that a certain amount of charge can be placed on the membrane of a neuron before the ions enter the cell through proteins called ion channels. The passive ion channels of the membrane are represented by resistance. Finally, the active ion channels are represented as the battery term that arises because the ionic concentrations differ inside and outside of the cell (i.e., Na^{+}, K^{+}, Ca^{2+}, Cl^{-}). Since Lapicque's original paper in 1907, theoretical neuroscientists have further developed the model to include more biological characteristics, for example: adaptation (Naud & Gerstner, 2012; Naud et al., 2008), excitatory and inhibitory conductance (Vogels & Abbott, 2009; Vogels, Sprekeler, Zenke, Clopath, & Gerstner, 2011), and synaptic connections (Thivierge & Cisek, 2011), and others. Moreover, LIF models are employed to classify single neuron recordings (Mensi

et al., 2012). The model can be further customized to suit the needs of researchers thus making it a powerful tool to generate predictions for further experimental work.

Here, a leaky integrate-and-fire model was employed whose parameters (explored below) were meant to capture fundamental aspects of single neuron activity and not the precise neurophysiology of OB cells. With our generic LIF neuron model, synaptic inputs contribute to the excitatory conductance of the cell and each incoming spike creates a slight depolarization, driving the neuron's membrane potential towards the threshold, and increasing the likelihood of observing a spike. Once the voltage crosses the threshold an action potential is generated, then the neuron is reset to resting potential, and hyperpolarized for a brief refractory period. There are two conditions: [1] the control condition, the neuron's soma will be injected with a constant current and stimulus patterns in the form of excitatory post-synaptic potentials (EPSPs); and, [2] the oscillation condition, the neuron's soma is injected with both constant and periodic (i.e., sine wave function – Equation 1.2) current as well as EPSP stimuli (Figure 1-2A & 1-2E). Spike time responses to stimuli were used to measure the neuron's ability to discriminate between incoming stimuli.

The LIF model neuron was characterized by a time constant $\tau = 20\,ms$, and the membrane potential V was updated at each time step ($dt = 0.05\,ms$). To set the scale for conductance and current, we used a membrane resistance, $R = 100\ \mathrm{m\Omega}$. Consistent with previous models (Vogels & Abbott, 2009; Vogels et al., 2011), synapses were conductance based and the neuron's membrane potential V evolved according to

$$\tau \frac{dV}{dt} = (V^{rest} - V) + g^{ex}(E^{ex} - V) + g^{inh}\left(E^{inh} - V\right) + I^{tonic} + I^{oscillation}, \qquad (1.1)$$

where V^{rest} is the resting potential (set to –65 mV), g^{ex} is the conductance for excitation, g^{inh} is the conductance for inhibition, E^{ex} is the reversal potential of excitation (set to 0 mV), E^{inh} is the reversal potential of inhibition (set to –80 mV), I^{tonic} is a constant input current (set to 30 pA for control, and 0 pA when oscillations were injected), and $I^{oscillation}$ is a time-varying oscillatory input current,

$$I^{oscillation} = A \times \sin\left(2B\pi \frac{t}{q}\right), \qquad (1.2)$$

where A is the amplitude, B is the frequency, and q is the sampling rate, set to 20 kHz. Similar to previous models (Vogels & Abbott, 2009), excitatory g^{ex} and inhibitory g^{inh} conductances decayed exponentially to zero with time constants $\tau^{ex} = 5ms$ and $\tau^{inh} = 10ms$, respectively. At the onset of each simulation, g^{ex} and g^{inh} were initialized to zero. To capture theta sniff cycles (Schaefer et al., 2006), we injected a 4 Hz oscillation with an amplitude of 60 pA; however, we show below that the main results are robust to a wide range of amplitudes.

At the beginning of each simulation, the membrane potential V (Equation 1.1) was set to V^{rest} $(-65\ mV)$. When the membrane potential reached $V^{threshold}$ $(-50\ mV)$, it was set to $V^{spike} = -40\ mV$ for 1 ms, then reset to V^{rest} $(-65\ mV)$ for an absolute refractory period of 8 ms, a feature that was added to the conductance-based model. We chose this duration of absolute refraction to precisely capture experimental results on the role of membrane potential oscillations (below); however, we acknowledge the wide variation of

absolute refraction in OB neurons reported in the literature (Carlson, Shipley, & Keller, 2000).

Consistent with experimental evidence of cellular heterogeneity in the nervous system (Buzsaki, Geisler, Henze, & Wang, 2004; Masland, 2001) we randomly selected several of the model's parameters by drawing them from normal distributions. The standard deviation ranged between 0.01 and 0.04 depending on the magnitude of the mean value (listed in Table 1). A total of 10 model neurons were generated in this fashion.

Conductance and EPSP stimuli

To match experimental protocols, stimulus patterns consisted of spike trains with a rate of ~100 Hz with Poisson-distributed spike times, and were convoluted to form an EPSP-like waveform (Schaefer et al., 2006),

$$V^{EPSP} = 0.3 * \left(1 - e^{-\frac{\Delta t}{10ms}}\right) \times e^{-\frac{\Delta t}{8ms}}, \quad (1.3)$$

where $\Delta t = F_{stim} - t$ is the difference between the current time-step t, and the previous stimulus input spike F_{stim}. Presynaptic input V^{EPSP} was computed for every time-step and added to the conductance by

$$\Delta \boldsymbol{g}^{ex} = \bar{g}^{ex}[V^{EPSP} + \Theta(\mu, \sigma)], \quad (1.4)$$

and

$$g^{ex} \rightarrow g^{ex} + \Delta \boldsymbol{g}^{ex}, \quad (1.5)$$

where $\Theta(\mu, \sigma)$ is a Gaussian function with μ = 0, and σ = 0.1 with a maximum frequency of 830 Hz. The resulting signal was multiplied by an amplitude constant of 0.2 (unit less). The inhibitory conductance was set to zero for simulations unless noted (Figure 1-4);

importantly, stimulus discrimination was similar when stimuli patterns were added to inhibitory conductance (data not shown). This noise was added to the EPSP waveform to model uncorrelated background network activity and thus generate differences across repetitions of stimulus patterns. The maximum amplitude of V^{EPSP} was set to $64\ pA$.

We devised two distinct conditions to examine the role of neural oscillations on pattern discrimination. First, in the "control" condition, the input consisted of EPSP stimuli combined with a tonic input (I^{tonic}). Second, in the "oscillation" condition, the input consisted of oscillations modeled as sine waves in addition to EPSP stimuli. While we injected oscillations as a current to capture relevant experiments, we also show that inputs can be added as a conductance signal (i.e., through the synapses) at a much lower amplitude.

Stimulus discrimination

We employed a distance metric to measure the difference between neuronal responses to EPSP stimuli. To measure stimulus discrimination, we began by convoluting all individual spikes using a Gaussian function with standard deviation of 3 ms. Then, we computed a mean peri-stimulus time histogram (PSTH) by averaging the convoluted spike trains across 10 repetitions of each stimulus. For each pair of stimuli, a and b, we estimated the mean stimulus discrimination by

$$d'_{(a,b)} = \sum_t \left| \frac{2(\bar{\boldsymbol{r}}_{t,a} - \bar{\boldsymbol{r}}_{t,b})}{\boldsymbol{\sigma}^2_{t,a} + \boldsymbol{\sigma}^2_{t,b}} \right|, \tag{1.6}$$

where $\bar{\boldsymbol{r}}_{t,a}$ and $\bar{\boldsymbol{r}}_{t,b}$ are the average PSTHs obtained at time-step t for stimulus a and b, respectively, and $\boldsymbol{\sigma}^2_{t,a}$ and $\boldsymbol{\sigma}^2_{t,b}$ denote the variance of the PSTHs. With respect to Equation 1.6, a small-sample bias arises with ten stimuli, wherein the variability is underestimated,

and thus our estimate of discrimination performance may be slightly higher. We generalized Equation 1.6 to an arbitrary number of stimuli by computing the matrix

$$\mathbf{D}^{`} = \begin{bmatrix} d^{`}_{(a,b)} & \dots & d^{`}_{(a,N)} \\ \vdots & \ddots & \vdots \\ d^{`}_{(N,b)} & \dots & d^{`}_{(N,N)} \end{bmatrix}, \tag{1.7}$$

where the elements of $\mathbf{D}^{`}$ include all pairs of stimuli $a, b, \dots, N = 10$. We recognize that one can obtain a measure of information in bits with the matrix of Equation 1.7; however, this was done in the original work, instead our analysis for networks used mean rank and results were consistent with $\mathbf{D}^{`}$ (Equation 1.7). We then obtained a discrimination score by taking the mean of lower triangular elements of $\mathbf{D}^{`}$ by

$$\overline{\mathrm{D}}^{`} = \frac{1}{h}\sum_{i<j} \mathbf{D}^{`}_{i,j}, \tag{1.8}$$

where h is the count of elements $i < j$ in $\mathbf{D}^{`}$. Equation 1.8 represents an estimate of stimulus discrimination taken over $N = 10$ stimuli presented for a repeated number of trials.

Each simulation was run using customized scripts in the Matlab computer language (2015a, The Mathworks, Natick, MA) on four Dell Precision 5600 desktops.

Results

Oscillation enhances stimulus discrimination

In a first series of simulations, we sought to capture experimental findings linking oscillations to the discrimination of EPSP stimuli in single *in vitro* mitral neurons of the OB (Schaefer et al., 2006). In the *in vitro* experiments of Schaefer et al. (2006), a single neuron was held in current clamp and injected with stimuli modeled as EPSPs (Figure 1-

2A, left panel, green and blue lines). During stimulus presentations, the neuron was also injected with an oscillation of 4 Hz frequency, directly mimicking the theta sniff cycle observed *in vivo* (Figure 1-2A, right panel, red lines). By comparison to the control condition, oscillations drastically increased the reliability of spikes across repetitions of stimuli (Figure 1-2B), resulting in sharp peaks and troughs, as well as lower variability in the PSTH averaged across trials (Figure 1-2C, shaded area). Thus, when the neuron was injected with oscillations, stimulus discrimination was markedly increased compared to the control condition without oscillations (Figure 1-2D).

As discussed in the 'Methods' section, to capture these results with the LIF model, we simulated the activity of a single conductance-based neuron that received stimuli in the form of Poisson spike trains (convolved into EPSPs, Equation 1.3 & 1.4), as well as subthreshold oscillations (4 Hz frequency and 60 pA amplitude, Equation 1.2) (Figure 1-2E). As with the *in vitro* experiments, the reliability of APs over repeated presentations of the stimuli markedly improved with the addition of oscillations (Figure 1-2F). This reliability resulted in sharp peaks and troughs, as well as lower variability in the PSTH obtained by averaging spikes across trials (Figure 1-2G, shaded area). Thus, the injection of oscillations yielded heightened stimulus discrimination compared to a control condition with no oscillation (Figure 1-2H). Thus, we closely replicated the experimental results where oscillations are beneficial compared to a neuron injected with constant current.

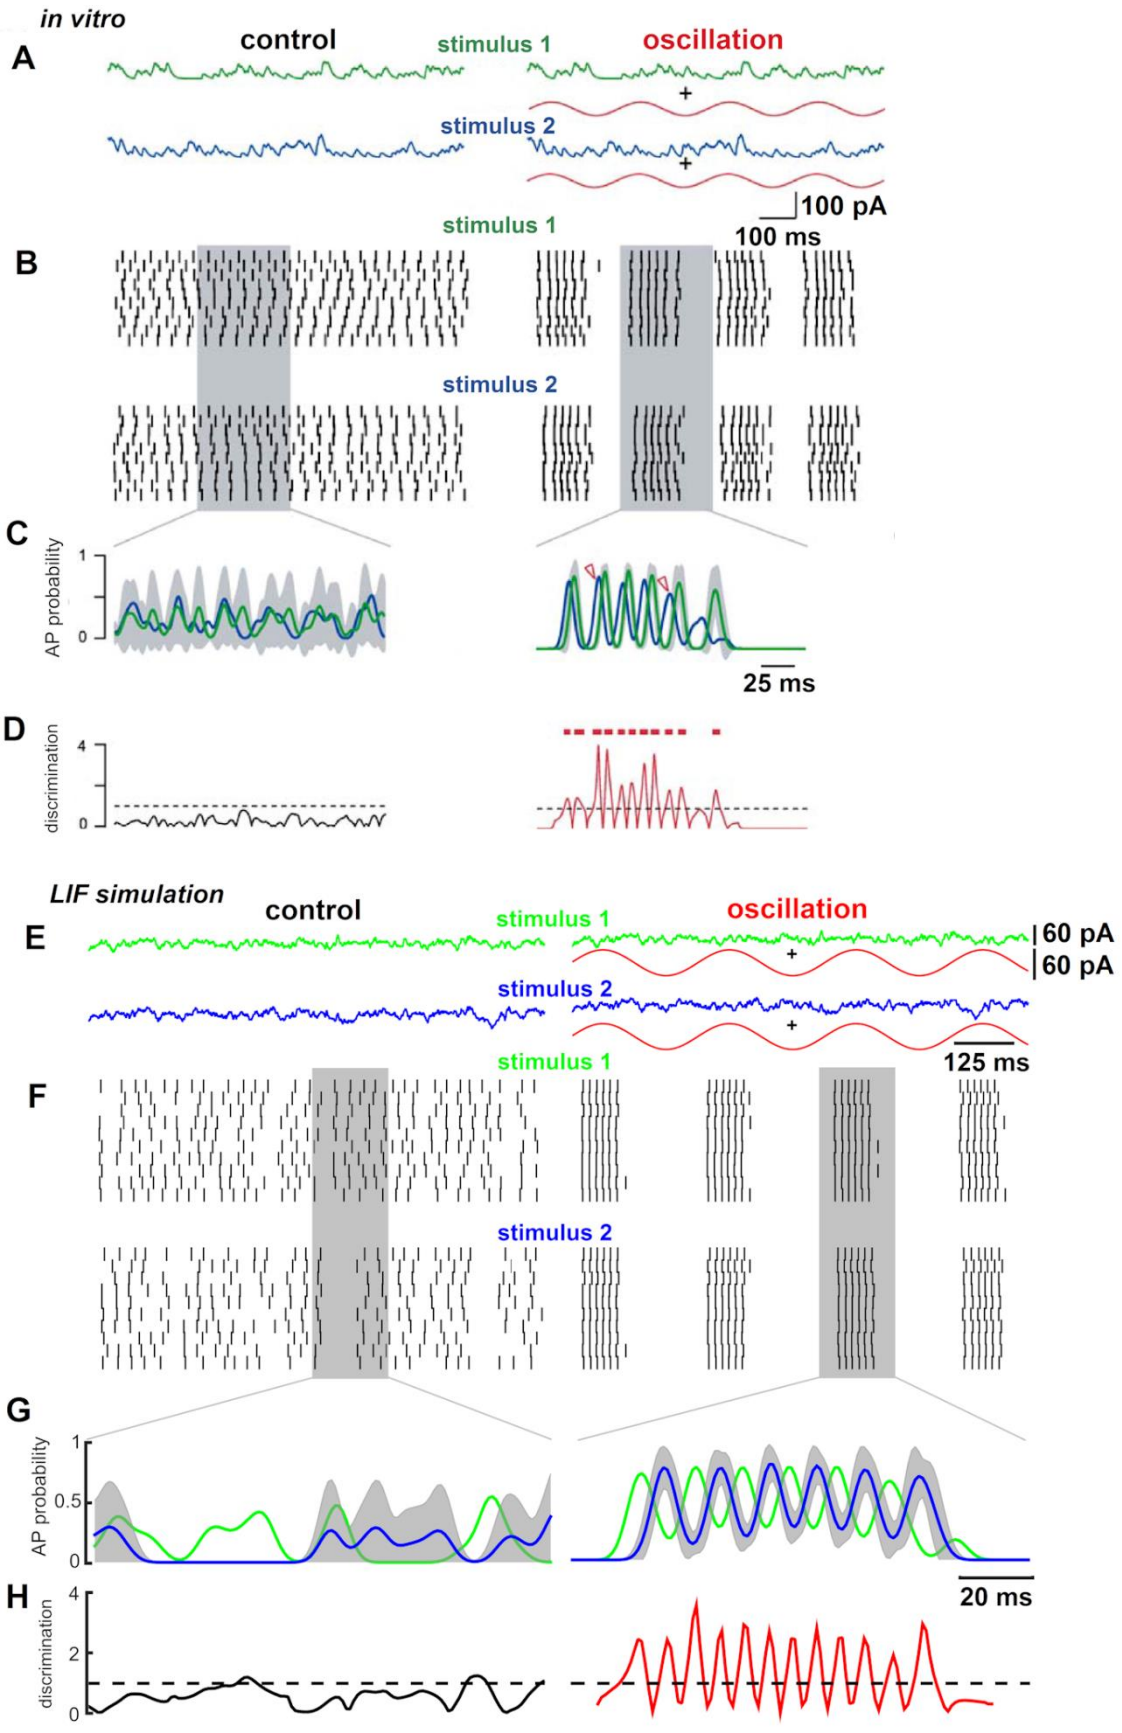
in vitro
A
control
stimulus 1
oscillation
stimulus 2
100 pA
100 ms
B
stimulus 1
stimulus 2
C
AP probability
1
0
25 ms
D
discrimination
4
0
LIF simulation
E
control
stimulus 1
oscillation
60 pA
60 pA
stimulus 2
125 ms
F
stimulus 1
stimulus 2
G
AP probability
1
0.5
0
20 ms
H
discrimination
4
2
0

Figure 1-2. Oscillations enhance stimulus discrimination in experiments and LIF model. A. Green and blue traces, examples of continuous stimuli from Schaefer et al. (2006). Red trace, 4 Hz membrane potential oscillation. **B.** Spike time responses to two stimuli each repeated ten times. Shaded area, zoomed in area in C and D. **C.** PSTHs for responses to two stimuli, corresponding to an average across ten trials. Shaded area, standard deviation of responses to stimulus 2. PSTHs were generated with a 5 ms time window. **D.** Stimulus discrimination (Equation 1.6). Dashed black line, discrimination value of 1. **E-H.** Simulation results that capture experimental results of A-D.

The advantage of oscillations for pattern discrimination was found across a range of simulation parameters, including the frequency and amplitude of oscillations (Equation 1.2; Figure 1-3A & 1-3B, respectively), the amplitude of the EPSP stimuli (Equation 1.3; Figure 1-3C), the firing rate of the stimulus spike trains (Equation 1.3; Figure 1-3D), the amplitude of Gaussian noise added to the stimuli (Equation 1.4; Figure 1-3E), the amplitude of oscillations added to the conductance of the neurons, a scenario that is closer to *in vivo* conditions (Figure 1-3F), and when inhibitory inputs were added to the stimulus pattern (Figure 1-4). Further, stimulus discrimination in the model was markedly superior to a null model where the firing rates were preserved; but spike times were randomly permuted between times $1 \ldots T$, where T equals 2 seconds sampled at 20 kHz (Figure 1-3A through 1-3F, solid gray lines). However, despite the advantage of oscillations across a broad range of parameters, we found several fundamental limitations on a neural code based on the timing of spikes relative to ongoing oscillations.

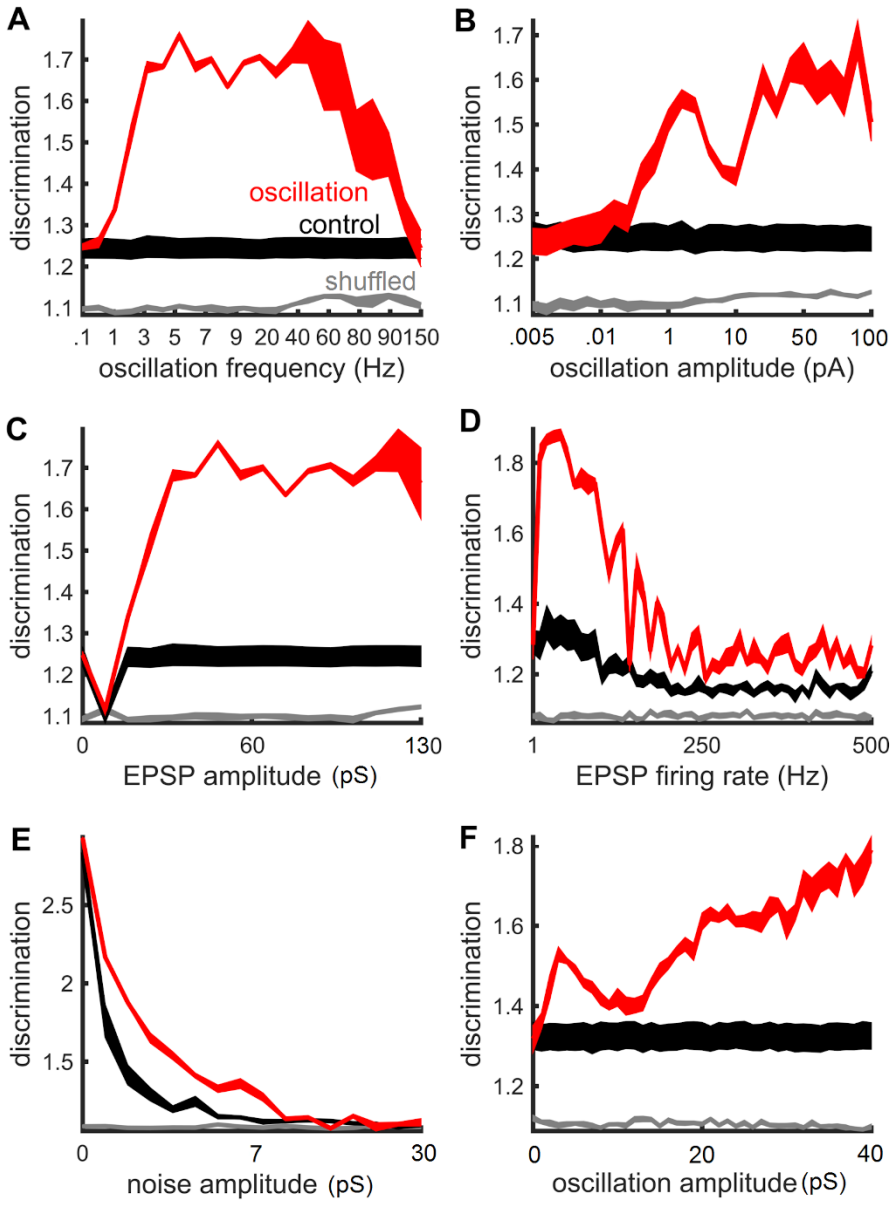

Figure 1-3. Stimulus discrimination is robust to a range of oscillation and stimulus parameters in the LIF model. A. Stimulus discrimination for various frequencies of oscillation. 'Control' neurons receive no oscillations. 'Shuffled', randomly permuted spike trains from the 'oscillation' condition. Thickness of lines is SEM. **B-F.** Stimulus discrimination for various amplitudes of oscillation added to the membrane voltage, EPSP stimulus amplitude, stimulus firing rate, amplitude of Gaussian noise added to the conductance (Equation 1.4), and amplitude of oscillations added to the

membrane conductance, respectively. In A-F, discrimination is measured by Equation 1.6. All values were averaged over the performance of 10 neurons with parameters drawn from Gaussian distributions (see Methods).

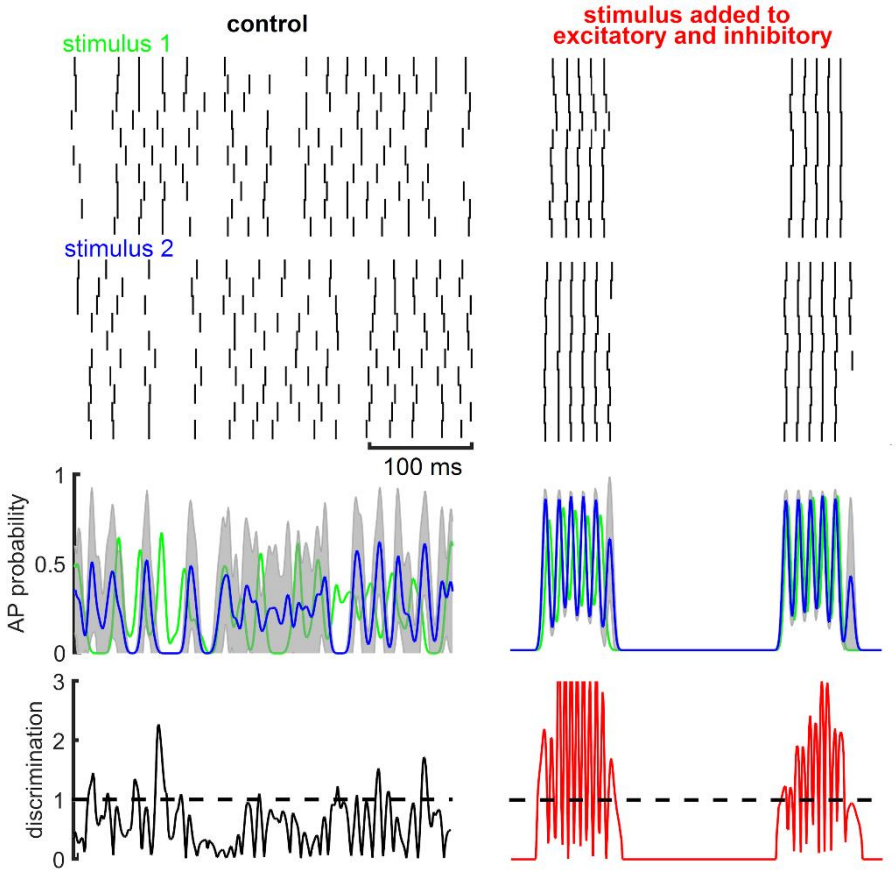

Figure 1-4. Oscillations enhance stimulus discrimination when inhibitory postsynaptic potentials (IPSPs) were added as stimuli. Stimulus discrimination when a neuron was injected with both EPSP and IPSP stimuli. Both control and oscillations (4 Hz, 32 pS) conditions are presented.

Phase of oscillations is critical to stimulus discrimination

The above results were obtained with continuous stimuli that covered the entire cycle of a given oscillation (i.e., all phases). In a realistic scenario, stimuli can be brief, lasting only

a few milliseconds, and thus would 'arrive' at a specific phase of oscillations. Further *in vitro* experiments with somatic voltage recordings of the OB have focused on brief stimulus patterns that covered only a portion of an oscillatory cycle, i.e., targeting specific phases of the oscillation (Figure 1-5A). In these experiments, the advantage of oscillations in discriminating between stimuli was limited to trials where the stimulus patterns were injected during the rising phase of a wave (Figure 1-5B). In fact, oscillations had no advantage when stimuli were injected during the falling phase, where discrimination was not possible (Schaefer et al., 2006).

To examine the relation between stimulus discrimination and the phase of oscillations in the neuron model, brief EPSP stimuli lasting 70 ms were injected (which translates into 100 degrees of the 4 Hz oscillation). In one condition, stimuli were delivered during the rising phase of oscillations (onset at the 288th degree); in a second condition, stimuli were delivered during the falling phase (onset at the 108th degree). These conditions mimicked those of the *in vitro* OB and hippocampal somatic recordings (Schaefer et al., 2006). Consistent with experiments (Figure 1-5B), stimuli presented during the rising phase of oscillations yielded markedly higher discrimination than stimuli presented during the falling phase (Figure 1-5C). In addition, the phase effect was robust to increases in both the strength of the noise added to stimulus patterns (Equation 1.4, Figure 1-6A *vs*. 1-6B), and the length of Gaussian windows used to generate the PSTHs (10 ms, as opposed to 3 ms, Figure 1-6C). In both experiment and simulation, only the rising phase of an oscillation was beneficial in generating reliable neuronal responses, highlighting a critical limitation of encoding stimuli with a spike-phase code.

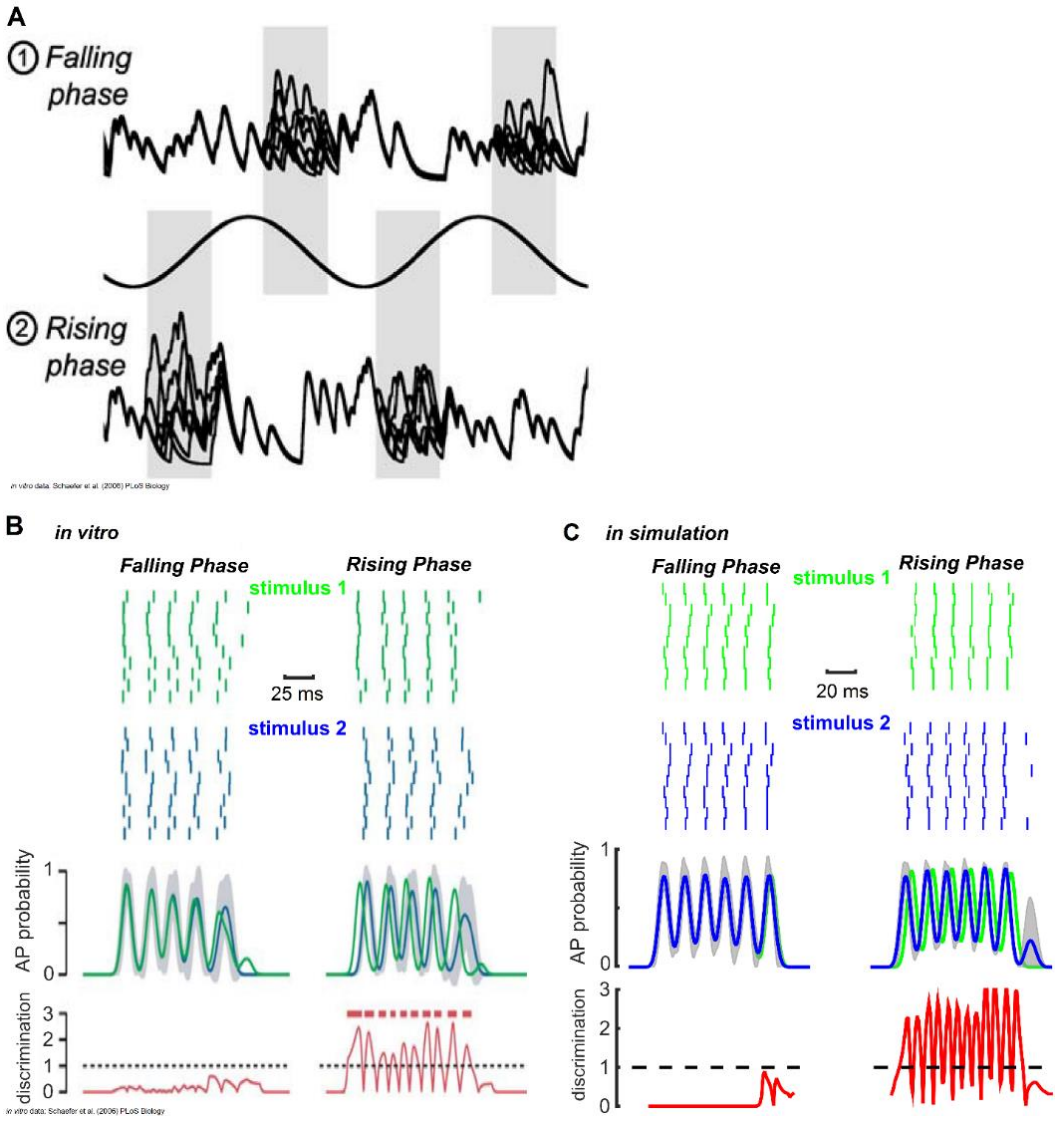

Figure 1-5. Stimulus discrimination depends critically upon the phase of oscillations. A. Schematic display of experimental data from Schaefer et al. (2006). A given cell of the OB receives one of two stimuli that differ only in the shaded area, aligned to either the rising or falling phase of ongoing oscillations. **B**. From Schaefer et al. (2006), OB cells in vitro show a phase preference effect whereby discrimination is only observed with stimuli presented during the rising phase of oscillations. Values depicted follow from Figure 1-2. **C**. Simulated neurons captured the phase preference

effect observed in the OB (see Methods for all default values employed to run these simulations).

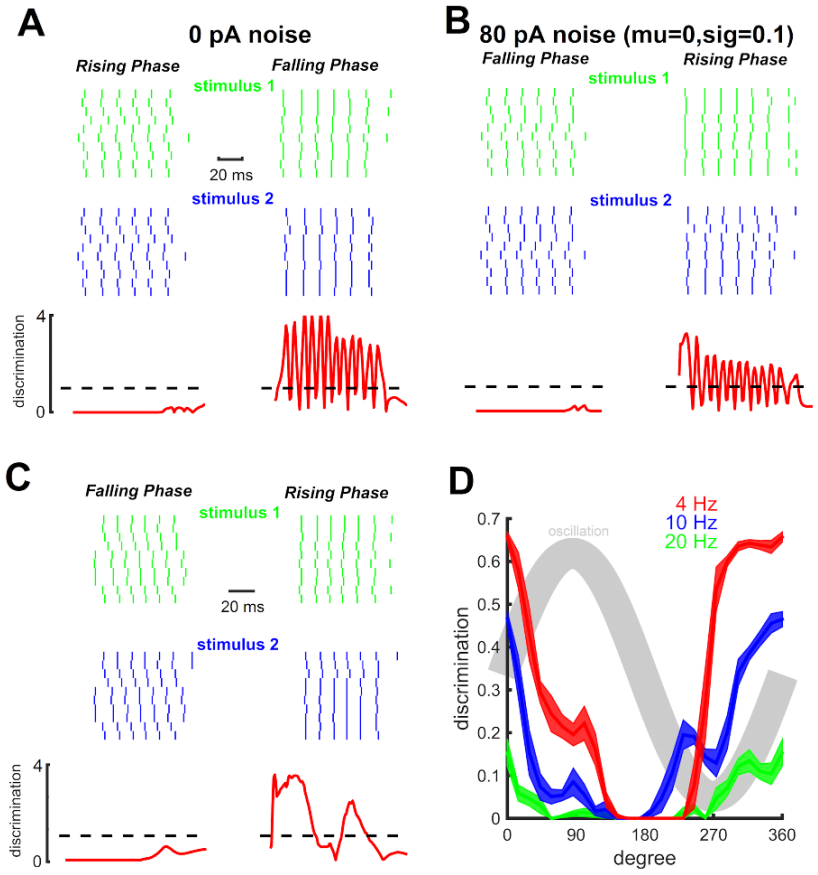

Figure 1-6. Phase effect is robust to changing the strength of noise, and duration of smoothing function. A. Stimulus discrimination functions for three different frequencies of oscillation, denoted by colours. In the background is an oscillation to denote the specific phases where discrimination was strongest. **B**., **C**., & **D**. Plots like that of Figure 1-5 (that had low noise) with: zero Gaussian noise; and strong noise (80 pA); and longer Gaussian windows used to generate PSTHs (10 ms).

We then examined this phase effect across the entire cycle of an oscillation for several different frequencies. This was achieved by taking the brief stimulus patterns above (i.e.,

70 ms duration) and sweeping the onset of presentation across the oscillation at intervals of the period of the wave divided by 25. For instance, a 4 Hz oscillation has a period of 250 ms, thus the interval between phases would be 10 ms. Simulations were run for 4, 10 and 20 Hz. Stimulus discrimination was strong during the rising phase for each frequency injected (Figure 1-6D). Further, the phase effect is significantly reduced for frequencies ≥20 Hz; however, this may be an artefact of the simulation parameters, because the duration of stimuli was set to 70 ms and this covers the entire period of oscillations faster than ~15 Hz. Overall, the phase effect matched the experimental data and robust to many changes in the model, including adjustments to the time-course of oscillation cycles. Matching experiments, these results show that only certain phases of an oscillation can contribute to stimulus discrimination; however, this was not the only limitation of encoding with a spike-phase code.

Phase of oscillation and spike-timing variability

We consider the hypothesis that the phase effect is due to differences in spike-timing variability across distinct phases of an oscillation, as reported in the OB and hippocampus (Schaefer et al. 2006). Specifically, variability (shown above in the spike trains) increases over the course of the oscillation due to the gradual accumulation effect (Figure 1-7A). To illustrate this effect, consider an oscillating current that is injected into a single neuron over several trials where each trial consists of a single cycle of oscillation. The gradual accumulation of variance is readily observed when aligning all trials of a given stimulus by the phase of oscillation (Figure 1-7B) and computing the variance of spike times for each consecutive spike (Figure 1-7C, left panel). Thus, spike time variability accumulates gradually between the rise and fall of oscillations, albeit in a nonlinear fashion.

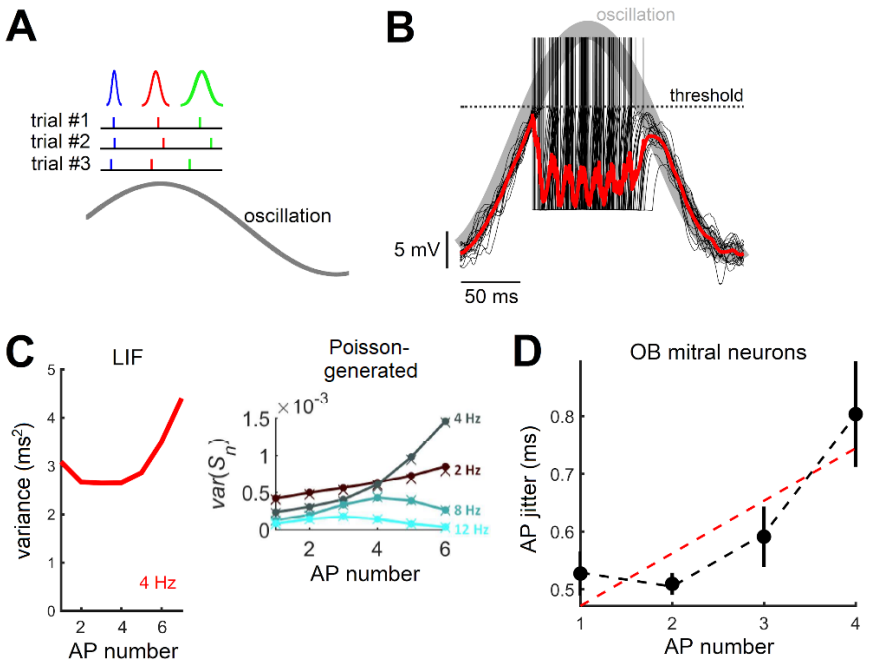

Figure 1-7. Variance accumulation in LIF simulation, a theoretical framework based on analytical model, and *in vitro* experiments. A. Cartoon illustration of jitter accumulation across consecutive spikes generated during a membrane potential oscillation. Colors indicate consecutive action potentials for a single neuron over 3 trials where activity is aligned to the onset of the oscillation (shown in black). **B.** In simulations of LIF neurons, we aligned spike across 10 cycles of oscillation according to the phase of the first spike. Dashed line, firing threshold of the neuron. Black, membrane potential across individual cycles of oscillation. Red, average subthreshold membrane potential. Grey, depiction of the oscillation to show the phase relative to individual spikes. **C**. Left, spike variance for consecutive spikes along a cycle of oscillation. Right, the variance across consecutive spikes obtained with an oscillatory firing rate of a given frequency (2-12 Hz). Values on the y-axis were rescaled by a constant factor for ease of visualisation. For both middle and right panel, spike times correspond to the average timing of action potentials from left panel. **D**. OB neurons display a nonlinear

accumulation of jitter across consecutive spikes. Black circles, mean jitter in OB cells (obtained from Schaefer et al. (2006) – Figure 2B). Dashed red line, best linear fit using a first-order polynomial. Dashed black line, best fit using a second-order polynomial. Vertical lines, SEM.

A similar effect was observed in recordings of mitral cells of the OB, where a U-shape function characterizes the relation between consecutive action potentials and variability (defined as the average distance of each spike to the mean timing of the n^{th} consecutive spiking event) (Figure 1-7D). To characterize this relation, we fitted the experimental data using either a first- or second-order polynomial. This statistic includes a penalty for the number of terms in the model and is appropriate for comparing how different models fit the same data. The adjusted correlation is defined as $r^2_{adjusted} = 1 - (SS_{resid}/SS_{total}) \times ([n-1]/[n-d-1])$, where n is the number of observations in the data, and d is the degree of the polynomial (a linear fit has a degree of 1 and a quadratic fit of 2), SS_{resid} is the sum of squared residuals from the regression and SS_{total} is the sum of squared differences from the mean of the dependent variable. An adjusted goodness-of-fit that accounted for the degree of the polynomial function showed that the data was better fit with a second-order (nonlinear) polynomial ($r^2_{adjusted} = 0.72$) compared to a first-order (linear) polynomial ($r^2 = 0.67$). Thus, both the LIF model and experimental data show a nonlinear accumulation of spike time variance across an oscillatory cycle; this increase in variance serves as an explanation for the phase dominance effect, whereby stimulus discrimination is more accurate when stimuli were presented during the rising phase.

We were curious whether the reduction in variance between the first and second spike was significant. Thus, we ran the same simulation for several trials, where the initial

conditions were different (i.e., connectivity), and other parameters were fixed (i.e., λ). The overlap in 95% confidence intervals ($p = 0.05$) suggest that indeed the difference in jitter from the first to the second spike is not statistically significant (Figure 1-7E). However, our current work focuses on the overall shape of the function relating AP number to jitter. Further, our model predicts that the difference in variance between first and second spike is dependent on oscillations frequency (Figure 1-7C, 12 Hz *vs*. other), a topic we revisit in the discussion below. Our main point is that this function is best captured by a quadratic than a linear fit, even when using an adjusted measure of fit that considers the degrees of freedom of this fit.

The idea that the accumulation of variance may be nonlinear, compared to linear, is an important distinction. The U-shaped function that we observed may result in the variability of spike times increasing at an exponential rate, much faster than a one-to-one linear accumulation, making it impossible to discriminate between stimuli with spikes that arrive after the peak phase. This issue would be intensified with higher firing rates than those observed above. Further, the predicted variance of the nonlinear function for the first spike is higher, compared to the linear, and this distinction will help highlight the limitations of encoding stimuli presented during the rising phase of oscillations.

Limitations of spike-phase coding

To thoroughly characterize the limitations of encoding information about stimuli in spike times relative to an oscillation, we employed a Poisson model that captured the main result of a U-shaped variance function (i.e., Figure 1-7A). Compared to the LIF model, with the Poisson model it is straightforward to both precisely control the spike times and increase firing rates. We chose to focus on Poisson spike statistics because they offer a

close fit to the experimental data considered here and have a minimal set of assumptions. With the Poisson model, we can place the spikes in time relative to an oscillation's phase and examine the variance of the distribution with either shifted times or heightened firing rates. We allowed this process to be inhomogeneous with time-dependent intensity $\Lambda(t)$. We denote the arrival time of spike n as S_n. To find the variance of S_n, we begin by defining $\Lambda(t)$ as the expected number of spikes between $(0, t)$:

$$\Lambda(t) = \int_0^t \lambda(s)ds\,. \tag{1.9}$$

The mean firing rate is then

$$r = \lim_{t:\infty} \frac{\Lambda(t)}{t} \tag{1.10}$$

The probability density function of S_n can be found by the product of the probability of having $n-1$ spikes up to time t and of having a spike in the interval $t, t+h$:

$$f_{S_n}(t) = \Lambda(t)e^{-\Lambda(t)}\frac{[\Lambda(t)]^{n-1}}{(n-1)!}, \tag{1.11}$$

From f_{S_n} it is straightforward to compute the mean and variance of S_n:

$$E[S_n] = \int_0^\infty t f(S_n)dt, \tag{1.12}$$

$$var(S_n) = \int_0^\infty t^2 f(S_n)dt - E[S_n]^2. \tag{1.13}$$

In the case of the homogeneous Poisson process, these integrals lead to the well-known expressions $E[S_n] = n/\lambda$ and $var[S_n] = n/\lambda^2$.

Our interest lies in the case of sinusoidal rate modulation, i.e., $\lambda(t) = A\sin(\omega t + \phi) + k$. More specifically, to prevent the rate from becoming negative, we consider the half wave

rectified version of this rate in which the phase ϕ and the offset k, i.e., $\lambda(t) = A\sin(2\pi f t)$ for $0 \leq t \leq 1/(2f)$ and zero otherwise. The spike count after half a period $T/2$ of the oscillation is given by $\Lambda(t/2) = A/(\pi f)$, and the average rate over the same duration is $r = 2A/\pi$. With this choice of intensity function, we can evaluate $f(S_n)$. However, this integral must be evaluated numerically, and this was done with Matlab using the trapezoidal integration function.

Like that of experimental recordings of the OB and our LIF model, variance of spike timing with the Poisson model increases non-linearly across consecutive spikes (Figure 1-7C, right panel). Like the theta sniff cycle and our LIF simulations, with an oscillation of 4 Hz, spike variance for the first few spikes ($n = 1,2,\&\ 3$) is markedly lower than the last spikes ($n = 6\ \&\ 7$).

We generalized the above results to different frequencies of oscillation ranging between 2 and 12 Hz. The analytical model thus predicts that the phase preference effect strongly depends upon the frequency of oscillations, which matches both the experimental data (i.e., Figure 8C2, from Schaefer et al. (2006)), and our LIF simulations (Figure 1-6). These results highlight another critical limitation of spike-phase coding, specifically that the benefit of this code may only generalize to frequencies of oscillation around 4 Hz, and this prediction may be realized through further experimentation.

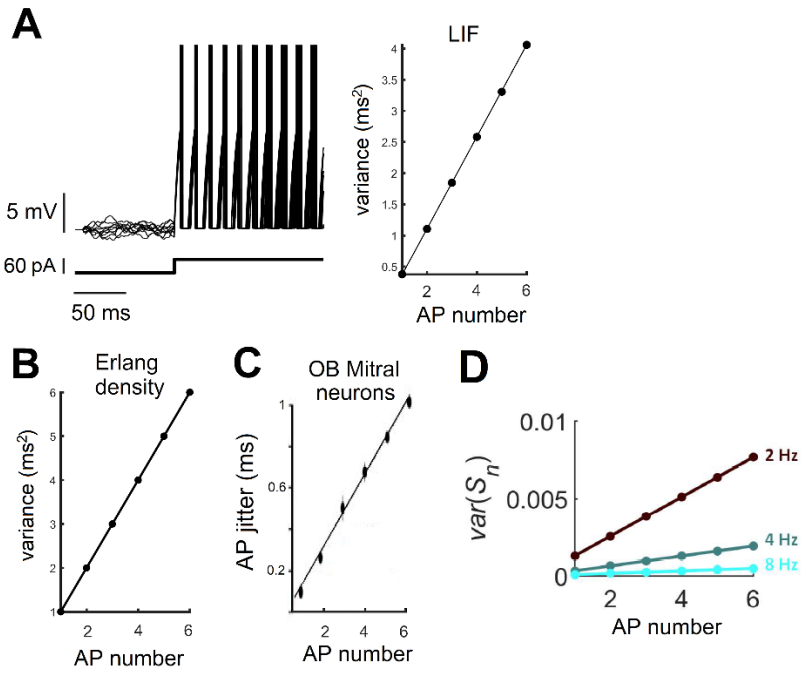

Figure 1-8. Variance accumulation with tonic (i.e., constant) current injections. A. Right, In the LIF model, a tonic input (60 pA) gives rise to a gradual accumulation in spike variance. Here, 10 simulations are shown, where membrane potential voltages are aligned by the first spike after stimulus onset. Left, Variance of spike times for the LIF model showing linear accumulation. **B**. Variance of an inhomogeneous Poisson process showing that a tonic input yields a linear accumulation of variance across consecutive spikes. **C**. Linear variance accumulation in OB cells receiving a tonic input (Schaefer et al., 2006). Black circle, average jitter across trials; solid line, best-fitting regression. **D**. Variance of spike time intervals in the analysis with constant firing rate (λ).

A U-shaped function relating spike time variance to consecutive spikes was no longer observed with the LIF model when substituting oscillations by a constant input (Figure 1-8A, left). Like the results of Schaefer et al. (2006), variance accumulation with the Poisson model followed a linear function (Figure 1-8B & 1-8C). In our analysis, this effect was a

direct result of variance under constant rate, whereby $\sigma^2 = n/\lambda^2$ increased linearly with consecutive spikes $n = 1,2,3,\ldots$ when λ remained constant. In this simple case, one prediction is that it would be possible to compensate for an increase in variance caused by the n^{th} spike interval by gradually scaling the firing rate according to $\sqrt{n}$, thus preventing the accumulation of variance altogether (Figure 1-8D).

Overall, our analytical analysis accounts for both constant and oscillating current injections using a simplified model of spike timing variance under the influence of a time-varying firing rate. Through this analysis we arrived at several critical limitations to the spike-phase coding hypothesis: [1] spike must be precisely timed; [2] firing rates must remain low; and, [3] may only be possible with low frequencies of oscillation.

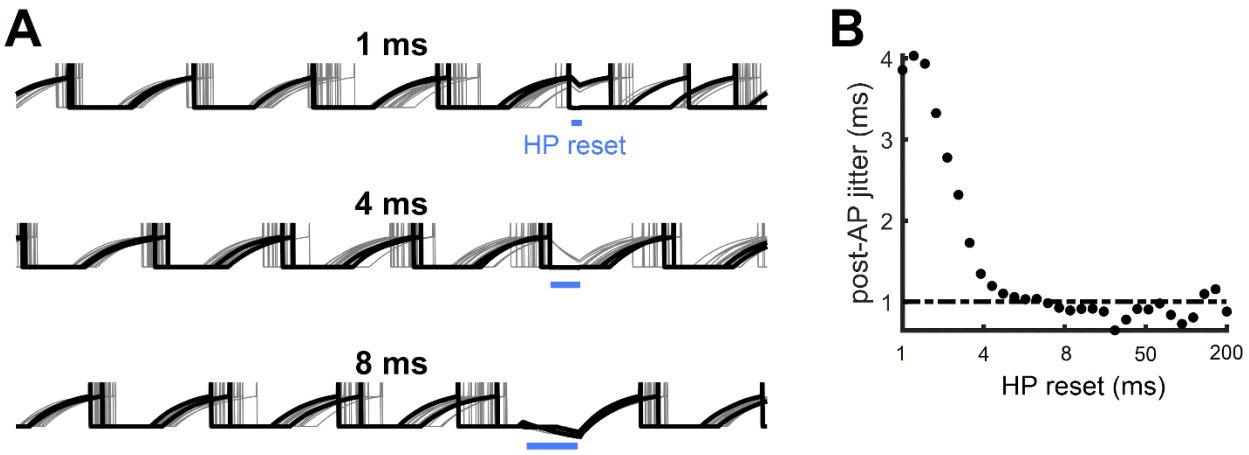

Figure 1-9. Effect of hyperpolarizing reset on the variance in spike time. A. Membrane potential in response to a tonic current (60 pA) that is reset to a value of -60 mV for 1 ms. Grey lines and two black lines, membrane voltages taken from 10 simulated neurons. **B.** Mean jitter of the first spike following the reset of the tonic current. Recovery intervals correspond to the duration of the tonic current reset. Dashed line, jitter of 1 ms (denoted as an accurate spike time in Schaefer et al. (2006)).

Hyperpolarizing reset critical to stimulus discrimination

Importantly, all the analytical results reported here assume that spike time variance does not continue to accumulate through the trough of oscillations. This is a reasonable assumption that is corroborated by *in vitro* and *in vivo* recordings (Schaefer et al., 2006). However, it is worth examining further how spike variance may be reset from one oscillatory cycle to the next. Experiments have considered the role of hyperpolarization (occurring in the trough of oscillations) in a simplified scenario where a step current was followed by a hyperpolarizing interval (Schaefer et al., 2006).

We captured this result with the LIF model, where the injection of a step current, strong enough to force firing activity, was injected followed by a short hyperpolarizing interval, strong enough to suppress firing (Figure 1-9A). Short hyperpolarizing intervals led to heightened variance in the first post-interval spike, while longer intervals beyond ~5 ms reduced the variability in spike timing (Figure 1-9B).

This value is comparable to the time constant of the model (i.e., the time necessary for the membrane potential to reach 63% of its maximal amplitude from its resting potential), calculated from Equation 1.1 to be 20 ms. Thus, the duration of the hyperpolarizing interval contributes to spike variance, a factor that plays a significant role in the ability of single neurons to perform pattern discrimination. Importantly, the duration of this hyperpolarizing interval must be long enough, considering the time constant of the neuron, to reset the membrane potential, thus resetting the spike-time variance.

Coupled frequencies of oscillation

The above results on the accumulation of spike variance from both the LIF and Poisson model depend upon the assumption that membrane oscillations can be approximated by

a simple sine function. More complex forms of oscillations have been reported in various studies, and combine lower- and higher-frequency components that are coupled by their phase or frequency (Jensen & Colgin, 2007; Williams, 1992). We examined this scenario by injecting cells with a coupled oscillation (phased coupled between 4 and 40 Hz) (Figure 1-10A). Here, simulations of the LIF model suggest variance does not accumulate over the course of the slower oscillation (Figure 1-10B). Thus, it may be possible to reduce the effect of variance accumulation by controlling the statistics of oscillations. While a more complete parameter exploration is needed to examine the robustness of this effect, these results nonetheless show that the dynamics of membrane potential oscillations influence spike variance, and that variance accumulation may be reduced under certain conditions.

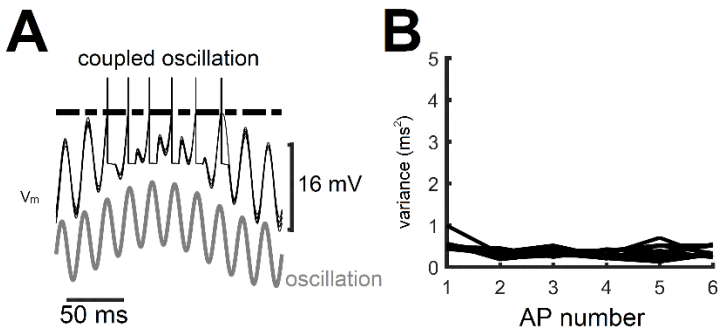

Figure 1-10. Effect of coupled oscillations on variance accumulation. A. Response of a neuron receiving a phase-coupled oscillation composed of a lower (4 Hz, 60 pA) and a higher (40 Hz, 60 pA) frequencies. Solid black line, membrane potential of a neuron; solid grey line, oscillation; dashed line, firing threshold. **B**. With coupled oscillations, there is no variance accumulation over consecutive spikes. Solid black lines, 10 trials where each trial is one period of the lower frequency oscillation, for a neuron receiving a coupled oscillation.

Two alternative hypotheses

We briefly consider two alternative hypotheses to explain the phase preference effect observed in both experiments and LIF simulations. First, the effect may be driven by differences in spike rates across rising and falling phases of an oscillation. This hypothesis, however, was not observed in our simulations, where spike rates were comparable across both phases that stimuli were presented (Figure 1-5, spike rasters).

Second, the amplitude of oscillations may account for the phase preference effect, given that this amplitude is lower at the onset of the rising phase compared to the falling phase. We tested this possibility by injecting a model cell with a tonic current of comparable amplitude to both the rising and falling phases (~30 and ~60 pA, respectively) at the time of stimulus onset. Results show an increase in pattern discrimination with higher amplitude of tonic input, contradicting the idea that lower amplitude at stimulus onset leads to increased discrimination (Figure 1-11, A *vs*. B). Furthermore, we found that stimulus discrimination was heighted for a delimited range of current inputs (Figure 1-11C).

Oscillation enhances stimulus discrimination in neuronal networks

Using a similar model to the one above (Kuebler, Bonnema, & Thivierge, 2013; Kuebler & Thivierge, 2014; Thivierge & Cisek, 2008, 2011), we extended the effect of oscillations on stimulus discrimination to neuronal networks. We sought to determine whether spike-phase coding was possible, thus, we presented stimuli only at either the rising phase, or the falling phase of oscillations. When stimuli were presented at the rising phase, we found that oscillations enhanced the ability of neuronal networks to discriminate between stimulus patterns for a wide range of frequencies and amplitudes of oscillation (Figure 1-

12A & 1-12B, respectively). However, when stimuli were presented at the falling phase of oscillations, stimulus discrimination was not possible for a wide-range of frequencies and amplitudes (Figure 1-12C & 1-12D). These results are similar to those reported above, where stimulus discrimination is enhanced for a wide-range of frequencies and amplitudes; however, only when stimuli are presented at the rising phase of oscillations.

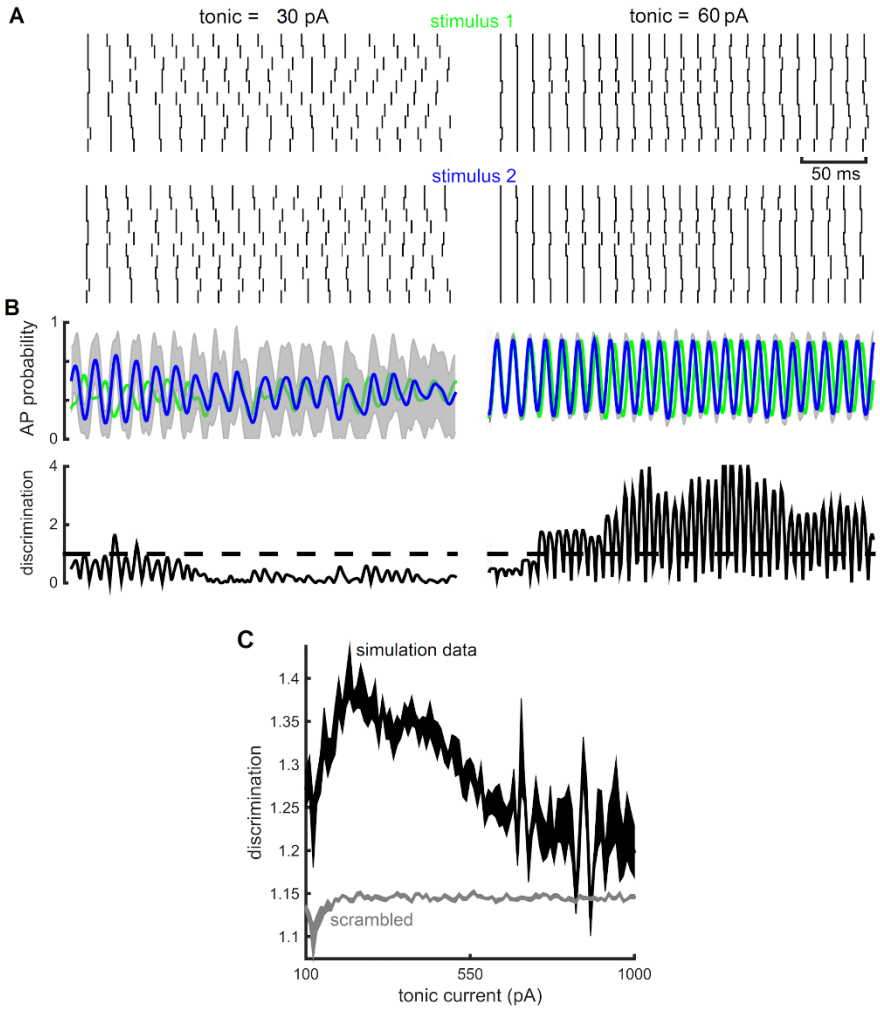

Figure 1-11. Effect of tonic input on pattern discrimination in the LIF model. A. In these simulations, oscillations were replaced with tonic input (60 pA). Top two panels show spike rasters over 10 trials in response to stimulus 1 and 2. Middle panel, PSTH obtained from the above rasters. Shaded area, SEM of stimulus 2. Bottom panel,

Pattern discrimination. **B.** Same as panel A, using a higher tonic input (120 pA). **C**. Stimulus discrimination as a function of the tonic input injected in the soma of the LIF neuron. Scrambled is a null control representing shuffled spike times.

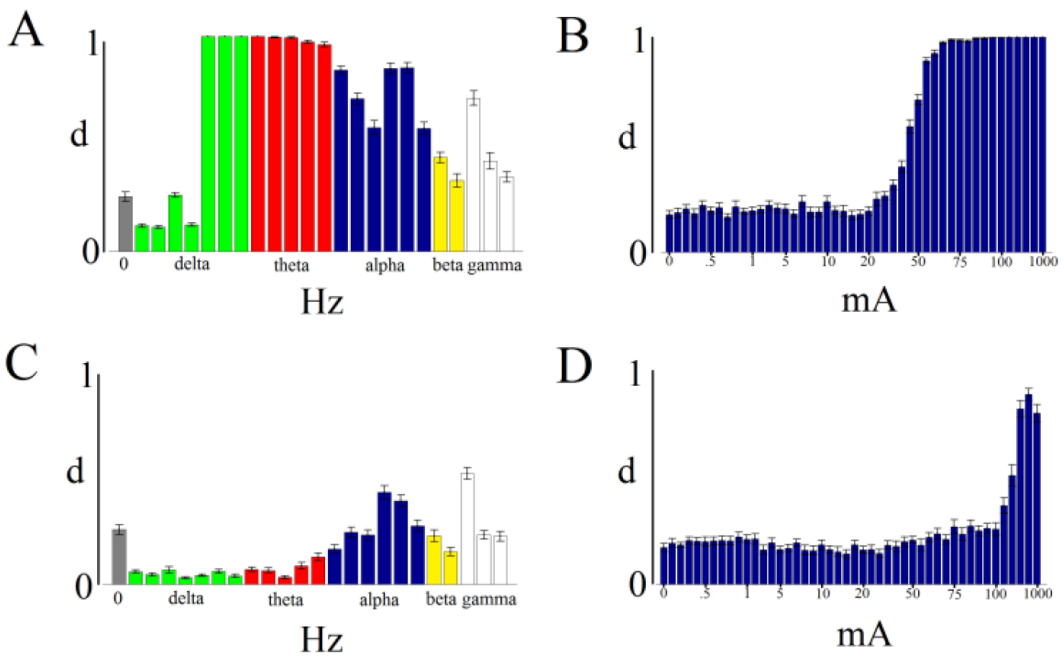

Figure 1-12. Discrimination as a function of frequency and amplitude in networks of spiking neurons. A. Discrimination as a function of frequency of oscillation with amplitude held constant (5 mA). Gray: no oscillation. Green: delta waves (0.2, 0.4, 0.5, 0.8, 1, 2, & 3 Hz, respectively) Red: theta waves (4, 5, 6, 7, & 8 Hz, respectively). Blue: alpha waves (9, 10, 12, 15, 16, and 18 Hz, respectively). Yellow: beta waves (20 & 25 Hz, respectively). White: gamma waves (30, 35 & 40 Hz, respectively). **B.** Discrimination as a function of amplitude with frequency held constant (10 Hz). **C.** Discrimination as a function of frequency (same Hz as 2A) with amplitude held constant (5 mA). **D.** Discrimination as a function of amplitude with frequency held constant (10 Hz). For A. & B., stimuli were presented at π (see Figure 1B of Kuebler et al. (2013)). For C. & D., stimuli were presented at 2π (see Figure 1B of Kuebler et al. (2013)).

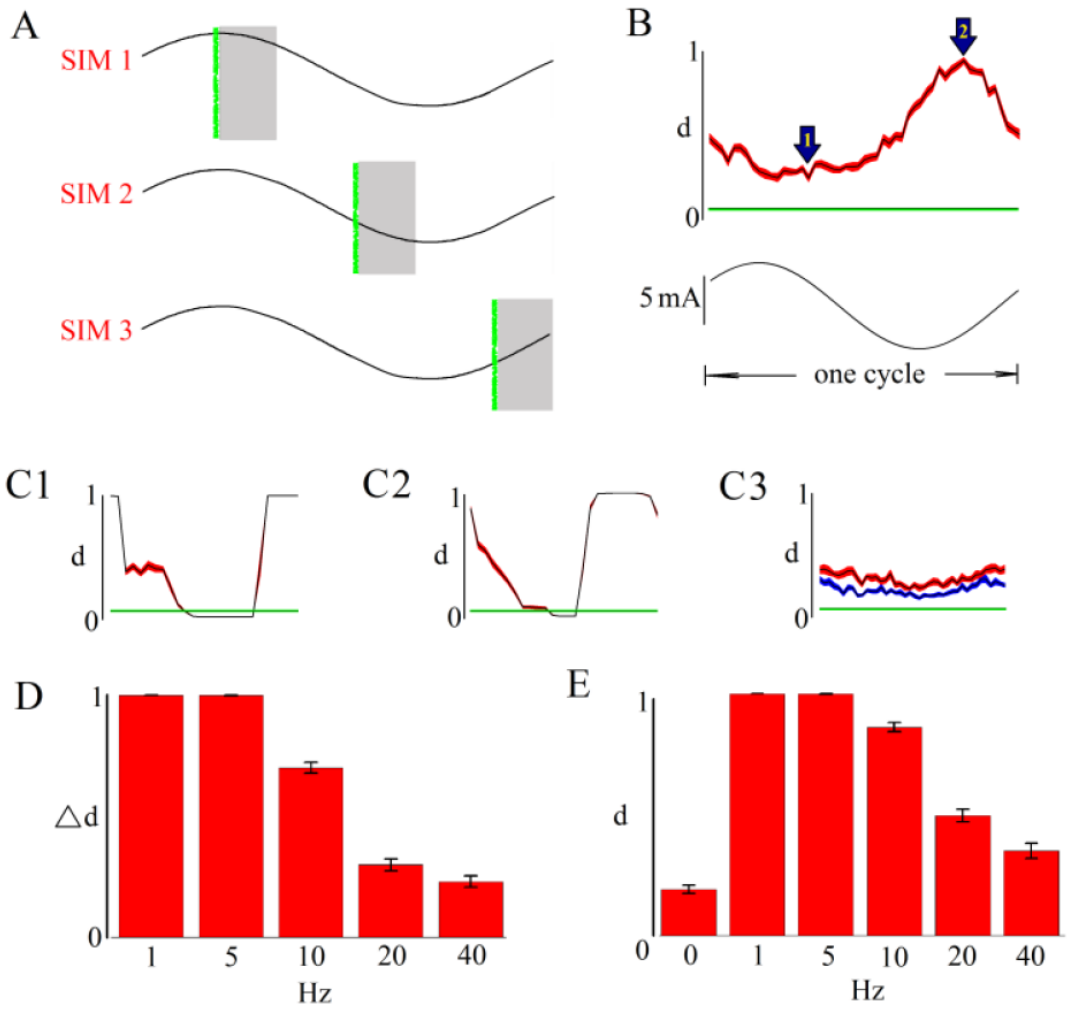

Figure 1-13. Network sensitivity to the timing of stimuli relative to the phase of oscillations – the phase dominance effect. A. Cartoon depicting how the timing of stimuli was shifted across an oscillation. This timing was held constant within a given simulation but changed between simulations. Green: timing of randomly generated stimuli (2 ms duration). Gray: post-stimulus response window (100 ms). **B.** Discrimination curve showing how the timing of stimuli relative to the phase of oscillations affects discrimination. Red curve (10 Hz alpha oscillation) represents the mean discrimination of the network and shading shows standard error. Amplitude was held constant (5 mA). Arrows: peak (#2) and trough (#1) of the discrimination curve.

Green: chance performance. **C.** Red represents the mean discrimination of the network and shading shows the standard error. Oscillations were injected at the following bands: 1. Delta (1 Hz), 2. Theta (5 Hz), and, 3. Beta (25 Hz). Blue curve: mean discrimination score with a 40 ms response window. Amplitude was held constant (5 mA). **D.** Difference between discrimination scores when discrimination was highest and lowest for each band, this effect is termed 'phase dominance'. **E.** Highest discrimination scores for each band of oscillation. 0 Hz corresponds to networks with no oscillations. Vertical bars: SEM.

Spike-phase coding with neuronal networks

With the same network of spiking neurons used above, we examined the effect of oscillation phase on discrimination, specifically which phases promoted the network's ability to discriminate between stimuli, and whether this effect is found with higher frequencies of oscillation. Simulations were carried out in the same fashion as above; however, here we slide the stimulus presentations across the entire period of the oscillations (i.e., each phase, Figure 1-13A). Results show that when neurons were injected with an oscillation of 10 Hz frequency the networks' strongest discrimination was between the through and rising phase (Figure 1-13B). We found a stronger phase effect for lower frequencies of 1 and 5 Hz, where discrimination scores ranged from 0 to 1 (Figure 1-13C1 & 1-13C2, respectively). Like experiments and our analyses above, with a higher frequency of 25 Hz, the phase effect was relatively weak, and discrimination was approximately equal for each phase (Figure 1-13C3). This trend was highlighted by the difference between the maximum and minimum discrimination scores for each frequency, where by comparison to high frequencies, lower frequencies had higher difference scores

(Figure 1-13D). Compared to the control condition (0 Hz), discrimination was heighted for several frequencies (Figure 1-13E). Overall, we suggest that oscillations are useful to a neuronal network's ability to discriminate between stimuli; however, like single neurons, the phase of oscillation is critical to discrimination.

Discussion

In this work, we investigated the advantages and limitations of a spike-phase neural code using experimental, simulation, and analytical analyses. With the LIF model we first showed that the benefit oscillations provide towards stimulus discrimination may be wide-spread across various parameters of the model, including the statistics of oscillations and EPSP stimuli. The LIF and Poisson model also captured the phase effect reported in experiments, whereby stimulus discrimination is improved with spikes that occur during the rising phase of oscillations (Schaefer et al., 2006; Turesson et al., 2012). An elegant and simple explanation of this phase preference effect was provided by an analysis of inhomogeneous Poisson spikes under a sine wave oscillation, and we found this analysis useful in revealing the constraints of a spike-phase neural code. Specifically, only a few of the spikes produced during a single oscillation period enhanced pattern discrimination compared to a model with no oscillations. In this scenario, there are considerable time intervals where oscillations confer no advantage for encoding stimuli, thus dampening the possibility for discrimination. We further explore the hyperpolarizing reset that oscillations provide, and two additional coding hypotheses. Finally, we suggest that the problem of non-linear variance accumulation may be reduced when oscillations are coupled (i.e., fast oscillation on top of a slower frequency). Overall, we highlight a few of the critical

limitations associated with a spike-phase neural code including: precise spike timing, low firing rates, and low frequencies of oscillation.

That oscillations generate a non-linearity in the variance of spike times carries significant functional consequences that, in turn, generate questions about our knowledge of hippocampal function. Our analyses show that when an oscillation is injected, only a few of the spikes emitted during a cycle have lower variance than a model with no oscillations, and these events typically occur during the rising and peak phase. Phase precession in the hippocampus involves a spike-phase code, whereby the timing of spikes is modulated according to the rodent's location in a 2D environment. Pairing these two ideas, our work suggests that spike times are only accurate when a rodent is in the 'preferred place' (i.e., spikes near the peak of oscillation), thus, only the neurons that fire at the peak of oscillations would contribute the rodent's knowledge of their location. To understand how a rodent hippocampal function, these two hypotheses need to be reconciled.

Our model predicts that the difference in variance between first and second spike is dependent on oscillations frequency. If we increase this frequency from 4 Hz (default) to 24 Hz, variance appears nearly parabolic (Figure 1-7C). As a result, the difference between variance for the first and second spikes is now markedly more pronounced. Therefore, our analysis predicts that at higher frequencies than the one tested by Schaefer et al (2006), we may see a marked drop in variance from the first to the second spike in an oscillation cycle. This predicts warrants experimental support from future studies.

The limitations of a spike-phase code may be overcome with more complex oscillatory functions such as phase-coupled oscillations, i.e., theta and gamma (Figure 1-10). Phase-

coupled oscillations, specifically a faster gamma on top of a slower theta oscillations, have been reported in the hippocampus (Belluscio, Mizuseki, Schmidt, Kempter, & Buzsaki, 2012; Canolty et al., 2006; Jensen & Colgin, 2007). However, these phenomena have not been observed in the OB to the best of our knowledge. Therefore, some regions of the brain may not make use of this solution to address the accumulation of spike variability. The frequencies we chose were 'hand-picked' to provide a simple proof of concept, the true benefit of coupled waves would be severely limited by the specific combinations of fast and slow frequencies and amplitudes. Given the restrictions of spike-phase coding with one oscillation, it may seem implausible that coupled oscillations can be useful; however, uncovering the limitations of coding with these complex functions requires further study.

Both our theoretical work and related experiments highlight important restrictions with spike-phase coding, thus it is useful to consider some alternative hypotheses. Most prominently, a neural code based on firing rates seems a likely candidate (Abbott, 1994; London et al., 2010). Rate codes are based on the idea that all information about the precise spike-timing is removed, leaving only the frequency of spikes, thus reducing the importance of considering the accumulation of variance with an oscillation. Indeed, our simulation results and the experiments we've modeled show that pattern discrimination in the OB is possible even when applying a broad (100 ms) filter to single spikes (Schaefer et al., 2006), implying that the specific spike times may not be important. Further firing rates are easily amenable to a simple linear decoder (Abbott, 1994; Salinas & Abbott, 1994), making it questionable whether the brain needs or uses a more complex neural

code. These two features pose a significant challenge to the spike-phase coding hypothesis.

Many of the theoretical and experimental results described here involve stimuli that are continuous in nature (Figure 1-2). It is worth considering what effect this may have on our conclusions on the role of oscillations. With a continuous stimulus spanning many cycles of oscillation, a portion of the stimulus will inevitably fall within the rising phase of the oscillation, where stimulus discrimination is enhanced. However, with a discrete stimulus spanning a limited time interval, this may not be the case. It is possible that the stimulus would not be present during the rising phase, thus in this scenario we expect that oscillations would be of little benefit to spike time variance. Unless the phase of oscillations is time-locked to the stimulus (Cui, Liu, McFarland, Pack, & Butts, 2016), it may be difficult to experimentally detect a positive effect of oscillations on spike variance and stimulus discrimination. The onset of stimuli in our environment can be highly unpredictable, making it less likely for precise time-locking to occur, and this poses a significant issue for the spike-phase code.

Our work provides a starting point from which several further questions may be addressed. First, our study did not examine the impact of oscillations in the context of a neuron where modifications in synaptic strength can alter spiking activity (Sussillo & Abbott, 2009). It is possible that synaptic plasticity enhances stimulus discrimination by adjusting connections between neurons to amplify differences between stimuli (Mazarakis et al., 2005). Further, synaptic plasticity may serve to shift spikes along the phase of an oscillation such that they become more reliable, as may be the case with phase precession (K. D. Harris, 2005). It is possible that synaptic plasticity may improve

the accuracy of spike-times, thus benefiting a spike-phase code; however, to our knowledge this has not been shown.

Second, our work started on the encoding of stimuli with single neurons and finished by extending these results to networks of spiking neurons. To simplify this process, we used an ensemble code that was generated by summing the activity of neurons (Naud & Gerstner, 2012). We did not, however, examine how various network population codes may help resolve some of the limitations of spike-phase coding discussed here. Neural networks are a complex issue, with many parameters to include or alter and numerous connectivity schemes to explore. Further, efficient population codes in the OB and other sensory regions remains an area of intense study and debate (Cury & Uchida, 2010; Friedrich & Stopfer, 2001; Miura, Mainen, & Uchida, 2012). Future research may uncover population codes that allow networks of neurons to overcome the limitations of non-linear variance accumulation.

Finally, we cannot yet provide a full theoretical treatment to examine why certain frequencies of oscillation lead to higher pattern discrimination than others (Figure 1-3A). The answer will likely involve a combination of spike timing, total number of spikes near the peak of oscillations, and length of hyperpolarizing interval. While we have examined each of these factors as individual components, we did not combine them in a way that would allow us to examine their optimal trade-off, which would involve an exhaustive search through a large space of potential solutions.

In conclusion, because neural oscillations are such a wide-spread phenomena throughout the central nervous system, researchers are eager to provide a mechanistic understanding of their functional role in sensory processing and motor control (Buzsaki &

Draguhn, 2004; Engel et al., 2001; Fries et al., 2007; X. J. Wang, 2010). Our work highlights critical limitations, generates key predictions, and provides a foundation for further research to explore the contribution of oscillations on brain function and behaviour. While the search for explanations has led to hypotheses on how the brain employs neural codes that combine spike timing and the phase of an oscillation, one must tread carefully in assessing the benefits and limitations of such coding strategies.

Predictions from our model

i. Spike-phase coding can be explained by a non-linear accumulation of spike time variance as a function of the phase of oscillations.

ii. The benefit of a spike-phase code is limited by the precise timing of action potentials relative to the phase of an oscillation. Specifically, shifting the spike times across the phases of an oscillation can result in changes in the variance of subsequent spike times.

iii. The benefit of spike-phase coding is limited by the firing rate of the neuron. Specifically, the variance of the spikes emitted during the falling phase will be greater for high firing rates compared to low firing rates.

Conclusion

Neurons, and or, networks of the brain may use various strategies to overcome the difficulty of performing computations in a noisy biological environment. Above, we have highlighted a scenario where oscillations are useful in generating spike *time* responses that help discriminate between stimuli. Importantly, we also show that spike *rate* responses were useful. In Chapter One, there is no assumption that information about stimuli was coded in a periodic (i.e., across several periods of oscillation) *versus* non-

periodic fashion. Thus, we consider the reliability of spiking activity within network bursts that occur in a non-periodic fashion, to uncover the other strategies that neurons may use.

Chapter 2 – Burst Prediction Neurons of *In Vitro* Cultures: Characteristics and Survival in Simulated Cerebral Ischemia

Abstract

Recordings of neuronal activity *in vivo* and *in vitro* exhibit network bursts characterized by brief periods of increased spike rates flanked by silent periods. Experimental and computational work shows that a subpopulation of neurons reliably predicts the occurrence of network bursts; yet relatively little is known about these neurons or their resistance to insult. Here, we examined some characteristics of burst predictors as well as their role in cultures undergoing an *in vitro* model of cerebral ischemia. We plated dissociated primary cortical neurons on multielectrode arrays and recorded the spontaneous electrical activity beginning at 17 days *in vitro* (DIV). Like related work, the spiking activity of DIV 17, 20 and 21 was characterized by neuronal avalanches where the statistics of bursts (i.e., counts, intervals, number of active electrodes) followed a power law with a slope of ~1.5. We then identified burst predictors, defined as electrodes that consistently fired immediately prior to network bursts. The spike timing of these predictors relative to bursts followed a skewed distribution that differed sharply from a null model based on a branching ratio. Results show that burst predictors may interact prior to bursts of activity, suggesting that there may be a presynaptic link between them. Further, by comparison to DIV 17, at DIV 21, predictors fired in closer temporal contiguity

to network bursts. This may be due to an increase in the density of synaptic connectivity as shown in other experimental work. Importantly, a portion of cultures were subjected to an excitotoxic insult at DIV 18. By pairing a propidium iodine stain with fluorescence imaging, we confirmed that ~65% of the neurons died in a high insult scenario. Importantly, the insult did not alter the distribution of avalanches. While there were alterations of overall spike rates; burst predictors, however, maintained baseline levels of firing activity. Burst predictors may contribute to neuronal computation, especially since the resilience of burst predictors following excitotoxic insult suggests a key role of these units in maintaining network activity following injury. These results have implications for both neuronal coding and the selective effects of ischemia in the brain.

Introduction

Network bursts are a common feature of several *in vivo* and *in vitro* preparations of the central nervous system. They are characterized by short transients of rapid spiking activity flanked by silent periods. Cultures of cortical neurons begin to emit network bursts around day *in vitro* (DIV) 6 and gradually develop a rich repertoire of activity during the course of maturation (Wagenaar et al., 2006). Experimental work has shown that network bursts may be useful in stimulus encoding and the propagation of information (Plenz, 2012, 2013; Shew & Plenz, 2013), as well as the formation of central pattern generators that significantly contribute to firing activity in cortex (Butera et al., 1999a). There are many statistics of network bursts including; but not exclusively: the number of channels active during a burst, the duration of bursts, and the number of action potentials in a burst. Each of these statistics of spontaneous network bursts are known to follow a power law distribution, where lower/small values of the abscissa will have higher frequencies

compared to higher/larger values. For instance, the frequency of short duration network bursts is greater by comparison to those of longer durations. Theoretical work has shown that the scaling exponents of power law distributions are suggestive of a critical state that promotes information processing (Beggs, 2008) (with some controversy surrounding the issue (Touboul & Destexhe, 2010).

Burst predicting neurons

Experimental and theoretical work has characterized subsets of neurons that reliably predict the occurrence of network bursts (Eckmann et al., 2008; Eytan & Marom, 2006; Zbinden, 2010, 2011). Experimental work first characterized burst predictors as 'early-to-fire' neurons, that spiked approximately tens of milliseconds prior to network bursts, and the authors suggest that these cells could have a role in the propagation of signals throughout the cortex (Eytan & Marom, 2006). In similar experiments, a subset of 'leader' neurons were found to encode information about network bursts, specifically they track the number of action potentials per neuron (Eckmann et al., 2008). Further, they show that 'leader' neurons have a half-life of 24-34 hours, suggesting they have relatively stable dynamics. In addition, the distribution of leaders may be altered with stimulation; however, the distribution returns within approximately 1 hour, implying that the functional state of the network is maintained despite transient changes (Eckmann et al., 2008). Computational work has focused on the synaptic connectivity of 'leader' burst predicting neurons, characterizing them as densely connected to postsynaptic excitatory targets, and sparsely connected to presynaptic excitatory, as well as postsynaptic inhibitory cells (Zbinden, 2010, 2011). Overall, neurons that predict bursts of activity may be central to

network function; however, their characteristics have not been fully elucidated, and their resilience to insult is unknown.

Because GABAergic interneurons play a role in the orchestration of network activity and help shape developing networks (Bonifazi et al., 2009), their survival following a neural injury may be crucial to the maintenance of bursts, and thus the propagation of information throughout the cortex. In the hippocampus, GABAergic interneurons with exceptionally dense axonal arborisation, termed ‘hub’ neurons, show a tendency to fire prior to the onset of network wide bursts (Bonifazi et al., 2009). Further, the authors show that stimulating a single hub cell can generate a network burst, which implies a causal role to these neurons (Bonifazi et al., 2009). By comparison to principal (pyramidal) cells, GABAergic neurons of the hippocampus can have greater resistance to an ischemic insult (Schwartz-Bloom & Sah, 2001; Wong et al., 2018). Taken together, these experimental results infer that burst predicting neurons of our *in vitro* cultures may be less susceptible to a neurotoxic insult; however, this link, to our knowledge, has never been directly established.

Here, we examined how *in vitro* network activity is altered by an excitotoxic glutamate insult that captures key features of oxygen-glucose deprivation during ischemia (Tauskela et al., 2008). Glutamate is an excitatory neurotransmitter found in great abundance in the central nervous system that is associated with plasticity (Meldrum, 2000). Glutamate can also have excitotoxic negative effects, whereby neurons can suffer from an over-abundance of glutamate concentrated in the synaptic cleft, and remain in a state of constant excitation, in turn, the neuron eventually succumbs to excitotoxicity and dies. When applied in solution to a cell culture, this insult results in wide-spread cell death that

was evident in the number of neurons that become positive for propidium iodide (PI – marker of cell death), and further supported by a reduction in firing rates.

We cultured dissociated cortical cells on multielectrode arrays (MEAs, Figure 2-1A) and extracted multiunit activity (MUA) from each channel (Figure 2-1B), which reflects the combined spiking activity of neurons surrounding the channel (Tauskela et al., 2008; Thivierge, 2014; Vincent et al., 2013; Vincent et al., 2012). We examined baseline neuronal activity at DIV 17 to identify channels that predict network bursts (Figure 2-1C). To assess the viability of burst predicting cells, a portion of the cultures were subjected to a glutamate insult, where on DIV 18 we added ouabain-TBOA to the extracellular fluid. While the ischemic insult caused marked cell death, both neuronal avalanches and burst predictor firing rates were largely maintained. These findings have implications for both neural coding and the functional recovery of neuronal networks following cerebral ischemia.

Hypotheses

i. There exists a subset of neurons that reliably fire action potentials ~50 ms prior to network-wide bursts, namely ‘burst predictor neurons’.
ii. Burst predictors fire co-ordinated action potentials prior to a network burst.
iii. Firing activity of burst predictors will occur in closer temporal contiguity to network bursts later in development (i.e., DIV17 *vs*. DIV21).
iv. The glutamate insult will cause wide-spread cell death in the cultures, measured by PI staining, and confirmed by a reduction in the firing rates on some electrodes.

v. The burst predicting electrodes will be active post-insult, suggesting that the burst predicting neurons remain alive; while others have perished.

Methods

Multielectrode arrays (MEA)

Ionic channels embedded in the membrane of neurons allow the transfer of ions in and out of the cells, in turn changing the voltage inside of the cell. Importantly, the voltage of neurons can be measured both intracellularly, for example with a patch clamp, or extracellularly, with a recording electrode in the extracellular matrix. Extracellular electrodes record the change in voltage that is generated by the ionic currents of single neurons in the surrounding region of about $50\ \mu m$ (Buzsaki, 2004). There are two types of MEAs that can be employed to make extracellular recordings: [1] *in vivo*, where electrodes are surgically implanted directly onto the brain tissue; and, [2] *in vitro*, where electrodes are in a dish and neurons are cultured in various conditions. We chose to focus on *in vitro* MEAs because we can easily and accurately control the conditions of insult by washing solutions in and out of the cultures after a given duration.

There are several factors that contribute to a recording: [1] the medium (i.e., extracellular fluid); [2] how tightly the neurons are packed around the electrodes (i.e., multi-layer, monolayer); [3] the electrode material (i.e., impedance); [4] signal processing (i.e., gain, bandwidth); and, [5] parameters of data acquisition (i.e., sampling rate), to mention a few. Each of the issues are addressed below, and have been addressed in previous work (Langlois et al., 2014; Tauskela et al., 2008; Thivierge, 2014; Vincent et al., 2013; Vincent et al., 2012). In each MEA, there are 60 electrodes, one is assigned as a ground (electrode #15), leaving $N = 59$.

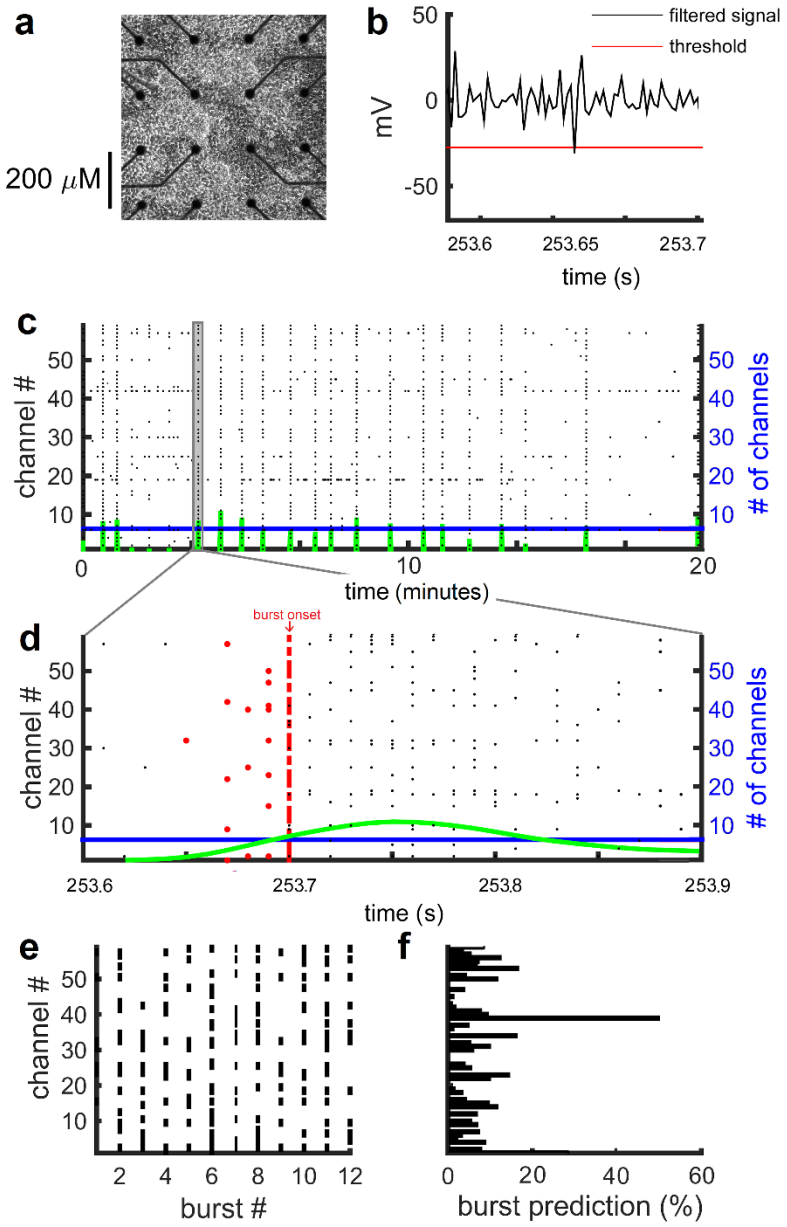

Figure 2-1. In vitro network bursts on multielectrode arrays. A. A subset of channels on a multielectrode array (MEA) imaged under phase contrast conditions to allow the discernment of neuronal somas. The 4 x 4 grid shown here was in the center of the cell culture. **B**. Typical recording of multi-unit activity (MUA) surrounding an electrode (high-pass filtered at 200 Hz). The red trace displays the MUA detection threshold for this channel. **C**. Raster plot showing MUA events recorded at each

channel for the entire 20 min recording. Grey shaded area is zoomed in below. **D**. Same as c; but zoomed in to show a single network burst (~300 ms). Red circles indicate channels identified as burst predictors (i.e., channels that spiked <50 ms before a network burst). Green, average population activity. Blue, threshold for the detection of network bursts (see Methods). **E**. Raster plot showing channels that recorded at least one MUA event before a burst in c. **F**. Burst predicting scores for the data in e.

Cultured neurons on multielectrode arrays

All experiments were approved by the Human Health Therapeutics Animal Care Committee at the National Research Council Canada (NRC, Ottawa, Canada) and carried out in accordance with approved guidelines. Culturing and plating of primary cortical/hippocampal neurons was performed as previously described (Vincent et al., 2013). Briefly, time-pregnant embryonic day 18 Sprague-Dawley rats (Charles River, St. Constant, Quebec, Canada) were anesthetized with halothane and culled by cervical dislocation. Following dissection of the cortical and hippocampal regions of the fetal brains, cells were centrifuged at 1,000 *g* for 3 min at 4^{o}C and were dispersed by gentle trituration. Cells were plated on MEAs at a high density of 1.8×10^6 cells/ml of medium (consisting of EMEM (Wisent) supplemented to 25 mM glucose), 10% fetal bovine serum (PAA), 10% horse serum (Sigma-Aldrich, St Louis, U.S.A.), and pen/strep (Gibco 1X). The entire surface of each MEA dish (Multi-Channel Systems, Reutlingen, Germany) was pre-treated with a high molecular weight poly-L-lysine (0.025mg/mL diluted in 1×PBS; Sigma-Aldrich) and laminin (0.02mg/ml; Gibco) to promote cell adhesion and minimize cell migration. To better approximate neurotoxic insult conditions previously employed in our laboratory, cell suspensions were plated at 1 ml/MEA so that the entire MEA surface

(and not just the inner channel region) was coated with cells (Tauskela et al., 2008; Vincent et al., 2013). Cultures were maintained in a humidified incubator at 37ºC with a 5% carbon dioxide / 95% air atmosphere. Osmolality was strictly controlled by daily addition of distilled water and by covering MEAs with a transparent, gas-permeable polydimethylsiloxane lid to prevent evaporation (Blau, Neumann, Ziegler, & Benfenati, 2009). To minimize glial cell proliferation, a mitotic inhibitor (15 µg/ml of 5-fluoro-2'-deoxyuridine and 35 µg/ml of uridine) was added at 4 *days in vitro* (DIV). In addition, MEAs were gently rinsed and media was filtered starting at 7 DIV and repeated every 4 days. A 50% media (containing 10% horse serum but not fetal calf serum) change was performed at 7 DIV, with 33% media changed every 3 DIV thereafter.

Recordings of spontaneous neuronal activity

Immediately prior to a recording session, MEAs were removed from the maintenance incubator, capped with a sterile vented lid and placed in the acquisition platform housed within a 37ºC incubator. A 20-min equilibration period was allowed for minimization of movement-stimulated activity. Spontaneous activity was recorded in 20 min sessions at 17-18 and 21-22 DIV using MC Rack software (Multi-Channel Systems) employing the following settings: unit-less amplifier gain (1100.0), input voltage range (+/– 2048 mV) and acquisition rate (5 kHz).

Data from the multi-unit recordings were analyzed offline using custom software written in MATLAB (Mathworks Inc., Natick, Massachusetts, USA). We preprocessed the raw voltage recordings as in previous work (Vincent et al., 2013; Vincent et al., 2012). First, each electrodes activity was down-sampled from 5 kHz (acquisition rate) to 1 kHz. Raw voltages were then stored in a matrix $\mathbf{X}$ of size N by T, where $N = 59$ is the number of

channels analyzed and $T = 1{,}114{,}000\,ms$ is the number of milliseconds in a single recording session (20 minutes). We applied a 2nd order high-pass Butterworth filter with a cut-off frequency of 200 Hz to retain high frequency deflections (Figure 2-1B, black trace). Consistent with related work (Quian Quiroga & Panzeri, 2009; Quiroga, Nadasdy, & Ben-Shaul, 2004), multi-unit activity (MUA) was detected by setting a threshold $\theta_i = 5 * \sigma_i$ (Figure 2-1B, red line) for a given channel i, where

$$\sigma_i = mdn\left[\frac{\mathbf{X}_{i,1..T}}{0.6745}\right], \tag{2.1}$$

where the matrix $\mathbf{X}$ contains bandpass filtered signals, each row is a channel indexed by i, and each column is a time-step (measured in ms). For each millisecond time-step, activation was detected when the i^{th} channel's voltage was lower than or equal to threshold θ_i. This process was repeated for all channels generating a matrix $\mathbf{Y}$ of the same size as $\mathbf{X}$, with '1's and '0's denoting activation or silence, respectively (Figure 2-1C). Channels with fewer than two threshold crossings or firing rates greater than the mean plus five standard deviations were removed.

Burst predictors

We sought to characterize the ability of different neurons to predict the upcoming occurrence of network bursts (Figure 2-1D). To detect bursting activity, we first down sampled the recorded data from 1 kHz recordings by a factor of 10. This operation resulted in a matrix $\mathbf{Z}$ of size N by T/B, where $B = 10\,ms$, and '1's and '0's that denote the presence or absence, respectively, of at least one spike in each bin of 10 ms. To detect network bursts, we summed network firing activity across channels, and employed a 100-fold bootstrapping method where the original data was compared to randomly

permuted data, where spike times were randomly shifted in time between 1 and T time-steps. For each random matrix, the firing rate of each channel was maintained; the firing times, however, were randomly permuted across $t = 1,2,\ldots T$ time-steps. This random matrix was summed across channels, producing a vector that was compared to the experimental data. Times when the experimental data exceeded 95% of the random vectors were marked as bursts (Figure 2-1C & 2-1D, blue line marks the 95% threshold). Cultures with less than 5 bursts for an entire baseline recording were removed from further analysis, resulting in the removal of 12 out of 42 cultures. Finally, we computed a burst prediction (BP) score for each channel (Figure 2-1F). This score is computed by counting the number of network bursts preceded by at least one spike by a maximum of 50 ms (Figure 2-1D, red dots), and dividing by the total number of network bursts. The resulting score was divided by the mean rate of each electrode to control for electrodes with high firing rates that did not emit spikes that were specific to network bursts. In the end, we have a score that measures the reliability of an electrode recording a spike in a window of 50 ms prior to a network burst.

Neurotoxic insult

Following baseline recording sessions (17 DIV), cultures at 18 DIV were exposed to a combination of a Na^+/K^+ ATPase inhibitor, 5 µM ouabain (Sigma-Aldrich, St. Louis, Missouri, USA), with a glutamate uptake inhibitor, 40 µM DL-threo-β-benzyloxyaspartic acid (TBOA; Tocris Bioscience, U.S.A.), for 15 min (‘low’ kill) or 20 min (‘high’ kill) or to the DMSO vehicle for 20 min (control). A combined concentration of ouabain and TBOA was added via pipette directly to the extracellular fluid of cultured cells in the insult condition. Ouabain inhibits the sodium-potassium pump (or Na^+/K^+ -ATPase), and

reduces the amount of both sodium pumped out and potassium pumped into the cells (Gao et al., 2002). TBOA is a glutamate transport inhibitor that prevents cellular uptake of glutamate thereby increasing the extracellular concentration of neurotransmitters (Bonde et al., 2003). Previous research has shown that this insult selectively kills neurons by an excitotoxic mechanism, specifically the build-up of extracellular glutamate causing neurotoxic activation of postsynaptic NMDA receptors (Tauskela, Aylsworth, Hewitt, Brunette, & Mealing, 2012). The ouabain-TBOA model of stroke causes over excitation throughout the neuronal network to the point of toxicity, and in some cases cell death.

Assessment of neuronal injury

Neuronal injury was assessed in two stages. First, at DIV 18 prior to insult application, phase contrast images that illuminated healthy cell bodies were acquired and neurons visible within the electrode field were manually counted, to provide a live neuron count (Figure 2-2A, top row). Second, at 19 DIV, 24 hours following neurotoxin exposure, cultures were exposed to the cell death marker propidium iodide (PI; 4.5 µM) that is membrane impermeable, and markedly increases its fluorescence by binding to the DNA of neurons whose plasma membrane has fragmented (Figure 2-2A, bottom row) (Tauskela et al., 2008). PI-positive fluorescent cells were manually counted using images obtained on a fluorescence microscope (Zeiss Axiovert 200 with Lambda DG-4). Images were processed using the ImageJ environment (National Institute of Health). The percentage of dead neurons was determined as the ratio of PI-positive neurons measured at 19 DIV (24 h post-insult) to the number of phase-bright neurons measured at 18 DIV (immediately prior to insult), is presented as the mean ± standard error of measurement for each experimental condition (control [$N = 9$], low kill [$N = 14$], high kill [$N = 7$]).

Statistical comparisons were made by ANOVA, with statistical significance inferred at $p < 0.05$.

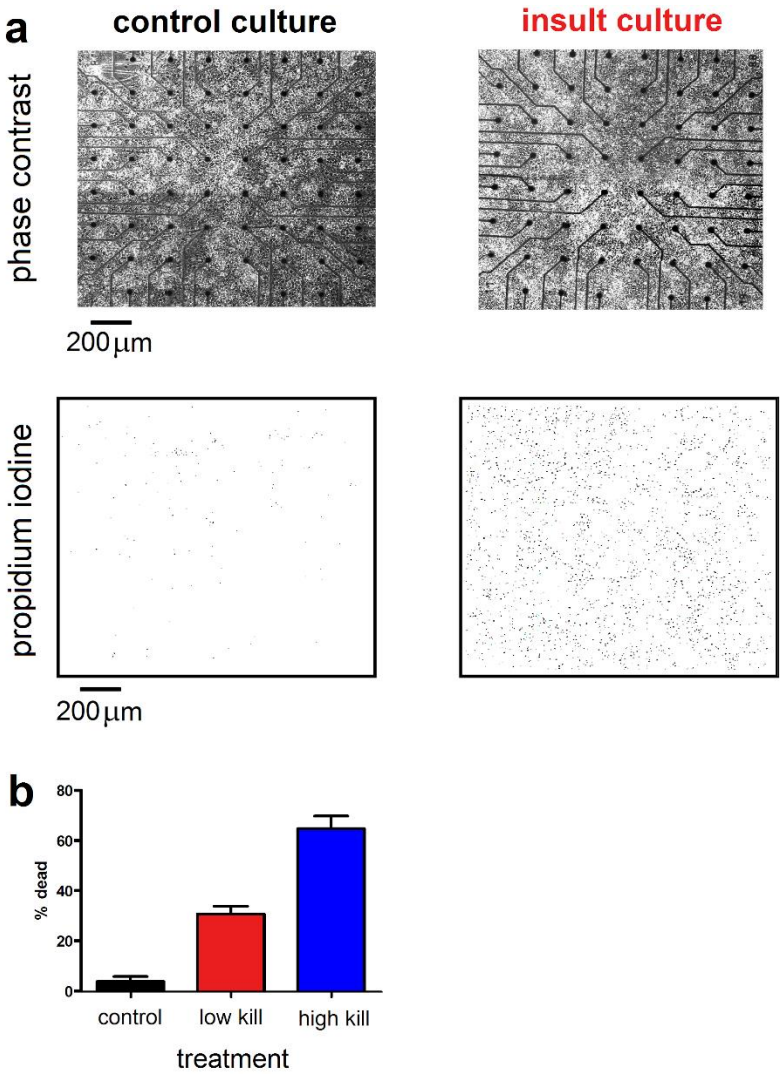

Figure 2-2. Excitotoxic insult causes widespread neuronal death. A. Two example phase contrast (top row) as well as propidium iodide (PI) and fluorescent images (bottom row). One staining image is shown for each condition: black for controls, red for insulted cultures. For the PI images, we have inversed the colour tones, black dots indicate the fluoresced PI and thus cells that are dead or dying. **B**. Mean alive versus

dead ratios (measured in %) for each treatment condition (see Assessment of Neuronal Injury section). Error bars are standard error of measurement (SEM).

Results

Impact of glutamate injury on neuronal cultures

We first evaluated the percentage of death caused by 10 min ('low' kill) and 20 min ('high' kill) exposures to ouabain-TBOA insults applied at 18 DIV, as determined by the ratio between the number of dead (Figure 2-2A, bottom row, PI-positive neurons measured 1-day post-insult at 19 DIV) and live neurons (Figure 2-2A, top row, phase-bright neurons measured immediately prior to the insult at 18 DIV). Representative same-field phase contrast and PI fluorescence images were acquired before and after 20 min ouabain/TBOA (Figure 2-2A). Control cultures (i.e., without insult) displayed low basal neurotoxicity, indicating high quality neuronal cultures. By comparison, the 10 and 20 min ouabain/TBOA insults yielded successive increases in neurotoxicity (Figure 2-2B). Overall, these results show that the insult caused wide-spread cell death within the insulted cultures.

Impact of injury on neuronal activity

Avalanches

Network bursts *in vitro* are often characterized by avalanches in which smaller (or shorter) events occurred more often than larger (or longer) events (Beggs & Plenz, 2003; Langlois et al., 2014). We measured the propagation of activity in both baseline and post-insult cultures by examining three distributions: [1] the duration of avalanches; [2] the total number of cells activated at least once during avalanches; and, [3] the number of MUA events generated during avalanches. Our results show that baseline distributions of all

three measures were approximately linear in log-log coordinates, suggesting a power law distribution (Figure 2-3). Consistent with previous work, we estimated the slope of power laws using a maximum likelihood method (Langlois et al., 2014; Thivierge, 2014). The estimated slope $\hat{\alpha}$ for a bounded discrete power law follows

$$\hat{\alpha} = \arg\max_{\alpha}[-\alpha(\sum_{i=1}^{n} x_i) - n \ln \zeta\,(\alpha, x_{\min}, x_{\max})], \tag{2.2}$$

where x_i is the ith element from the vector of all data $X = [x_1, x_2, \ldots, x_n]$ and n is the size of the data vector. Techniques for estimating lower and upper bounds from raw data are covered elsewhere (Bauke, 2007); here, we set $x_{\min}$ and $x_{\max}$ to the minimum and maximum values observed in the data. The nth element of X corresponds to the duration of time (measured in ms), number of units active during the nth avalanche recorded, and the number of MUA events (Fig. 3, left, middle, and right panels, respectively). The Hurwitz zeta function, $\zeta(\alpha, x_{\min}, x_{\max})$ is given by

$$\zeta(\alpha, x_{\min}, x_{\max}) = \zeta(\alpha, x_{\min}) - \zeta(\alpha, x_{\max}), \tag{2.3}$$

where

$$\zeta(\alpha, x) = \sum_{i=0}^{\infty} \frac{1}{(i+x)^{\alpha}}. \tag{2.4}$$

Consistent with previous work (Langlois et al., 2014; Thivierge, 2014), the estimated slope $\hat{\alpha}$ of all power law distributions was ~1.5 (albeit slightly above, see Figure 2-3, coloured insets). Although the insult yielded marked cell death (Figure 2-2), the propagation of activity was largely maintained, with a slight drop in power law exponents for low kill as well as high kill MEAs (Figure 2-3). The presence of avalanches in the activity of post-

insult cultures implies that spontaneous neuronal dynamics were largely intact despite elevated levels of neuronal death.

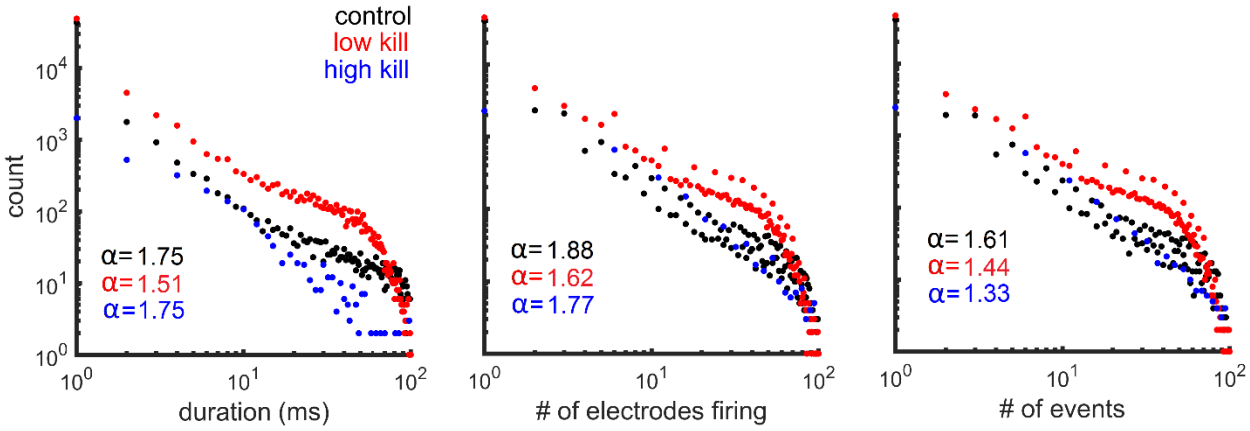

Figure 2-3. Impact of insult on neuronal avalanches. Left panel: histogram in log-log coordinates of the duration of avalanches. Middle panel: number of electrodes active during an avalanche. Right panel: number of MUA events. Maximum likelihood estimates of power law slopes (α) for each distribution are shown on each plot. Control data correspond to DIV 17 and both low kill and high kill data correspond to DIV 21.

We compared these results with those obtained from a branching model that produces avalanches of network bursts (Beggs & Plenz, 2003). First, we computed the critical branching parameter for all baseline recordings. This branching parameter (σ) quantifies the ratio of downstream electrodes that become active as a function of upstream activation. The branching parameter is computed by

$$\sigma = \sum_{d=0}^{n_{max}} d \times p(d), \qquad (2.5)$$

where $p(d)$ is the probability of observing d descendants, and n_{max} is the maximal number of active electrodes. Descendants d are computed by

$$d = round\left(\frac{n_d}{n_a}\right), \tag{2.6}$$

where n_a is the number of ancestors observed in the first time bin, n_d is the number of active electrodes in the second time bin, and $round$ is the rounding operator to the nearest integer (Beggs & Plenz, 2003; Vincent et al., 2012). The probability of observing d descendants was computed by

$$p(d) = \sum_{avalanches}\left(\frac{n_{\sum a|d}}{n_{\sum a}}\right)\left(\frac{n_{max}-1}{n_{max}-n_a}\right), \tag{2.7}$$

where $n_{\sum a|d}$ is the total number of ancestors when n_d descendants were observed, $n_{\sum a}$ is the total number of ancestors observed, and

$$\left(\frac{n_{max}-1}{n_{max}-n_a}\right), \tag{2.8}$$

is a factor that provides a correction for the number of electrodes available in the next time bin because of refractoriness (Beggs & Plenz, 2003). Mean propagation was $\sigma = 0.97$ ($s.d. = 0.04$). This means that on average, every spike was followed by ~1 spike, as consistent with a 3/2 slope of avalanches (Vincent et al., 2012).

Next, we paired the result of the above critical branching analysis with a propagation model. We designed a randomly connected network of $N = 100$ binary (on or off) units that made $C = N$ synapses (i.e., all-to-all connectivity). All synapses had a probability p_i of propagation that was randomly chosen, but constrained so that the sum of all presynaptic connections was equal to

$$\sigma = \sum_{i=1}^{C} p_i, \tag{2.9}$$

where $0 \leq p_i \leq 1$ and $0 \leq \sigma \leq C$. With a propagation constant $\sigma = 1$, the model generated irregular activity with transients of network synchronization like the *in vitro* experimental data (Figure 2-4b, black). The best-fitting slope for avalanche activity (number of spikes) was $\alpha = 1.5$. Increasing the branching parameter ($\sigma = 1.2$) resulted in larger network bursts (Figure 2-4b, red), in turn, the best-fitting slope was slightly lower than 1.5 ($\alpha = 1.3$). Conversely, decreasing the branching parameter ($\sigma = 0.8$) led to smaller network bursts (Figure 2-4b, green) whose distribution was fitted with a steeper slope ($\alpha = 1.7$).

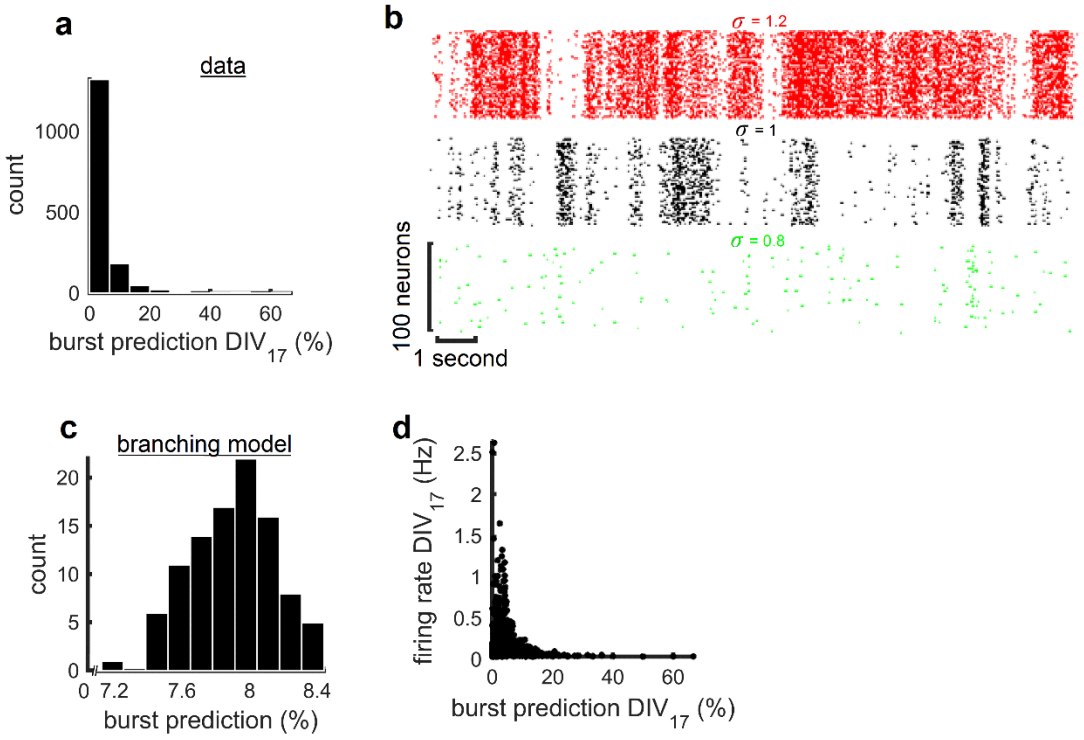

Figure 2-4. Burst predictors and firing activity. a. Histogram of burst predictor scores for all cultures at baseline (DIV 17). **b**. Spike raster depicting network activity generated by the branching model ($N = 100$) with various branching parameters (σ). **c**. Distribution

of burst prediction scores obtained from pooling 100 branching models ($\sigma = 1$). **d**. Scatter plot of firing rates versus burst prediction scores (DIV 17).

Impact of injury on burst predictors

Burst prediction

To further characterize the impact of ischemic injury on spontaneous neuronal activity, we identified channels that reliably predicted network bursts (i.e., Figure 2-1D). We computed a burst prediction (BP) score for each channel at baseline (DIV 17) that measures the propensity of a given channel to be active immediately prior to a network burst (within a pre-burst window of 50 ms), normalized by the mean firing rate (see Methods). The distribution of BP scores was positively skewed, with a heavy tail towards higher values, suggesting that a small subset of channels was responsible for a sizeable proportion of burst prediction (Figure 2-4A).

By comparison, BP scores obtained with the branching model were low, normally distributed, and had a narrow range between 7 and 8.5% (Figure 2-4C). In the branching model, the propagation of activity is based on a random process, and therefore serves as a null model against which experimental scores may be compared. The marked difference in BP scores between the model and experiments suggests that strong burst predictors do not emerge naturally from a process of random propagation.

Are the BP scores obtained from cortical recordings related to firing rates? Plotting BP scores against firing rates reveals that the highest burst predictors (i.e., values greater than 20%) had low firing rates (< 0.5 Hz). Conversely, lower burst predictors (values below 20%) had a broad range of firing rates, including both higher (> 0.5 Hz) and lower

(< 0.5 Hz) values. Overall, firing rates were not strongly correlated with BP scores (Pearson correlation, $r^2(413) = -0.0004, p > 0.372$).

Alterations in firing rate

To determine whether burst predictors were susceptible to insult, we examined the difference in firing rate ($\Delta rate$) between baseline (DIV 17) and post-insult recordings (DIV 21) across individual channels (Figure 2-5A). For each channel of a given MEA, we subtracted the mean firing rate at baseline from the mean firing rate post-insult. Values of $\Delta rate$ that fall below zero indicate a post-insult decrease in firing rate. Across each condition, strong burst predictors (BP scores greater than 20%) maintained firing rates that were comparable to baseline. By comparison, weak burst predictors (BP scores lower than 20%) in the control and low kill conditions showed mainly positive (and some slight negative) changes in firing rate. In the high kill condition, on the other hand, weak burst predictors showed mainly negative changes in rate. Baseline firing rates were not strongly related to $\Delta rate$ values ($r^2(413) = 0.02, p > 0.002$). Most importantly, a small subset of strong burst predictors displayed firing rates that were highly resistant to changes following the excitotoxic injury.

Burst predictors post-insult

Was there any neuronal reorganization in terms of the burst predictors because of ouabain/TBOA? A Wilcoxon test for cumulative distributions shows that for each experimental condition the DIV 17 burst prediction distribution was like that of DIV 21: $Z_{control} = -0.41, p_{control} = 0.678$; $Z_{low\ kill} = 1.62, p_{low\ kill} = 0.106$; $Z_{high\ kill} = 1.14, p_{high\ kill} = 0.256$, respectively. In addition, for the low kill condition we found that the correlation between the percentage of dead channels (based on spiking activity: $SP_i =$

$0/SP_i > 0$) and the mean score for burst predicting channels was not significant, $r_{low\ kill}(293) = 0.14, p = 0.628$. Overall, the ratio of channels that detected burst predictors did not change because of the ouabain/TBOA insult, and on average burst prediction did not decrease as a function of neuronal death.

Are the predictor channels the same one after ouabain-TBOA treatment? A pairwise Pearson correlation between BP score for each channel on DIVs 17 and 21 showed that scores were not strongly correlated: $r_{control}(205) = -0.02, p_{control} = 0.824$; $r_{low\ kill}(293) = -0.07, p_{low\ kill} = 0.249$; $r_{high\ kill}(204) = 0.05, p_{high\ kill} = 0.468$. Using the DIV 17 data, we then grouped channels into high (i.e., top five) or low BP scores and measured the mean for each group on DIV 21. We found that channels with the greatest burst prediction on DIV 17 yielded significantly higher scores on DIV 21 compared to those that scored low on DIV 17 (Figure 2-5B). In addition, we found that channels that had BP scores equal to 0 on DIV 17 emerged with scores > 0 after the insult, in a similar fashion to control cultures (Figure 2-5C). Finally, a principal component analysis (PCA) of BP scores across 5 DIVs (15, 16, 17, 21, & 22) for one insulted culture, shows that the first PCA component was associated with mean burst prediction scores across the DIVs (Figure 2-5D). Overall, although there was a degree of neuronal reorganization after the insult, the most important predictors were resilient to the insult.

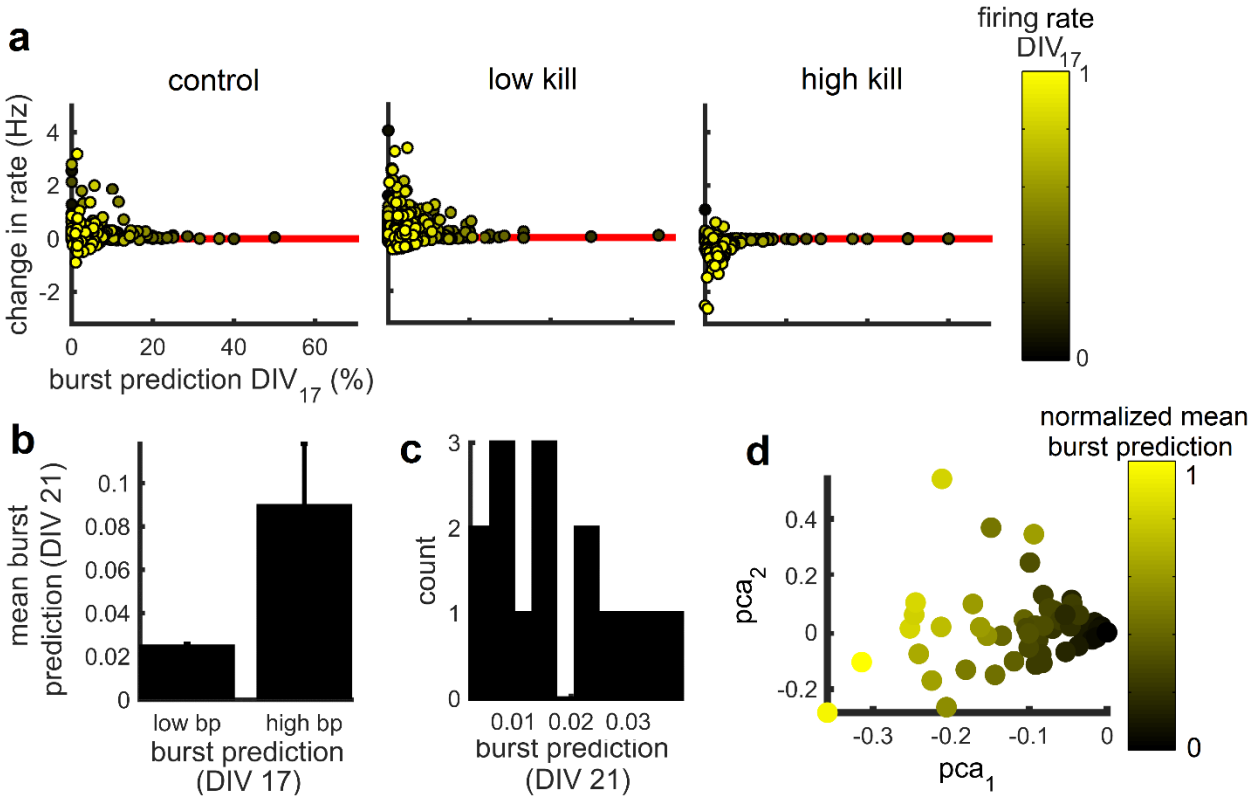

Figure 2-5. Strong burst predictors remain stable after insult. A. Burst prediction scores versus change in firing rate (DIV 17 vs. 21) for one representative culture of control, low kill, and high kill conditions. Firing rates were normalized between [0,1]. Each dot corresponds to a single channel. **B**. Burst prediction scores of low kill cultures pre- (DIV 17) and post- (DIV 21) treatment. High BP, electrodes with the top five scores. Low BP, all other electrodes. Vertical bars, SEM. **C**. Burst prediction scores obtained at DIV 21 for electrodes whose scores at DIV 17 was zero. **D**. Principal components analysis of burst prediction scores across 5 days *in vitro* (DIV 15, 16, 17, 21, and 22) for a representative culture of the low kill condition.

Timing of burst predictor activity

Experimental work shows that the interval between cells firing prior to the onset of a burst is shorter in mature compared to immature cultures 21. Because our recordings took

place over several days *in vitro*, this effect may influence our ability to detect burst predictors. To examine this possibility, we measured the interval between each of the top five burst predicting channels (DIV 17 and 21 of controls) and the onset of bursts (Figure 2-6A). Cultures at DIV 21 had shorter intervals than DIV 17, with a distribution of intervals that was positively skewed (Figure 2-6B). While these distributions were not significantly different (Wilcoxon rank test, $n = 347, p > 0.423$), it is worth considering how changes in the timing of burst predictors would affect our results. These changes would be particularly concerning if, instead of intervals becoming shorter, these intervals became longer over the course of maturation. In this case, strong burst predictors identified at DIV 17 might fall outside of our pre-burst time window (50 ms) at DIV 21. However, if intervals become shorter (as shown here), strong predictors identified at DIV 17 would still fall within the same 50 ms time window at DIV 21 (albeit with an activation that was closer to burst onset). Burst predictors that might emerge at DIV 21 due to shorter time intervals than DIV 17 are not considered in the above analyses, where all burst predictor scores are computed at DIV 17.

Spatial distribution of burst predictors

We examined whether strong burst predictors were spatially clustered on the recording array. First, we identified the channels with the top five BP scores (Figure 2-7A). Then, we measured the Euclidean distance between each pair of these channels (interelectrode distance, 200 μm). The distribution of distances taken over all cultures at baseline (DIV 17) followed a broad distribution with a mean of 738.36 μm ($s.d. = 177.6\,\mu m$) (Figure 2.7B). To determine if strong burst predictors were clustered on the array, we computed a Pearson correlation between mean BP score and physical distance between pairs of

channels. The lack of a strong relationship ($r^2(35) = 0.01, p > 0.958$) shows that strong predictors were not systematically clustered in space, but instead were broadly distributed across the array.

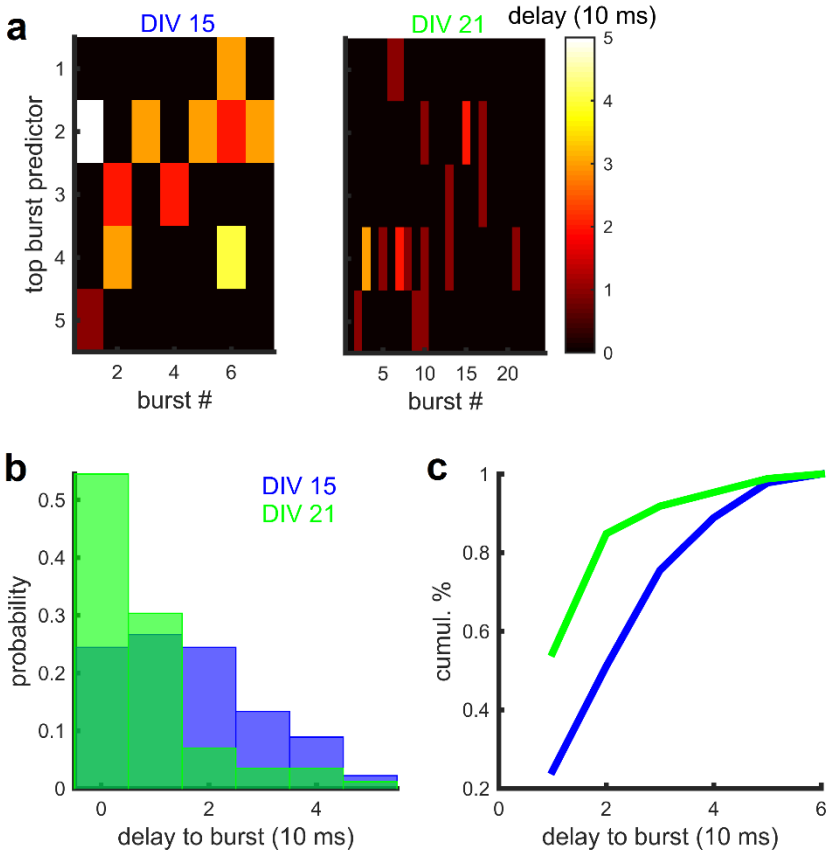

Figure 2-6. Time interval between the activity of burst predictors and burst onsets. A. Heat map of the time interval (in 10 ms resolution) between the five strongest burst predictors and network burst onset in one representative control culture (DIV 15 and 21). **B**. Distributions of time intervals for control cultures (DIV 15 and 21).

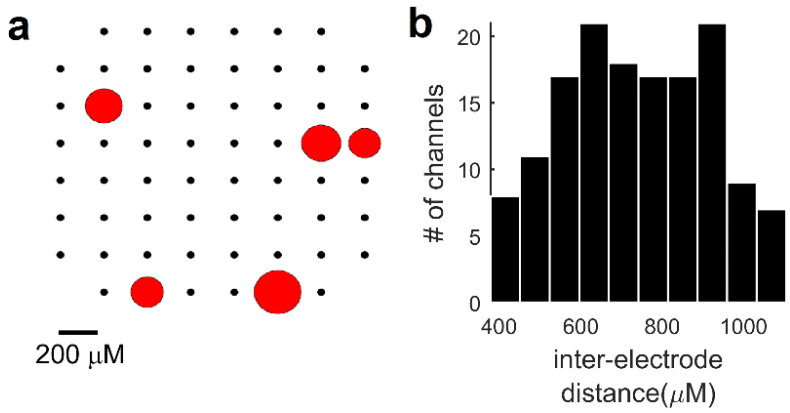

Figure 2-7. Spatial distribution of the top five burst predictors. A. Image of an example MEA's top five burst predictors (red circles) and their spatial location's superimposed over the MEA electrode field (black circles). Red circles indicate the magnitude of BP_i scores. **B**. Histogram of physical distances (measured in μm) between each of the top five burst predictors for each MEA.

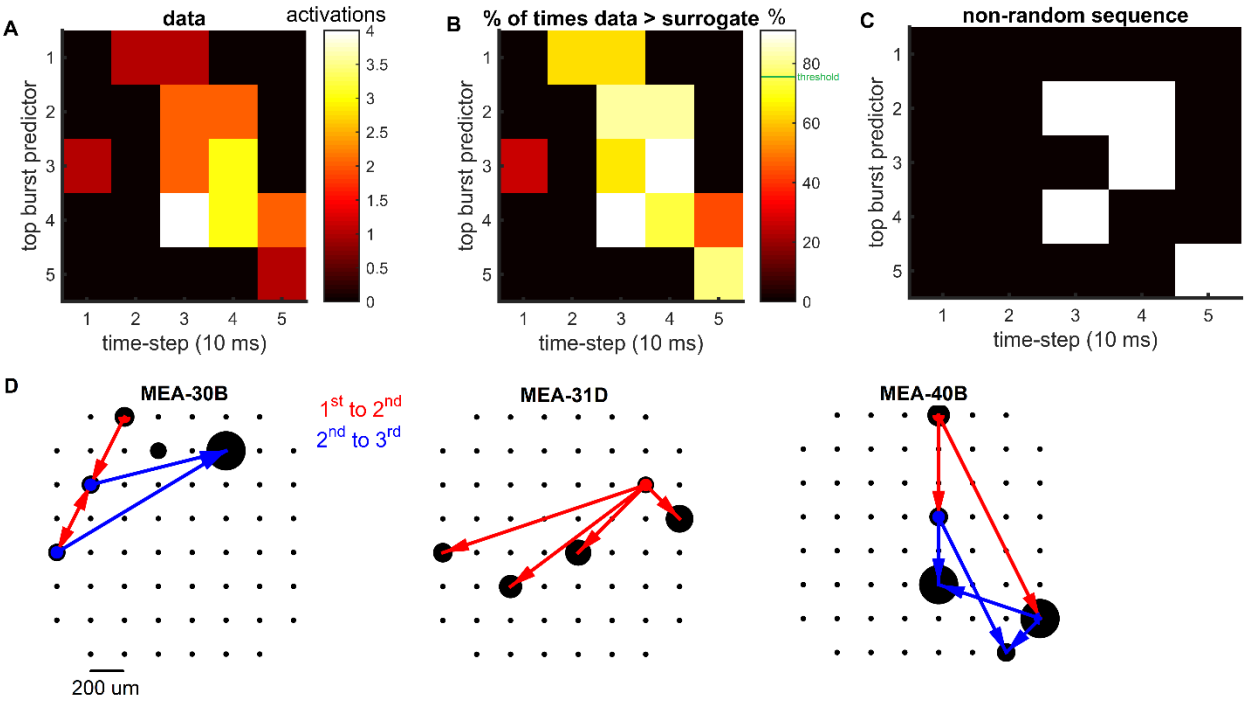

Figure 2-8. Sequential activation of pre-burst electrodes. A. Heat map representing the number of activations in a 50 ms window prior to network bursting activity for one

representative MEA. **B**. Heat map representing the same as A, except, we have generated a ratio based on randomly permuted spike times versus the true data, reported as a percentage. **C**. Non-random binary pattern of pre-burst firing activity, set by threshold (i.e., green line) in B. **D**. Non-random pre-burst firing activity of three MEAs. Red lines represent the first to first, blue represents the second to fire, prior to network bursts.

Correlated pre-burst activity between burst predictors

In a final series of analyses, we examined whether burst predictors fire together before a network burst; or tend to work alone. Here, we examined the firing activity of the top five burst predictors in a 50 ms window before each network burst occurred (Figure 2-8A). We counted the number of spikes recorded 50 ms prior to a network burst, divided this by the total number of network bursts, and compared this to a surrogate data set where spike times were random permuted within the 50 ms window (Figure 2-8B). After comparing true *versus* surrogate data, we found that a small subset of burst predictors was active prior to a burst (Figure 2-8C). Finally, we analyzed other MEAs and found similar patterns of co-activation among burst predictors (Figure 2-8D). These results indicate that burst predictors may interact prior to a burst in a ‘rally-the-troops’ type of fashion, and this activity may be beneficial for the propagation of information throughout the cortex.

Discussion

Although neuronal network bursts are a common feature of brain recordings, it is possible that small subsets of neurons contribute to the neural signature in diverse ways. Here, we examined a few characteristics of burst predictors, and their role in the activity of

cultures undergoing an *in vitro* model of cerebral ischemia. Dissociated primary cortical neurons were plated on multielectrode arrays and spontaneous activity was recorded at 17 days *in vitro* (DIV). As shown in previous work, firing activity was characterized by neuronal avalanches where the statistics of bursts followed a power law, such that large/long duration network events occurred less often compared to small/short duration bursts of activity. We identified burst predictors as electrodes that consistently fired < 50 ms prior to network bursts. The timing of these predictors relative to bursts followed a skewed distribution that differed sharply from a null model based on the branching ratio. Results also show that the activity of burst predictors co-occurs prior bursts of activity, suggesting they may interact. Further, at DIV 21 predictors fired in closer temporal contiguity to network bursts by comparison to DIV 17, possibly due to an increase in the density of synaptic connectivity as shown in other experimental work. Most importantly, a portion of cultures were subjected to an excitotoxic insult (DIV 18). Propidium iodine and fluorescence imaging confirmed cell death in these cultures. While the insult did not alter the distribution of avalanches, it resulted in alterations in overall spike rates. Burst predictors, however, maintained baseline levels of firing activity. The resilience of burst predictors following excitotoxic insult suggests a key role of these units in maintaining network activity following injury, with implications for the selective effects of ischemia in the brain.

The aim of the current study was to examine whether cells that predicted the occurrence of network bursts were more (or less) susceptible than others to an *in vitro* model of ischemia. We computed BP scores that reflected the ability of individual channels to predict upcoming network bursts. Overall, these scores showed a highly-skewed

distribution that could not be reproduced by a branching model with random propagation. Further, BP scores could not be accounted for strictly based on firing rates, changes in propagation delays, or spatial clustering on the arrays. Following an excitotoxic insult, we found that channels with high BP scores (i.e., strong burst predictors) maintained firing rates that were comparable to baseline, while weak burst predictors showed marked alterations (either a decrease or increase) in activity. Importantly, we observed an extent of neuronal reorganization, in that some channels became new predictors post-insult; by comparison, the strongest predictors, remained strong despite extensive cell death in the local network. Together, our results highlight the resilience of strong burst predictors to an excitotoxic insult and open the possibility that these units contribute to the maintenance of network activity following cortical insults.

The idea that neuronal injury may target various aspects of brain networks while leaving other aspects intact has been examined across a broad range of studies. In adult mice, for instance, the rabies virus may be selectively taken up by sensory neurons and then distributed to the dorsal root ganglia (Velandia-Romero, Castellanos, & Martinez-Gutierrez, 2013). In humans, different subgroups of neurons have a distinct immune signature that contributes to susceptibility to infection (Cho et al., 2013). This differential susceptibility may have consequences (albeit largely unknown) on information processing in local networks (Srinivas, Jain, Saurav, & Sikdar, 2007), and our work is a step in this direction.

It remains unclear what might explain the resilience of strong burst predictors to an ischemic insult. In hippocampus, GABAergic interneurons with long or short range dense axonal arborisations have been showed to generate network activity (Bonifazi et al.,

2009). In several brain regions, GABAergic interneurons have greater resistant to cell death than principal neurons (Schwartz-Bloom & Sah, 2001; Wong et al., 2018). Ischemia does, however, play a role in the activation of GABAergic neurons through a rise in Cl^- that reduces the Cl^- influx of these cells (Schwartz-Bloom & Sah, 2001). Further work is required to resolve this question and examine the role of GABA transmission in strong burst predictors both *in vivo* and *in vitro* as well as their link to ischemia.

The identity of burst predictors remains unknown. A calcium imaging study with the hippocampus shows that GABAergic interneurons can generate large network activations (Bonifazi et al., 2009); however, this work was performed with a developing hippocampus with young GABAergic cells that were excitatory (caused by higher intracellular chloride concentration in immature neurons). In addition, an MEA study suggests that excitatory neurons are better predictors of network bursts compared to inhibitory neurons (Tajima, Mita, Bakkum, Takahashi, & Toyoizumi, 2017). Pairing these results with those of our lab, it is not well understood whether burst predicting neurons are excitatory or inhibitory. It may be that some areas of the brain use inhibitory predictors, while others use excitatory, and maybe a mix of the two. Future research may resolve this discrepancy and identify neurons that are central to network bursting.

One question we did not examine here is the possibility of a delayed effect of ischemia on strong burst predictors. Because our recordings were limited to three days' post-insult, it is possible that we did not detect effects of ischemia that occur over a longer duration. Future work that extends our results by recording activity several days *in vitro* following excitotoxic insult would allow us to clarify whether network activity (and particularly burst predictors) remain stable across a longer period. Finally, it is unclear from our results

whether burst predictors merely anticipate network bursts (i.e., association) or whether they drive the burst (i.e., causal) (Bonifazi et al., 2009). Addressing this question is technically challenging, as it requires the identification of strong burst predictors during ongoing recordings. This would constitute a crucial step in understanding the relationship between burst predictors, network activity, and ischemic injury.

Conclusion

Neurons of the brain may use various means of information processing. In Chapter One, we saw how single neurons can utilize oscillations (i.e., periodic signals) to generate reliable responses to incoming stimuli. In Chapter Two, we have highlighted a scenario where neurons are contributing uniquely to the neural signature emitted by a local network that consists of non-periodic activity. However, there is no built-in assumption that the network bursts we have examined are useful to downstream neurons, specifically whether this firing activity can be decoded using a linear model remains unknown. Thus, in Chapter Three we sought to determine if network bursts are a likely candidate for optimal decoding with linear downstream units.

Chapter 3 – Linear Readout of Cortical Activity Suggests a Role for Neural Criticality

Abstract

Spontaneous neuronal activity both *in vivo* and *in vitro* is often characterized by network bursts, whereby a substantial proportion of neurons are active in close temporal contiguity. How the network bursting activity observed in experiments contributes to linear decoding remains an unresolved question in neuroscience. Here, we recorded from dissociated cortical neurons using multi-electrode arrays (MEAs) and computed a 'spike-triggered average' of population activity for each electrode ($N = 59$), by measuring the probability of co-occurrences between the spiking activity on the electrode of interest and that recorded on each of the other electrodes, termed the 'population coupling'. Results show that despite fluctuations in spontaneous activity over time, population activity over all electrodes can be described by a low-dimensional attractor with $N - 1$ parameters, substantially fewer than the number of parameters required for pairwise correlations (N^2). To test whether activity across different networks could be accurately discriminated, population couplings from the first half of recordings (10 minutes) were fed into a linear model trained with a Fisher criterion, namely linear discriminant analysis (LDA). Then, the model was tested by presenting it spiking activity from six networks sequentially. We found that this model was useful in successfully discriminating amongst different networks

with less than 3% error rate. Further, the linear model was robust to adjustments in the number of electrodes included for input to the LDA, suggesting that fewer than $N-1$ parameters can be useful to accurately discriminate between networks. Using simulations of neural activity in a branching model, we show that network activity near a critical regime (but not necessarily at the exact critical point) is optimally discriminated by a linear readout. Overall these results highlight a role for spontaneous bursting activity in neuronal function.

Introduction

Networks of the brain emit highly variable patterns of activity that may be better understood from the perspective of neuronal populations as opposed to individual neurons. One way to characterize a population is by examining the spontaneous bursting activity of neuronal cultures. Networks bursts are momentary increases in the firing rate of units in the network, flanked by silent periods, and are observed *in vivo* and *in vitro*. Bursting activity is highly variable across time; yet, may confer functional benefits in the form of encoding information about stimuli. A burning question in neuroscience is whether such variable patterns of activity are amenable to downstream decoding function, and specifically whether critical dynamics play a role.

Self-organized criticality (SOC) is a theory of dynamical systems that has garnered the attention of researchers worldwide for many decades. For SOC to work, a system must have a ‘critical point’ that rests between two extremes: order and disorder. In the brain, the analog for order and disorder are synchronous and asynchronous activity, respectively. Both states are commonly observed across many regions and preparations both *in vivo* and *in vitro*. Through a series of phase transitions (i.e., synchronous to

asynchronous), a system is drawn towards the critical point (or 'attractor'). The system never actually reaches the critical point, instead it fluctuates around the attractor transitioning between different degrees of order and disorder. One appeal of working with self-organizing critical systems is that they eventually tune their own parameters to reach criticality.

Independent of self-organization, there are many computational benefits to a critical system. First, when brain networks are in the critical state, both the response to stimuli is optimized, and their information processing capability is maximized (Shew & Plenz, 2013). Second, the distribution of neuronal bursts of activity recorded from cortical networks follows a power law, whereby small events (i.e., a few neurons) are observed more often then large events (i.e., many neurons) (Beggs & Plenz, 2003; Plenz, 2012). Third, there are long-range spatial and temporal correlations in neuronal activity (Palva et al., 2013). Fourth, the fluctuations in neuronal activity scale with network size (Tagliazucchi et al., 2012). Many studies investigating criticality in brain networks examine the encoding capabilities; however, whether the activity of these networks can be decoded by downstream neurons has received little attention.

We begin with a thought experiment whereby we imagine the processes involved in a human sensing and perceiving various stimuli in the environment (i.e., cat, flower, or house – Figure 3-1). In this scenario, visual inputs are synaptically transmitted to networks that must discriminate between the various stimuli. These networks are likely the agents that generate a straightforward readout for downstream neurons to decode. The work of downstream neurons is to utilize incoming signals to differentiate between stimuli. The most straightforward representation of decoding employs linear models, whereby points

on a plot, representing the activity of neurons or networks, are separated by a linear function. Taken together with our knowledge of the variable activity inherent of critical systems, we question whether decoding is possible. We further examined a reduced representation of the activity emitted by cortical networks, namely 'population couplings' (Okun et al., 2015). The thought experiment is a simplification of processes in the brain; however, the question of how/whether the dynamics of cortical activity impact linear classification is amenable to investigation with *in vitro* cortical networks.

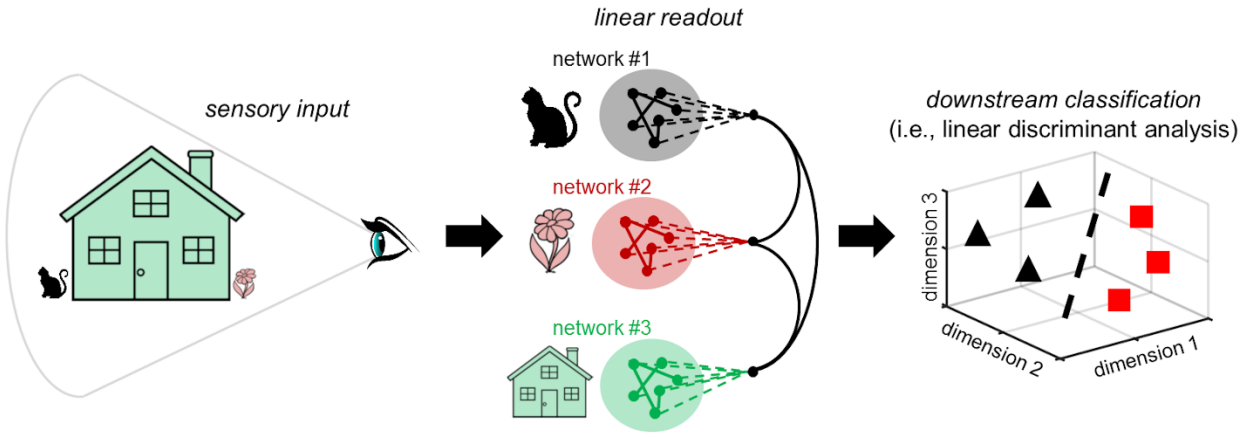

Figure 3-1. Model using cortical activity for linear classification. Sensory input is first received via the visual system. For illustration purposes, the activity of each network represents an abstract concept (i.e., cat, flower, or house). Downstream readout units must discriminate between stimuli in the environment (i.e., cat versus flower).

Here, we examined how the dynamics of *in vitro* cortical network activity effects the downstream readout by linear units. We began by characterizing the firing activity of networks by a 'population coupling'. Here, we computed a spike-triggered average of network activity measured on the entire MEA and found that these reduced representations of networks were amenable to classification with a linear model. We

further characterized cortical cultures as critical systems, measured by propagation constants and the slope of power law distributions. Importantly, we captured the statistics of *in vitro* bursting activity with a model that can be tuned to generate the critical branching phenomena. The model predicted that only systems near or at a critical point are optimal for linear decoding. Our results have implications for the role of bursting activity in neuronal networks and highlight the interplay between critical systems and decoding with linear models.

Hypotheses

i. Population couplings are reliable across both individual electrodes and DIVs; yet different across cultures.

ii. Population couplings can be classified according to the network of origin and outperform ‘shuffled’ spike times.

iii. The activity of *in vitro* networks can be characterized as close to a critical point, and this may be advantageous to classification.

iv. The activity of a critical branching model, like our *in vitro* data, is optimal for classification.

v. Population couplings are useful in predicting the spike times of individual electrodes, and thus recreating the firing activity of the entire culture.

    a. The activity of only surrounding local networks can be useful in predicting the spike times of individual electrodes.

Methods

Cultured neurons on multi electrode arrays

All experiments were approved by the Human Health Therapeutics Animal Care Committee at the National Research Council Canada and carried out in accordance with approved guidelines. Culturing and plating of primary cortical/hippocampal neurons was performed as previously described (Vincent et al., 2013). Briefly, time-pregnant embryonic day 18 Sprague-Dawley rats (Charles River, St. Constant, Quebec, Canada) were anesthetized with halothane and culled by cervical dislocation. Following dissection of the cortical and hippocampal regions of the fetal brains, cells were centrifuged at 1,000 *g* for 3 min at 4ºC and were dispersed by gentle trituration. Cells were plated on MEAs at a high density of 1.8×10^6 cells/ml of medium (consisting of EMEM (Wisent) supplemented to 25 mM glucose), 10% fetal bovine serum (PAA), 10% horse serum (Sigma-Aldrich, St Louis, U.S.A.), and pen/strep (Gibco 1X). The entire surface of each MEA dish (Multi-Channel Systems, Reutlingen, Germany) was pre-treated with a high molecular weight poly-L-lysine (0.025mg/mL diluted in 1×PBS; Sigma-Aldrich) and laminin (0.02mg/ml; Gibco) to promote cell adhesion and minimize cell migration. To better approximate neurotoxic insult conditions previously employed in our laboratory, cell suspensions were plated at 1 ml/MEA so that the entire MEA surface (and not just the inner electrode region) was coated with cells (Tauskela et al., 2008; Vincent et al., 2013). Cultures were maintained in a humidified incubator at 37ºC with a 5% carbon dioxide / 95% air atmosphere. Osmolality was strictly controlled by daily addition of distilled water and by covering MEAs with a transparent, gas-permeable polydimethylsiloxane lid to prevent evaporation (Blau et al., 2009). To minimize glial cell proliferation, a mitotic

inhibitor (15 µg/ml of 5-fluoro-2'-deoxyuridine and 35 µg/ml of uridine) was added at 4 *days in vitro* (DIV). In addition, MEAs were gently rinsed and media was filtered starting at 7 DIV and repeated every 4 days. A 50% media (containing 10% horse serum but not fetal calf serum) change was performed at 7 DIV, with 33% media changed every 3 DIV thereafter.

Recordings of spontaneous neuronal activity

Immediately prior to a recording session, a MEA was removed from the maintenance incubator, capped with a sterile vented lid and placed in the acquisition platform housed within a 37ºC incubator. A 20-minute equilibration period allowed for minimization of movement-stimulated activity. Spontaneous activity was recorded in 20 min sessions at 17-18 and 21-22 DIV using MC Rack software (Multi-Channel Systems) employing the following settings: unit-less amplifier gain (1100.0), input voltage range (+/– 2048 mV) and acquisition rate (5 kHz).

Data from the multi-unit recordings were analyzed offline using custom software written in MATLAB (Mathworks Inc., Natick, Massachusetts, USA). We preprocessed the raw voltage recordings as in previous work (Kuebler et al., 2015; Vincent et al., 2013; Vincent et al., 2012). First, each electrode's activity was down-sampled from 5 kHz (acquisition rate) to 1 kHz. Raw voltages were then stored in a matrix $\boldsymbol{X}$ of size N by T, where $N = 59$ is the number of electrodes analyzed and $T = 1{,}114{,}000$ ms is the number of milliseconds in a single recording session (20 minutes). We applied a 2nd order high-pass Butterworth filter with a cut-off frequency of 200 Hz to retain high frequency deflections. Consistent with related work (Quian Quiroga & Panzeri, 2009; Quiroga et al., 2004), multi-unit activity (MUA) was detected by setting a threshold $\theta_i = 5 \times \sigma_i$, for a given electrode *i* (Equation

2.1). Briefly, for each millisecond time-step, activation was detected when the ith electrode's voltage was lower than or equal to threshold θ_i. This process was repeated for all electrodes generating a matrix $\boldsymbol{Y}$ of the same size as $\boldsymbol{X}$, with '1's and '0's denoting activation or silence, respectively (Figure 3-2A). Electrodes with fewer than two MUA events or firing rates greater than five standard deviations away from the mean were removed from further analyses.

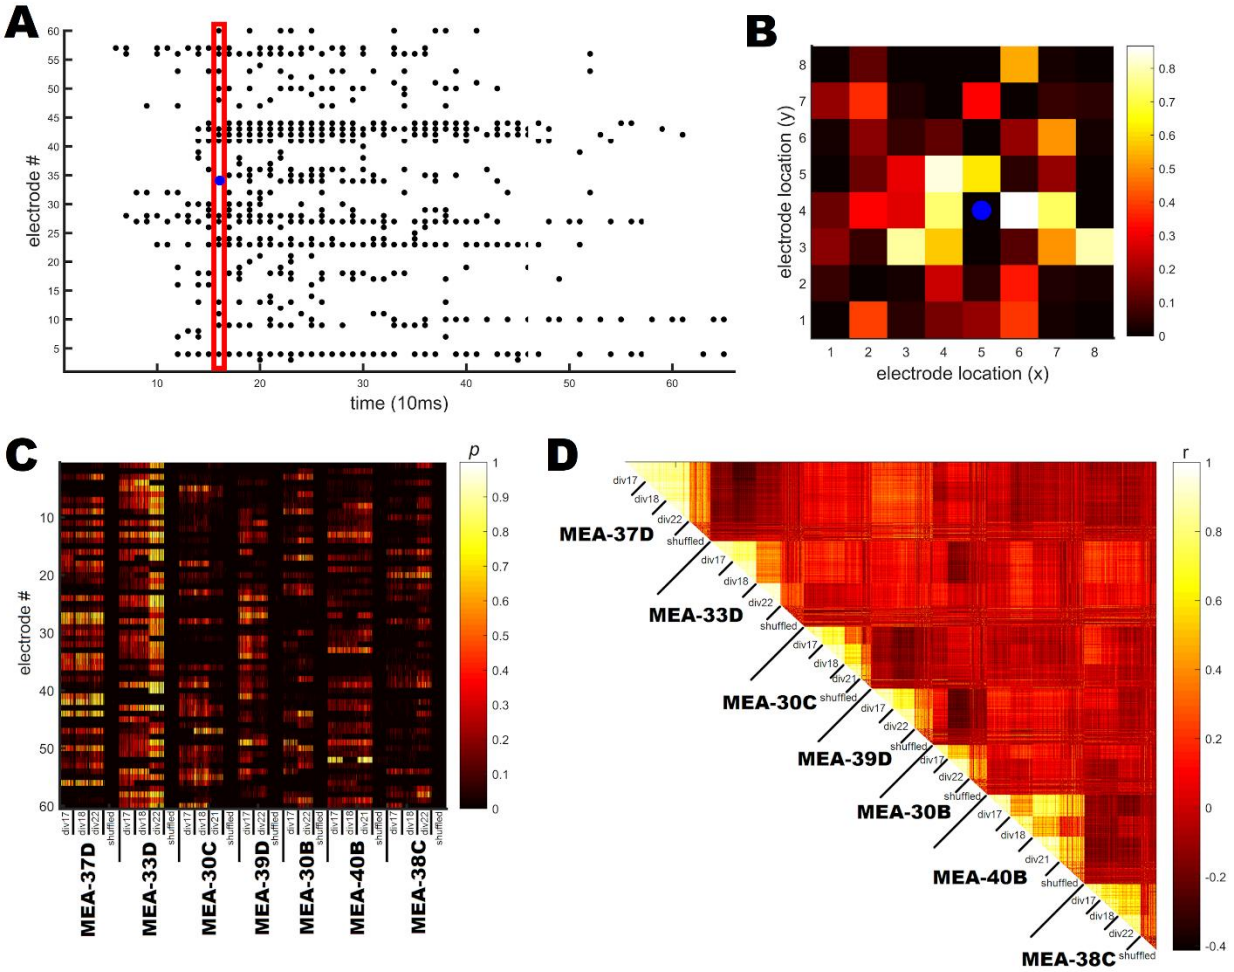

Figure 3-2. Neuronal network bursting activity characterized as population couplings (PC). A. A single network burst spanning several hundred milliseconds recorded from an *in vitro* cortical culture. Black dots represent time-steps where the

electrode's voltage exceeded threshold (Eq.1). Red box denotes firing activity when an electrode of interest (EOI) recorded a spike (i.e., blue dot). **B**. Heat map represents the population coupling for EOI 54 (denoted by blue dot). Dark colours indicate low probability of firing on an electrode; light indicates high likelihood of observing a spike. **C**. PCs for seven cortical cultures across DIV 17, 18, and 19. Shuffled data represents randomly permuted spike times. Dark and light colours denote same as B. **D**. Correlations between PCs of seven cultures. Dark and light colours indicate magnitude and direction of correlation.

Population couplings

Like related work (Tsodyks, Kenet, Grinvald, & Arieli, 1999), we measured population coupling by examining the firing activity of other electrodes (Figure 3-2A, red box) when an electrode of interest (EOI) recorded a spike (Figure 3-2A, blue dot). For each electrode, we counted the number of co-occurrences between spikes on the EOI and those recorded on each of the other electrodes in the array, these counts were then divided by the total number of spikes on the EOI, thus generating a single vector of length $N-1$ representing the likelihood of co-occurring spikes across the entire network (Figure 3-2B). A PC was generated for each electrode that recorded greater than five action potentials in the first 10 minutes of the recording. PCs were computed for each day *in vitro* 17, 18, and 22 generating a grand total of three PCs for each electrode. This model allowed us to examine correlations in the reliability of population couplings across DIVs, and diversity across cultures ($K=7$).

Results

Population couplings of MEAs

In a first series of analyses, we examined the population couplings both within and between different cortical cultures (i.e., MEAs). PCs were largely the same across electrodes, specifically, within a given MEA, a small subset of electrodes were observed emitting a heightened number of coincidental spikes (Figure 3-2C). Importantly, the PC of each culture was unique (Figure 3-2C). Furthermore, reliable PCs were observed over several DIVs including 17, 18 and 22 (Figure 3-2C). Finally, when the recorded spike times were randomly permuted, termed ‘shuffled’, the likelihood of coincidental firing between channels was markedly reduced, thus providing a baseline comparison (Figure 3-2C). Overall, each cortical network had a PC that was reliable across both electrodes and DIVs; yet PCs were varied when comparing across unique cultures. Downstream neurons may be able to take advantage of reliable; yet discriminate inputs, an option that will be discussed in depth in the sections below.

We further examined the population couplings by computing the correlations both within and between neuronal cultures, as well as DIVs. Population couplings were highly correlated within a single DIV, and the magnitude of these correlations was maintained across several DIVs (Figure 3-2D). Importantly, consistent with previous results, correlations between the population couplings of diverse cultures were markedly low. Finally, the correlation between observed population couplings and shuffled was low. Overall, PCs were reliable across both electrodes and DIVs; yet the PC of each culture was unique when comparing across networks. From a computational perspective, reliable

and discriminate variation in the activity of upstream networks may be useful for classification by downstream neurons.

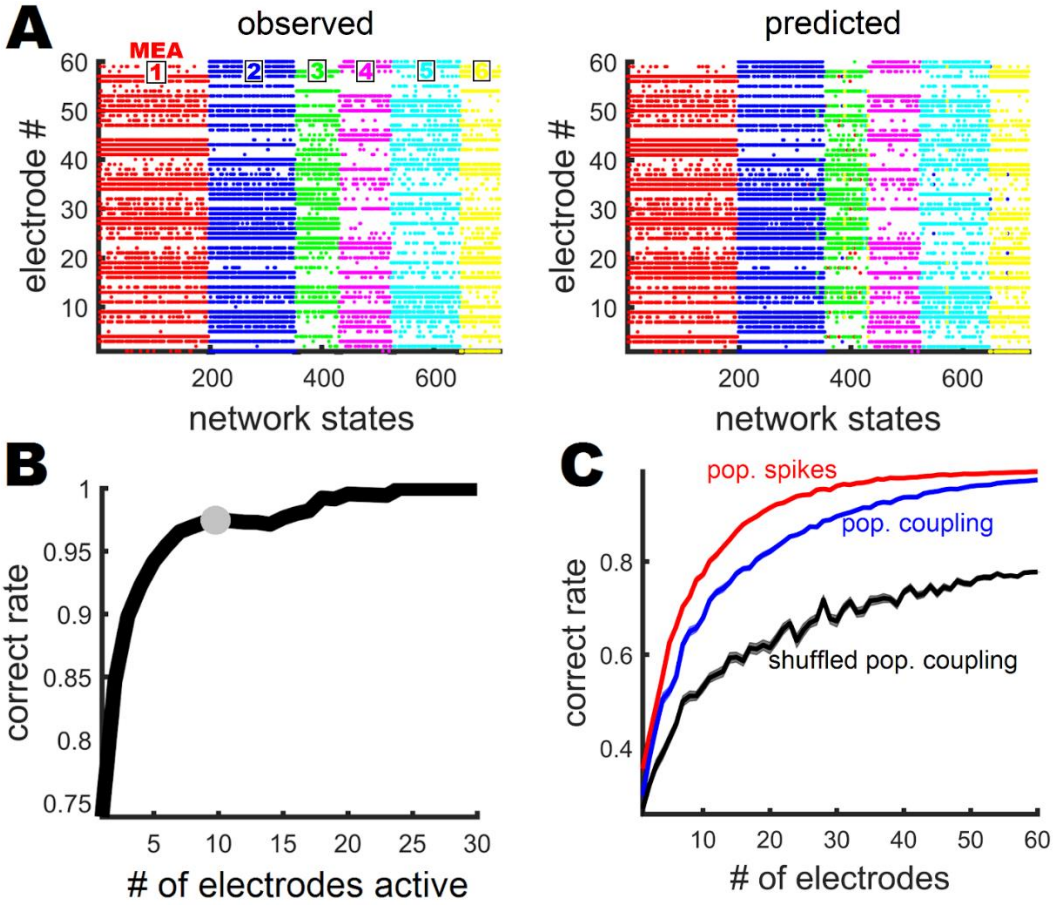

Figure 3-3. Classification of neuronal network activity with a linear model. A. Top panel: ground truth population activity. Bottom panel: ground truth population activity labeled by linear classifier. Spike rasters showing consecutive population bursts. Each network is labeled by color. **B**. Classification accuracy with various numbers of electrodes active on an MEA. Grey dot denotes the data in A. **C**. Performance of the linear readout with population bursts spikes versus population couplings versus shuffled population couplings. Shaded area denotes standard error of mean (SEM).

Decoding population couplings with a linear model

We next sought to determine whether this spontaneous bursting activity may be amenable to classification with downstream units. We chose a straightforward way to perform classification, we separated points on a 2-dimensional surface by finding a linear boundary that most clearly separated the data. Linear models have proven useful for decoding neuronal activity in many different scenarios (Klampfl, David, Yin, Shamma, & Maass, 2012; Meyers, Freedman, Kreiman, Miller, & Poggio, 2008; Nienborg & Cumming, 2010; Shadlen, Britten, Newsome, & Movshon, 1996), including learning models (Buonomano & Maass, 2009). Using a strongly supported method of classification (Berberian, MacPherson, Giraud, Richardson, & Thivierge, 2017), we employed a Fisher linear discriminant analysis to find a linear function that best separates the population couplings. Briefly, the performance of the model was measured using the 10-fold cross-validation method, where the data is binned into 10 bins, 9 were used for training and the 10^{th} for testing. This binning occurs a total of ten times, thus each combination of bins was used to train, and each bin was used for testing (Berberian et al., 2017). Our linear model was trained on the firing activity extracted from the first 10 minutes of the recording and tested on the last 10 minutes of the recording.

Our first step was to characterize the accuracy of classification when various numbers of electrodes were included in the analysis. As a first pass, we detected each of the time-steps where at least n electrodes were active ($n = 1,2, \dots, N$ & $N = 60$), and the network activity at each of these time-steps was concatenated into a matrix (Figure 3-3A, left panel, red dots). This procedure was repeated for each MEA, and this matrix was used as input for the linear model (Figure 3-3A, left panel). For example, the model classified

PCs according to their network of origin with a ~95% accuracy when at least 10 electrodes were active (Figure 3-3A & 3-3B). Time-steps where zero electrodes recorded an action potential were not included. The linear model performed with above 95% accuracy when at least 10 electrodes were included in the training and testing data (Figure 3-3B). Together, these results show that a linear model can be accurately trained to classify neural activity based on a compact representation derived from population couplings.

We tested three different hypotheses for linear decoding by downstream units. This was achieved by training the model with three types of firing activity from the first 10 minutes of recordings: [1] population spikes, defined as the activity of all time-steps where at least 10 electrodes were active; [2] population coupling (see Methods); [3] shuffled population coupling, the same as population coupling based on randomly permuted spike times. Each of these hypotheses were evaluated by randomly selecting a subset of electrodes, of sizes 2 through 60, and computing the correct rate of classification. Fifty subsets of electrodes were generated for each group size (i.e., 2 through 60). For example, we randomly selected 2 electrodes with replacement a total of 50 times, computed the correct rate for each, then generated an average across the 50 subsets.

We found that population spikes slightly outperformed PCs in the accuracy of decoding (Figure 3-3C, red *vs*. blue, respectively). The slight advantage in accuracy conferred by population spikes likely occurs because the activity of all N electrodes was used to train the classification model, compared to $N - 1$ with population couplings. Importantly, both population spikes and PCs achieved greater accuracy compared to the shuffled data (Figure 3-3C). Although performance of the model was slightly lower when considering population bursts compared to PCs; PCs conveyed substantial and useful information

about their network of origin. Thus, highlighting the possibility that a reduced representation of cortical networks may be useful to downstream cells.

Criticality, a branching model, and decoding

To capture the experimental data, we turned to a simple phenomenological model of critical branching that has been employed to describe the power law distribution of neural avalanches (Benayoun, Cowan, van Drongelen, & Wallace, 2010) and $1/f$ scaling dynamics (Ihlen & Vereijken, 2010). In this model, each spike causes a given number of spikes in downstream neurons, and this number is determined by the branching ratio, σ (Figure 3-4A). Values of σ that are lower than 1 generate activity that is sparse; while values greater than 1 produce dense firing patterns where nearly all units are active (Figure 3-4B, blue vs. red, respectively). Importantly, our analysis below shows that with a value of $\sigma = 1$ the firing activity of the branching model is a close match to our *in vitro* data.

As a first step, we computed the branching ratio σ of our experimental data. Like our previous work (Kuebler et al., 2015), the branching parameter is computed by

$$\sigma = \sum_{d=0}^{n_{max}} d \times p(d), \qquad (3.1)$$

where $p(d)$ is the probability of observing d descendants, and n_{max} is the maximal number of active electrodes. Since σ was computed using a function presented in Chapter Two, please refer to Equation 2.5 for a full discussion on how the parameters are computed.

In related work (Kuebler et al., 2015), in a set of control cultures we reported an average branching parameter of $\sigma = 0.97$ $(s.d. = 0.04)$, and this was closely matched by the

current dataset (Figure 3-4C). Importantly, a value of $\sigma = 1$ generates firing activity like the *in vitro* cultures (Figure 3-4B, black). This means that on average, every spike was followed by ~1 spike, as consistent with a 3/2 slope of avalanches (Vincent et al., 2012). This value $\sigma = 1$ is significant because this type of firing activity is associated with a *critical* branching process, whereby local dynamics of networks are coupled to a global variable (Zapperi, Baekgaard Lauritsen, & Stanley, 1995). These results provide one piece of evidence that the bursting activity of our *in vitro* networks may be characterized as a system near a critical point.

To find further evidence of a critical system, we analyzed the statistics of bursting activity. As in related work (Beggs, 2008; Beggs & Plenz, 2003; Kuebler et al., 2015; Vincent et al., 2012), activity was characterized by network bursts that extended over several hundred milliseconds, and the distribution of the number of electrodes active during network bursts followed a power-law, whereby smaller events occurred more often compared to larger events (Figure 3-4D). We computed a maximum likelihood estimate for the slope of this distribution (Langlois et al., 2014), yielding $\alpha = 1.52$. This value is important because it is close to the expected mean field exponent of the critical state (LeBlanc, Angheluta, Dahmen, & Goldenfeld, 2013; Zapperi et al., 1995). Pharmacological alterations in neural activity shifted this slope (PTX: $\alpha = 1.47$; APV/DNQX: $\alpha = 1.63$). Thus, the activity of our *in vitro* neuronal networks may be described as a system close to a critical point.

Distributions of the count of active units during avalanches were obtained for branching ratios $\sigma = 0.5$, $\sigma = 1$, and $\sigma = 1.5$. Each distribution was fitted to a power law using maximum likelihood estimation (Figure 3-4E). Importantly, as in experiments (Haldeman

& Beggs, 2005; Kuebler et al., 2015), best-fitting slopes varied from $\alpha = 1.3$ for subcritical networks (branching ratio $\sigma = 0.5$), to $\alpha = 1.5$ for critical networks (branching ratio $\sigma = 1$) and $\alpha = 1.7$ for supercritical networks (branching ratio $\sigma = 1.5$). Thus, with a branching ratio equal to 1, spiking activity was a close match to the bursting of our *in vitro* cortical cultures.

To better understand the interplay between criticality and decoding, we examined the impact of adjusting the branching parameter on the linear readout. To be clear, by a critical system we refer to neuronal networks with $\sigma = 1$ and $\alpha = 1.5$, and acknowledge that there are other ways to characterize criticality. As with the *in vitro* networks, we extracted population bursts from branching model data using a threshold of 10 concurrently active units within one time-bin. We then computed PCs and trained a linear readout to classify them according to their network of origin. Different networks were simulated with both different initial conditions and randomly assigned connectivity. Classification accuracy of the readout was highest near $\sigma = 1$, and this result was robust to changes in the discharge probability for spontaneous activity (Figure 3-4F) as well as changes to the density of connections in the model (Figure 3-4G). Thus, it seems that critical networks generated activity that is the most amenable to downstream classification.

Considering exclusively the critical condition $\sigma = 1$, we examined the performance of the readout subject to simultaneous input from increasing number of networks. In a model with $N = 100$ units, 7 networks were classified before error rose above 50% (chance). (Figure 3-4H). However, by adding more units to each network, this performance could be markedly improved: with $N = 500$ units, classification error of 10 networks remained

below chance. In all cases, error increased approximately exponentially with the number of subnetworks to classify. In sum, the branching model captured *in vitro* results showing that a regime of $\sigma \sim 1$ is optimal for linear readout of activity.

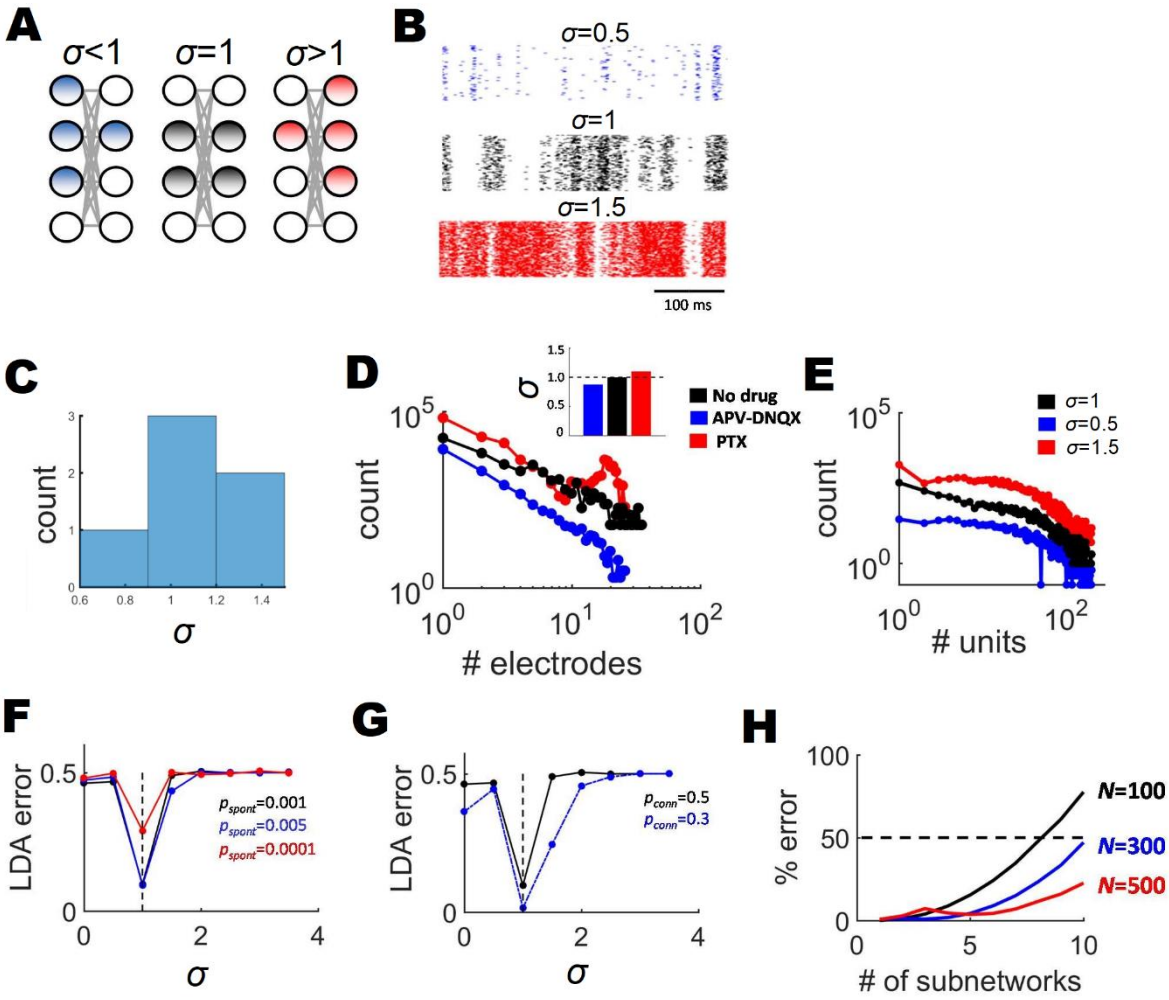

Figure 3-4. Classification with a critical branching model. A. Activity propagation in a branching model with synapses (grey lines) connecting neurons (circles). When the branching ratio (σ) is less than 1, activation fades across layers, while $\sigma = 1$ leads to sustained activation and $\sigma > 1$ leads to amplification. **B**. Spiking activity generated by the branching model with different branching ratios. **C**. Branching parameter of *in vitro* experimental data. **D**. Number of electrodes active per avalanche for control (no drug),

APV-DNQX, and PTX conditions. **E**. Branching model data of the number of units active during avalanches follows a power law when $\sigma = 1$ and deviates from a power law otherwise. **F**. Optimal linear readout of activity with various probabilities of spontaneous activity (p_{spont}). Dashed vertical line indicates $\sigma = 1$. **G**. Linear readout with various network densities obtained by altering the probability of pairwise connections in the branching model from sparser ($N/3$) to denser ($N/2$). **H**. Classification error increases with the number of subnetworks entered in the readout. However, classification performance improves by increasing the number of units (N) in the branching model.

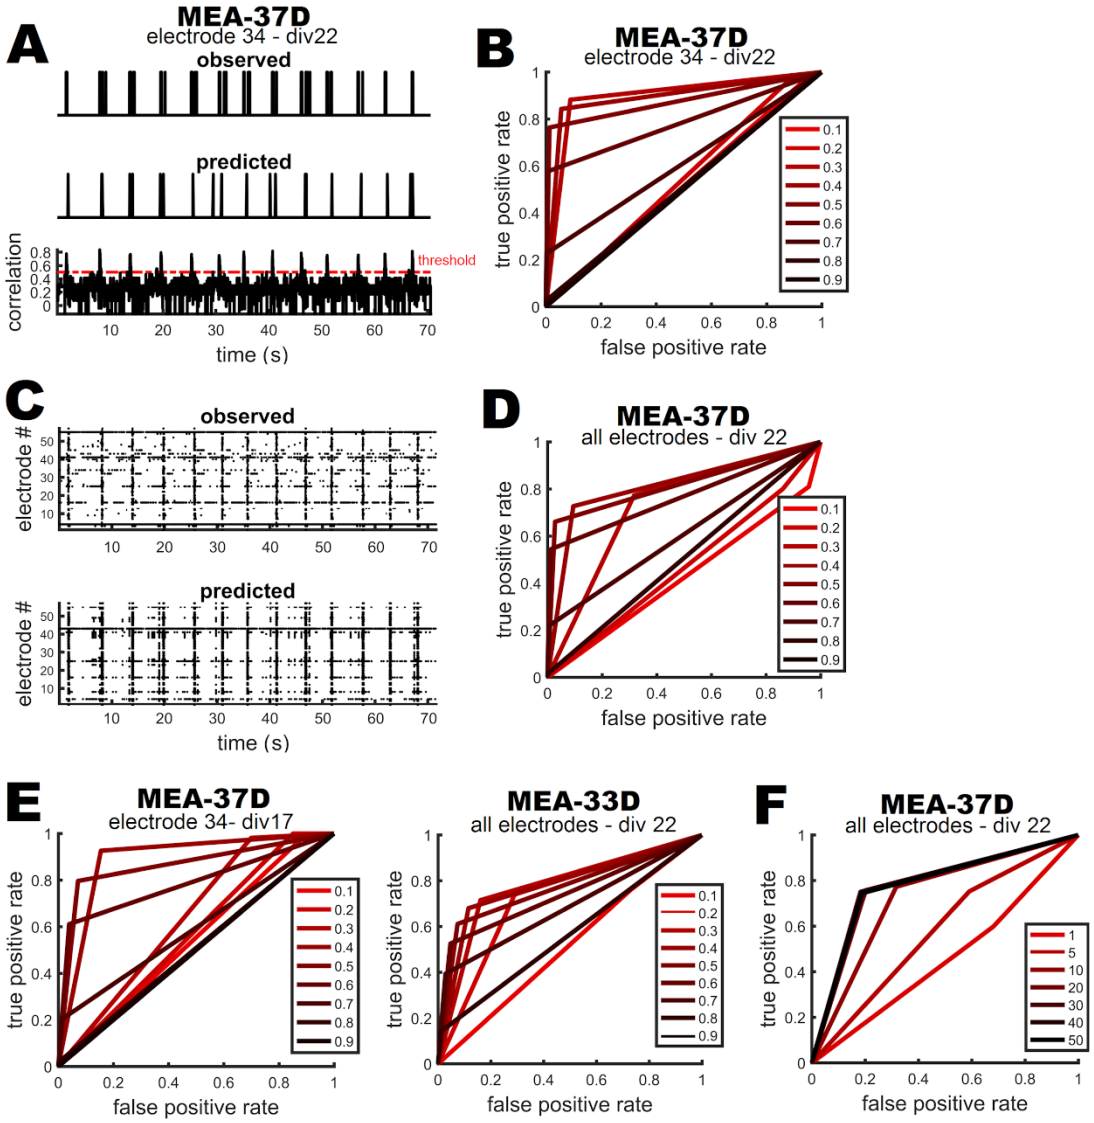

Figure 3-5. Predicting individual electrode and network activity with PCs. A. Observed spikes compared to predicted spike times based on a correlation between the PC of the EOI and network activity. Red dashed line denotes a threshold set on the correlation between a PC and network activity. **B**. Receiver operating characteristics (ROC) curves, representing the accuracy of prediction for the firing activity of a single electrode. Dark colours denote high correlation between PC and network activity; light colours represent low correlations. **C**. Observed versus predicted network activity based on correlations above. **D**. ROC curves representing the prediction accuracy of network

activity. **E**. ROC curves for the same electrode as B; but a different DIV, as well as another neuronal culture. **F**. ROC curves for the same culture as D; but using different numbers of electrodes for prediction.

Predicting spike times with PCs

In our final set of analyses, we considered whether PCs were useful in predicting the spike times of individual electrodes, and entire neuronal networks. To generate a prediction for single electrodes, we correlated the PC of an electrode with the activity in each time-step of the final 10 minutes of the recording, generating a correlation function across time (Figure 3-5A, bottom trace). A predicted spike was counted when the magnitude of the correlation surpassed a threshold value (Figure 3-5A, red dashed line, set to $r = 0.5$). To examine the accuracy of prediction for various thresholds, we employed a receiver operating characteristic (ROC) curve analysis (i.e., Fig. 3-5B, assorted colours). ROC curves are useful for examining the accuracy of prediction by plotting the trade-off between false-positives (x-axis) *versus* true-positives (y-axis). Ideally a classifier will have many true-positives, and just a few false-positives, thus the closer to the upper left area of the plot the better for classification. For single electrode prediction of spike times, accuracy was optimal with a modest correlation of ~0.5 (Figure 3-5B & 3-5E left panel). Thus, PCs were useful in accurately predicting the firing activity of single electrodes; however, this accuracy was largely dependent on the correlation threshold.

We next used an aggregate of single electrode predicted spike times to generate a rather straightforward model of network activity based on PCs. Importantly, the PC of each electrode was slightly different, thus generating diverse spike patterns across the network (Figure 3-5C). Overall, the aggregate of the predicted firing patterns closely matched the

observed network pattern, and the average accuracy of prediction was optimal with a correlation value of ~0.5 (Figure 3-5D & 3-5E right panel). PCs were useful in predicting the spike times of individual channels, and this was robust to adjustments in the number of electrodes, as well as reductions in the size and location of networks used for prediction (Figure 3-5H & 3-6, respectively).

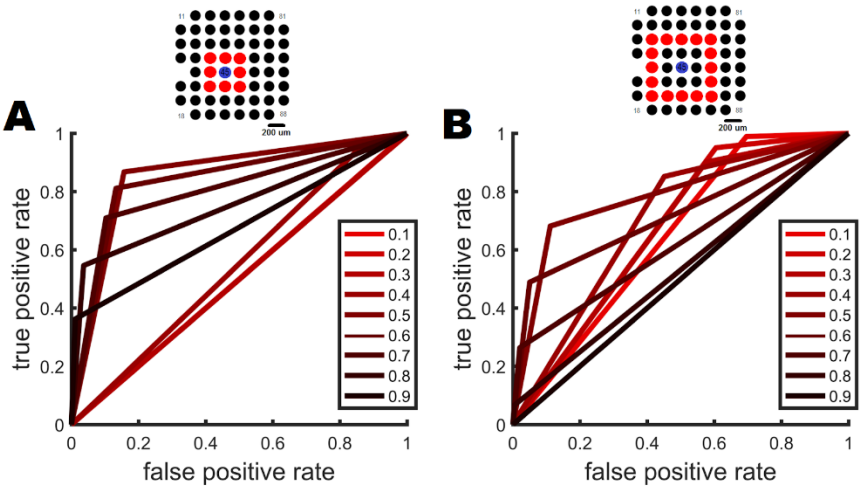

Figure 3-6. Predicting individual electrode spikes with local network activity. A. Accuracy of prediction for a single electrode using only the immediately surrounding electrodes (i.e., red electrodes inset). **B**. Same as A, except using electrodes at a further distance.

Discussion

Although neuronal activity is variable, it is possible that some forms of variability contribute to the ability of downstream neurons to decode incoming information about the environment. Here, we examined the spontaneous activity of *in vitro* neuronal networks, specifically bursts of firing activity, to determine whether their dynamics contribute to decoding by downstream units. Dissociated primary cortical neurons were plated on multielectrode arrays and spontaneous activity was recorded at 17, 18, and 21 days *in*

vitro (DIV). We characterized the firing activity as population couplings (PCs), where we measured coincidental spiking on other electrodes when an electrode of interest recorded a spike. PCs were reliable across both electrodes and DIVs; yet different when compared between MEAs. Correlations confirmed the similarity within and differences between PCs. Both PCs and population spikes proved useful for decoding with a straightforward linear model. We found that the firing activity of our *in vitro* cultures was characteristic of a critical system because of: [1] branching parameters $\sigma \sim 1$, such that on average 1 upstream spike generated 1 downstream spike; and, [2] neuronal avalanches where the statistics of bursts followed a power law, with an exponent $\alpha \sim 1.5$, such that large/long duration network events occurred less often compared to small/short duration bursts of activity. We further found that a critical branching model captured the activity of *in vitro* cultures and allowed us to test the impact of critical branching on linear decoding. We found that decoding was optimal if networks were at or near a critical state, and this was robust to many adjustments in the model. The usefulness of spontaneous neuronal network activity towards decoding has direct implications for criticality and computation in the brain.

The target of this work was to examine whether spontaneous network bursts could be used for decoding with linear downstream models. As a first step, we showed that correlations were high within a culture and across DIVs (i.e., reliable); yet, markedly lower when compared across networks (i.e., discriminant). We hypothesized that characteristics such as these may enable downstream units to decode the activity of several networks. Employing a strongly supported linear model (Berberian et al., 2017), we found that classification was accurate with both a full network representation (i.e., population spikes) and a reduced snapshot (i.e., population couplings). Thus, the bursting

activity of cortical networks was accurately classified by a linear model, implying that network bursts may be useful for downstream neurons.

We further found that the cortical networks could be characterized as critical systems. We determined whether each network was a system near a critical point according to two measures: [1] the propagation constant, or branching parameter σ; and, [2] the slope of distributions representing the bursting activity of cultures, α. We found that the activity of cortical cultures was comparable to that of critical systems with $\sigma \sim 1$ and $\alpha \sim 1.5$. This distinction is important because critical systems are said to be optimal for information processing. Importantly, our work supports this hypothesis, and further highlights that reduced representations of critical systems may be optimal for classification by downstream units. Why critical systems, known to have such variable activity, are optimal for information processing remains an open question.

Our population couplings are a function of correlated firing among neurons in a local network, and this implies an anatomical and/or functional 'connection' among neurons in a culture. Experimental work shows that different network architectures can give rise to similar spiking activity (Prinz, Bucher, & Marder, 2004). These results seem incompatible. However, the work of Prinz et al. (2004) was performed with central pattern generator neurons, while we employed cortical neurons. Pattern generator neurons have strong intrinsic oscillations that generate reliable firing patterns. On the other hand, cortical neurons exhibit greater variability with respect to spike timing responses to sensory stimuli both *in vivo* and *in vitro*. Moreover, while global spiking patterns may be similar, the unique neuronal signature of different neuronal networks may differ.

Our work shows that population couplings are useful to both linear classification and the prediction of spike times for individual electrodes. Whether reduced representations, like network bursts and population couplings, are useful for neural processes such as spike timing dependent plasticity remains unknown. Future work may bridge the gap between network dynamics and synaptic plasticity.

Our work is based on the firing activity of neurons within *in vitro* cortical cultures, and researchers have noted significant differences between *in vitro* and *in vivo* recordings, specifically with the bursting activity of networks. Whether the utility of bursting activity can be generalized to *in vivo* networks remains an open question. Future research may determine whether *in vivo* networks use similar strategies compared to *in vitro* preparations.

Coming back to the original thought experiment (Figure 3-1), we have shown evidence of the similarities and dissimilarities in the activity of cortical networks, and these characteristics may provide the substrate for downstream neurons to classify stimuli. Our work further suggests that the activity of networks can be characterized as critical, which is optimal for linear classification. Overall our work supports the view that networks emits reliable and discriminant activity, which can be characterized as system at or near a critical point, and this may give rise to optimal information processing in the brain.

Conclusion

Neurons of the brain use various strategies to communicate in a noisy biological environment. In Chapters One and Two, some of the mechanisms that single neurons may employ to perform computations were explored, i.e., phase of oscillations and hyperpolarizing resets. Here, we considered the fundamental question of whether the

spontaneous neural activity of large networks, that is characteristic of a system near a critical point, carries important functional consequences for information processing in the brain. Taken together, results point to a useful role of critical brain dynamics in neural decoding and show that a compact, low-dimensional representation of neural activity based on population couplings may serve as a basis for optimal linear readout, provided that this readout is sensitive to the correlational structure of ongoing neural dynamics. Overall, this thesis work highlights three potential candidates for computation in the brain based on both experimental and computational methodologies.

General Discussion

Since the 1920s we have made recordings of neurons of the brain and shown that both spontaneous and stimulus evoked spiking activity can be variable across time. Our aim was to advance our understanding of the strategies that neurons or networks of the brain might use to perform computations. To this end, we examined a few plausible possibilities for harnessing neuronal variability.

Neuronal oscillations and stimulus encoding

Oscillations, in the membrane potential of a single neuron, that results in a periodic ramping up of neural firing activity, are widely observed across many brain regions and preparations. Oscillatory activity in the brain has been a marker for differences on task performance with several different paradigms. Given the prevalence of oscillations, many researchers have sought to uncover the mechanisms behind their functional benefits (Buzsaki & Draguhn, 2004). One possibility that has been experimentally investigated is the spike-phase coding hypothesis, whereby individual spikes encode information about stimuli relative to an oscillation (Kayser et al., 2009; Schaefer et al., 2006). However, the limitations of encoding information about stimuli in this fashion had not been fully elucidated.

In the Chapter One, we examined the critical limitations of the spike-phase hypothesis. Like the *in vitro* experiments (Schaefer et al., 2006), the phase of oscillations played a significant role, whereby stimulus discrimination was heighted for a delimited ranges of phases. Specifically, during the rising phase discrimination was strong, by comparison, neurons were not able to discriminate between stimuli presented during the falling phase

of oscillations. We considered whether this phase effect was due to an accumulation in variance, and indeed results show a non-linear function, where subsequent spikes occurred at increasingly variable times. Our analytical analysis, suggests that spike-phase coding is extremely sensitive to both: [1] the precise timing of spikes with respect to phase; and, [2] the firing rate of the neuron. Taken together, these limitations pose a significant issue for spike-phase coding and add to our understanding of how neurons may utilize oscillations. Future research may show how neurons can overcome these issues.

Recordings of coupled oscillations, characterized by a fast wave phase locked with a slower oscillation, are common in neuronal recordings (Buzsaki & Draguhn, 2004; Canolty et al., 2006; Csicsvari, Hirase, Czurko, Mamiya, & Buzsaki, 1999; Jensen & Colgin, 2007). Given the amount of experimental evidence for oscillation coupling, we chose to use theta and gamma oscillations as a proof of concept that this idea is not far-fetched. Our work shows that it is possible to reduce the variability in subsequent spikes with coupled oscillations. A thorough investigation may reveal all the combinations of frequencies of oscillation that reduce variability, and possibly show that a hyperpolarizing reset during a set of spikes helps track stimuli, thus increases stimulus discrimination, and the possibility of downstream classification.

Finally, oscillations are commonly recorded in networks of cortical neurons (Buzsaki & Draguhn, 2004; Buzsaki, Horvath, Urioste, Hetke, & Wise, 1992). We further discussed the possibility of encoding information about stimuli with oscillations in networks of spiking neurons. Here we show that stimulus discrimination was heightened for a wide-spread range of frequencies and amplitudes of oscillation, as well as stimulus parameters. In

addition, our results show that stimulus discrimination was limited to a portion of the trough and rising phase. Overall, when networks of spiking neurons were injected with oscillations stimulus discrimination was heightened; however, the phase of oscillations played a critical role. A limitation of this work is that we did not examine the specific impact of synaptic weights and connectivity on stimulus encoding. Future research may show the parameter regime that heightens stimulus discrimination for neural networks.

Encoding of bursting activity

Network bursting activity, transient increases in the firing rate of neurons in a network flanked by silent periods, are a typical feature of *in vitro* recordings (Eckmann et al., 2008; Kuebler et al., 2015; Shaukat & Thivierge, 2016; Vincent et al., 2012; Zbinden, 2010, 2011). A subset of neurons, termed ‘burst predictors’ or ‘first-to-fire’, are known to emit spikes reliably prior to network bursts, and these cells maybe be especially important to network function. However, the characteristics of these neurons have not been fully clarified, and despite their apparent priority in network dynamics, their resilience to neuronal insult remains unknown.

In the Chapter Two, we examined how the activity of single neurons may be harnessed to generate network activity. Like related experiments (Eckmann et al., 2008), we examined spontaneous *in vitro* firing activity to find a set of burst predictors, and indeed found that the distribution of scores was such that many were low and a few were high. Further, we found evidence that burst predictors interact prior to a burst; yet, that they are not necessarily physically clustered on the surface of the MEA. At first glance these results may appear to be contradictory; however, it may be that these neurons are distributed throughout the network and have longer range connections (Bonifazi et al.,

2009). Future research, using higher resolution recordings, may be able to better track the location of predictors, and thus track their interactions.

Experimental work has shown that burst predicting neurons of the hippocampus are GABAergic and have dense axonal arborizations (Bonifazi et al., 2009). GABAergic interneurons are also known to have an elevated resistance to insult (Schwartz-Bloom & Sah, 2001). For the first time, we have characterized the resistance of burst predictors to a simulated ischemic event, and indeed, the firing activity of these neurons was stable days after the insult. Thus, our results are consistent with the hypothesis that burst predictors are GABAergic interneurons. Future research may find direct causal evidence of burst predictors by pairing optogenetics with neural recordings.

Branching models are commonly used to capture the dynamics of avalanche activity in neuronal cultures (Beggs & Plenz, 2003; Haldeman & Beggs, 2005). Our results show that a straightforward branching model did not accurately predict the behaviour of the burst predicting neurons. This occurred because the probability of burst initiation was equal and fixed across neurons. Further, we show that predictors may interact prior to a burst. Thus, there are likely more complex dynamics at play that the branching model did not capture. Future research may show that it is possible for a more complex model, such as the LIF, to reproduce the prediction behaviour observed within our *in vitro* cultures.

Overall, our work highlights key characteristics of burst predicting neurons, most importantly their resistance to neuronal insult, thus advancing our understanding of the neurons that track network bursting activity.

Neuronal critical systems and decoding

In contrast to global oscillations, there is experimental and theoretical evidence that the brain is a critical system (Beggs, 2008; Beggs & Timme, 2012; Tagliazucchi et al., 2012), and there are functional benefits to criticality (Plenz, 2013; Shew & Plenz, 2013). Many studies have focused on the ability of networks to encode information (Beggs, 2008; Haldeman & Beggs, 2005). However, whether the firing activity of critical networks is amenable to downstream decoding remains unknown.

In the Chapter Three, we examined whether the bursting activity of *in vitro* cultures, characterized as critical, was useful for classification with a linear model. Bursting activity was described as population couplings, a measure of coincidental firing among electrodes in the array. Population couplings were reliable across both electrodes and DIVs; yet, different when comparing across MEAs. We found that bursting activity was characteristic of critical systems, amenable to linear decoding, and this was robust to reduced representations of firing activity. Based on a branching model that captured the statistics of firing activity, we predict that neural networks at or near the critical point are optimal for linear classification.

Models of plasticity utilize coincidence spiking between pre- and post-synaptic targets to explain how synapses are modified in cortex. Further, our work shows that correlations in the firing activity of neurons in a network can predict individual spike times. Future work may reconcile the gap between network bursting activity and single neuron plasticity.

Overall, these results increase our understanding of the interplay between neural dynamics and successful computation in the brain.

Do results hold *in vivo*?

In chapter 1, we consider whether oscillations can provide a benefit to stimulus discrimination. Whether these results hold for *in vivo* preparations remains an open question. The results of Schaefer et al. (2006) show oscillations in the membrane potential of single OB neurons; however, the ability to discriminate between stimuli *in vivo* was not measured, and it is unknown if this would be possible. Cortical neurons *in vivo* can exhibit increased variability in spike time responses to stimuli, we would expect this to be the case in the OB. Variability in spike timing would play a role in stimulus discrimination, thus it may be more difficult to discriminate between stimuli *in vivo*. However, oscillations may help neurons attain greater stimulus discrimination via phase coding.

In chapter 2, we examine the ability of single neurons to predict the onset of networks bursts of activity. A longstanding debate in neuroscience is whether network bursts are a by-product of *in vitro* preparations, or functionally important. Network bursting activity is typically not observed *in vivo* other than with anesthetized preparations. However, more recently, bursting activity has been observed *in vivo* without anesthesia and during task performance (Bellay et al., 2015; Yu et al., 2017). Importantly, whether individual neurons predict the onset of *in vivo* bursts remains unknown.

In chapter 3, we measured the classification of population activity from different local cortical circuits, based on network bursting activity. It remains unknown whether population coupling would be useful *in vivo*. Previous experimental work has shown that spontaneous and evoked activity *in vitro* are correlated, suggesting some underlying functional architecture (Tsodyks et al., 1999). Moreover, experimental work shows that

spontaneous and evoked population couplings are correlated and can be evoked *in vivo* by visual and optogenetic stimuli (Okun et al., 2015). Based on these results it seems that population couplings may be useful strategy for downstream networks to employ *in vivo*.

Concluding statement

Neurons and networks of the brain may use various strategies to overcome problems such as variability. Uncovering the methods that neurons use to perform computations in cortex is one of the foremost burning questions in neuroscience. In Chapter One, we showed how the variability in the responses of single neurons and networks can be harnessed by oscillations to discriminate between stimuli. Specifically, we examined the limitations of a spike-phase code and show that there are critical constraints on encoding information about stimuli in spike times relative to an oscillation. One alternative to computation with single neurons is encoding within the activity of neuronal networks. In Chapter Two, we discuss the contribution and characteristics of neurons that seem to encode the onset of bursts in the population activity of *in vitro* cortical networks. Here, we show that a small subset of neurons fire prior to a network burst, suggesting that they are important to network function, and examined their resiliency to simulated neuronal insult. We further found that a straightforward branching model with random propagation did not capture the burst predicting behaviour, and we provide some additional evidence that prediction occurs as a function of several neurons interacting. In the Chapter Three, we investigate the likelihood that bursts of network activity, characterized as population couplings, are useful for decoding with linear models. We show that the bursting activity of *in vitro* networks is indeed amenable to linear classification, even with reduced representations of population activity. We further show that our *in vitro* cortical networks can be characterized as systems at or near a critical point, and that this type of activity is optimal for classification. Overall, we highlight a few of the methods that neurons and networks of the brain may use to communicate.

References

Abbott, L. F. (1994). Decoding neuronal firing and modelling neural networks. *Q Rev Biophys, 27*(3), 291-331.

Abbott, L. F. (1999). Lapicque's introduction of the integrate-and-fire model neuron (1907). *Brain Res Bull, 50*(5-6), 303-304.

Abbott, L. F., & Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput, 11*(1), 91-101.

Adrian, E. D., & Bronk, D. W. (1928). The discharge of impulses in motor nerve fibres: Part I. Impulses in single fibres of the phrenic nerve. *J Physiol, 66*(1), 81-101.

Akam, T. E., & Kullmann, D. M. (2010). Oscillations and filtering networks support flexible routing of information. *Neuron, 67*(2), 308-320. doi:10.1016/j.neuron.2010.06.019

Akam, T. E., & Kullmann, D. M. (2012). Efficient "communication through coherence" requires oscillations structured to minimize interference between signals. *PLoS Comput Biol, 8*(11), e1002760. doi:10.1371/journal.pcbi.1002760

Akam, T. E., & Kullmann, D. M. (2014). Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nat Rev Neurosci, 15*(2), 111-122. doi:10.1038/nrn3668

Bal, T., & McCormick, D. A. (1993). Mechanisms of oscillatory activity in guinea-pig nucleus reticularis thalami in vitro: a mammalian pacemaker. *J Physiol, 468*, 669-691.

Bartos, M., Vida, I., Frotscher, M., Meyer, A., Monyer, H., Geiger, J. R., & Jonas, P. (2002). Fast synaptic inhibition promotes synchronized gamma oscillations in

hippocampal interneuron networks. *Proc Natl Acad Sci U S A, 99*(20), 13222-13227. doi:10.1073/pnas.192233099

Bauke, H. (2007). Parameter estimation for power-law distributions by maximum likelihood methods. *The European Physical Journal B, 58*(2), 167-173.

Beggs, J. M. (2008). The criticality hypothesis: how local cortical networks might optimize information processing. *Philos Trans A Math Phys Eng Sci, 366*(1864), 329-343. doi:10.1098/rsta.2007.2092

Beggs, J. M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *J Neurosci, 23*(35), 11167-11177.

Beggs, J. M., & Timme, N. (2012). Being critical of criticality in the brain. *Front Physiol, 3*, 163. doi:10.3389/fphys.2012.00163

Bellay, T., Klaus, A., Seshadri, S., & Plenz, D. (2015). Irregular spiking of pyramidal neurons organizes as scale-invariant neuronal avalanches in the awake state. *Elife, 4*, e07224. doi:10.7554/eLife.07224

Belluscio, M. A., Mizuseki, K., Schmidt, R., Kempter, R., & Buzsaki, G. (2012). Cross-frequency phase-phase coupling between theta and gamma oscillations in the hippocampus. *J Neurosci, 32*(2), 423-435. doi:10.1523/JNEUROSCI.4122-11.2012

Benayoun, M., Cowan, J. D., van Drongelen, W., & Wallace, E. (2010). Avalanches in a stochastic model of spiking neurons. *PLoS Comput Biol, 6*(7), e1000846. doi:10.1371/journal.pcbi.1000846

Berberian, N., MacPherson, A., Giraud, E., Richardson, L., & Thivierge, J. P. (2017). Neuronal pattern separation of motion-relevant input in LIP activity. *J Neurophysiol, 117*(2), 738-755. doi:10.1152/jn.00145.2016

Black, A. H., Nadel, L., & O'Keefe, J. (1977). Hippocampal function in avoidance learning and punishment. *Psychol Bull, 84*(6), 1107-1129.

Blau, A., Neumann, T., Ziegler, C., & Benfenati, F. (2009). Replica-moulded polydimethylsiloxane culture vessel lids attenuate osmotic drift in long-term cell cultures. *J Biosci, 34*(1), 59-69.

Bonde, C., Sarup, A., Schousboe, A., Gegelashvili, G., Zimmer, J., & Noraberg, J. (2003). Neurotoxic and neuroprotective effects of the glutamate transporter inhibitor DL-threo-beta-benzyloxyaspartate (DL-TBOA) during physiological and ischemia-like conditions. *Neurochem Int, 43*(4-5), 371-380.

Bonifazi, P., Goldin, M., Picardo, M. A., Jorquera, I., Cattani, A., Bianconi, G., . . . Cossart, R. (2009). GABAergic hub neurons orchestrate synchrony in developing hippocampal networks. *Science, 326*(5958), 1419-1424. doi:10.1126/science.1175509

Brody, C. D., & Hopfield, J. J. (2003). Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron, 37*(5), 843-852.

Brunel, N., & van Rossum, M. C. (2007). Lapicque's 1907 paper: from frogs to integrate-and-fire. *Biol Cybern, 97*(5-6), 337-339. doi:10.1007/s00422-007-0190-0

Brunel, N., & Wang, X.-J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *J Neurophysiol, 90*(1), 415-430.

Buonomano, D. V., & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat Rev Neurosci, 10*(2), 113-125. doi:10.1038/nrn2558

Butera, R. J., Jr., Rinzel, J., & Smith, J. C. (1999a). Models of respiratory rhythm generation in the pre-Botzinger complex. I. Bursting pacemaker neurons. *J Neurophysiol, 82*(1), 382-397.

Butera, R. J., Jr., Rinzel, J., & Smith, J. C. (1999b). Models of respiratory rhythm generation in the pre-Botzinger complex. II. Populations Of coupled pacemaker neurons. *J Neurophysiol, 82*(1), 398-415.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nat Neurosci, 7*(5), 446-451. doi:10.1038/nn1233

Buzsaki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science, 304*(5679), 1926-1929. doi:10.1126/science.1099745

Buzsaki, G., Geisler, C., Henze, D. A., & Wang, X. J. (2004). Interneuron Diversity series: Circuit complexity and axon wiring economy of cortical interneurons. *Trends Neurosci, 27*(4), 186-193. doi:10.1016/j.tins.2004.02.007

Buzsaki, G., Horvath, Z., Urioste, R., Hetke, J., & Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science, 256*(5059), 1025-1027.

Buzsaki, G., & Wang, X. J. (2012). Mechanisms of gamma oscillations. *Annu Rev Neurosci, 35*, 203-225. doi:10.1146/annurev-neuro-062111-150444

Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., . . . Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science, 313*(5793), 1626-1628. doi:10.1126/science.1128115

Carlson, G. C., Shipley, M. T., & Keller, A. (2000). Long-lasting depolarizations in mitral cells of the rat olfactory bulb. *J Neurosci, 20*(5), 2011-2021.

Chiappalone, M., Bove, M., Vato, A., Tedesco, M., & Martinoia, S. (2006). Dissociated cortical networks show spontaneously correlated activity patterns during in vitro development. *Brain Res, 1093*(1), 41-53. doi:10.1016/j.brainres.2006.03.049

Cho, H., Proll, S. C., Szretter, K. J., Katze, M. G., Gale, M., Jr., & Diamond, M. S. (2013). Differential innate immune response programs in neuronal subtypes determine susceptibility to infection in the brain by positive-stranded RNA viruses. *Nat Med, 19*(4), 458-464. doi:10.1038/nm.3108

Clayton, M. S., Yeung, N., & Cohen Kadosh, R. (2015). The roles of cortical oscillations in sustained attention. *Trends Cogn Sci, 19*(4), 188-195. doi:10.1016/j.tics.2015.02.004

Cohen, M. R., & Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat Neurosci, 14*(7), 811-819. doi:10.1038/nn.2842

Cressman, J. R., Ullah, G., Ziburkus, J., Schiff, S. J., & Barreto, E. (2009). The influence of sodium and potassium dynamics on excitability, seizures, and the stability of persistent states: I. Single neuron dynamics. *J Comput Neurosci, 26*(2), 159-170.

Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., & Buzsaki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. *J Neurosci, 19*(1), 274-287.

Cui, Y., Liu, L. D., McFarland, J. M., Pack, C. C., & Butts, D. A. (2016). Inferring Cortical Variability from Local Field Potentials. *J Neurosci, 36*(14), 4121-4135. doi:10.1523/JNEUROSCI.2502-15.2016

Cury, K. M., & Uchida, N. (2010). Robust odor coding via inhalation-coupled transient activity in the mammalian olfactory bulb. *Neuron, 68*(3), 570-585. doi:10.1016/j.neuron.2010.09.040

Destexhe, A., McCormick, D. A., & Sejnowski, T. J. (1993). A model for 8–10 Hz spindling in interconnected thalamic relay and reticularis neurons. *Biophysical journal, 65*(6), 2473-2477.

Doiron, B., Chacron, M. J., Maier, L., Longtin, A., & Bastian, J. (2003). Inhibitory feedback required for network oscillatory responses to communication but not prey stimuli. *Nature, 421*(6922), 539.

Domnisoru, C., Kinkhabwala, A. A., & Tank, D. W. (2013). Membrane potential dynamics of grid cells. *Nature, 495*(7440), 199-204. doi:10.1038/nature11973

Eckmann, J.-P., Jacobi, S., Marom, S., Moses, E., & Zbinden, C. (2008). Leader neurons in population bursts of 2D living activity. *New Journal of Physics, 10*, 19.

Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci, 2*(10), 704-716. doi:10.1038/35094565

Eytan, D., & Marom, S. (2006). Dynamics and effective topology underlying synchronization in networks of cortical neurons. *J Neurosci, 26*(33), 8465-8476. doi:10.1523/JNEUROSCI.1627-06.2006

Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nat Rev Neurosci, 9*(4), 292-303. doi:10.1038/nrn2258

Ferguson, K. A., Chatzikalymniou, A. P., & Skinner, F. K. (2017). Combining Theory, Model and Experiment to Explain How Intrinsic Theta Rhythms Are Generated in

an in Vitro Whole Hippocampus Preparation without Oscillatory Inputs. *eNeuro*, ENEURO. 0131-0117.2017.

Friedrich, R. W., & Stopfer, M. (2001). Recent dynamics in olfactory population coding. *Curr Opin Neurobiol, 11*(4), 468-474.

Fries, P., Nikolic, D., & Singer, W. (2007). The gamma cycle. *Trends Neurosci, 30*(7), 309-316. doi:10.1016/j.tins.2007.05.005

Fukunaga, I., Herb, J. T., Kollo, M., Boyden, E. S., & Schaefer, A. T. (2014). Independent control of gamma and theta activity by distinct interneuron networks in the olfactory bulb. *Nat Neurosci, 17*(9), 1208-1216. doi:10.1038/nn.3760

Gan, J., Weng, S.-m., Pernía-Andrade, A. J., Csicsvari, J., & Jonas, P. (2017). Phase-Locked Inhibition, but Not Excitation, Underlies Hippocampal Ripple Oscillations in Awake Mice In Vivo. *Neuron, 93*(2), 308-314.

Gao, J., Wymore, R. S., Wang, Y., Gaudette, G. R., Krukenkamp, I. B., Cohen, I. S., & Mathias, R. T. (2002). Isoform-specific stimulation of cardiac Na/K pumps by nanomolar concentrations of glycosides. *J Gen Physiol, 119*(4), 297-312.

Gautam, S. H., Hoang, T. T., McClanahan, K., Grady, S. K., & Shew, W. L. (2015). Maximizing Sensory Dynamic Range by Tuning the Cortical State to Criticality. *PLoS Comput Biol, 11*(12), e1004576. doi:10.1371/journal.pcbi.1004576

Gerstner, W., & Naud, R. (2009). Neuroscience. How good are neuron models? *Science, 326*(5951), 379-380. doi:10.1126/science.1181936

Hafting, T., Fyhn, M., Bonnevie, T., Moser, M. B., & Moser, E. I. (2008). Hippocampus-independent phase precession in entorhinal grid cells. *Nature, 453*(7199), 1248-1252. doi:10.1038/nature06957

Haldeman, C., & Beggs, J. M. (2005). Critical branching captures activity in living neural networks and maximizes the number of metastable States. *Phys Rev Lett, 94*(5), 058101. doi:10.1103/PhysRevLett.94.058101

Hanslmayr, S., & Staudigl, T. (2014). How brain oscillations form memories--a processing based perspective on oscillatory subsequent memory effects. *Neuroimage, 85 Pt 2*, 648-655. doi:10.1016/j.neuroimage.2013.05.121

Harris, K. D. (2005). Neural signatures of cell assembly organization. *Nat Rev Neurosci, 6*(5), 399-407. doi:10.1038/nrn1669

Harris, K. D., Henze, D. A., Hirase, H., Leinekugel, X., Dragoi, G., Czurko, A., & Buzsaki, G. (2002). Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature, 417*(6890), 738-741. doi:10.1038/nature00808

Harris, T. E. (1989). *The Theory of Branching Processes*. New York: Dover Publications.

Hutcheon, B., & Yarom, Y. (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends Neurosci, 23*(5), 216-222.

Ihlen, E. A., & Vereijken, B. (2010). Interaction-dominant dynamics in human cognition: beyond 1/f(alpha) fluctuation. *J Exp Psychol Gen, 139*(3), 436-463. doi:10.1037/a0019098

Jensen, O., & Colgin, L. L. (2007). Cross-frequency coupling between neuronal oscillations. *Trends Cogn Sci, 11*(7), 267-269. doi:10.1016/j.tics.2007.05.003

Jensen, O., & Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. *J Neurophysiol, 83*(5), 2602-2609.

Johnson, M., & Chartier, S. (2017). Spike neural models (part I): The Hodgkin-Huxley model. *The quantitative methods for psychology, 13*(2), 105-119. doi:10.20982/tqmp.13.2.p105

Johnson, M., & Chartier, S. (2018). Spike Neural Models Part II: Abstract Neural Models. *The quantitative methods for psychology, 14*(1), 1-16. doi:10.20982/tqmp.14.1.p001

Jolivet, R., Schurmann, F., Berger, T. K., Naud, R., Gerstner, W., & Roth, A. (2008). The quantitative single-neuron modeling competition. *Biol Cybern, 99*(4-5), 417-426. doi:10.1007/s00422-008-0261-x

Kamondi, A., Acsady, L., Wang, X. J., & Buzsaki, G. (1998). Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials. *Hippocampus, 8*(3), 244-261. doi:10.1002/(SICI)1098-1063(1998)8:3<244::AID-HIPO7>3.0.CO;2-J

Kayser, C., Montemurro, M. A., Logothetis, N. K., & Panzeri, S. (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron, 61*(4), 597-608. doi:10.1016/j.neuron.2009.01.008

Kepecs, A., & Lisman, J. (2003). Information encoding and computation with spikes and bursts. *Network, 14*(1), 103-118.

Kepecs, A., Wang, X. J., & Lisman, J. (2002). Bursting neurons signal input slope. *J Neurosci, 22*(20), 9053-9062.

Klampfl, S., David, S. V., Yin, P., Shamma, S. A., & Maass, W. (2012). A quantitative analysis of information about past and present stimuli encoded by spikes of A1 neurons. *J Neurophysiol, 108*(5), 1366-1380. doi:10.1152/jn.00935.2011

Kuebler, E. S., Bonnema, E. M., J., & Thivierge, J. P. (2013). *Stimulus discrimination in networks of spiking neurons.* Paper presented at the International Joint Conference on Neural Networks, Dallas, Texas.

Kuebler, E. S., Tauskela, J. S., Aylsworth, A., Zhao, X., & Thivierge, J. P. (2015). Burst predicting neurons survive an in vitro glutamate injury model of cerebral ischemia. *Sci Rep, 5*, 17718. doi:10.1038/srep17718

Kuebler, E. S., & Thivierge, J. P. (2014). Spiking variability: Theory, measures and implementation in Matlab. *The quantitative methods for psychology, 10*(2), 131-142.

Langlois, D., Cousineau, D., & Thivierge, J. P. (2014). Maximum likelihood estimators for truncated and censored power-law distributions show how neuronal avalanches may be misevaluated. *Phys Rev E Stat Nonlin Soft Matter Phys, 89*(1-1), 012709.

Lapicque, L. (1907). Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen., 9*, 620-635.

Latham, P. E., & Lengyel, M. (2008). Phase coding: spikes get a boost from local fields. *Curr Biol, 18*(8), R349-351. doi:10.1016/j.cub.2008.02.062

Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci, 3*(11), 884-895. doi:10.1038/nrn964

Laurent, G., & Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science, 265*(5180), 1872-1875. doi:10.1126/science.265.5180.1872

LeBlanc, M., Angheluta, L., Dahmen, K., & Goldenfeld, N. (2013). Universal fluctuations and extreme statistics of avalanches near the depinning transition. *Phys Rev E Stat Nonlin Soft Matter Phys, 87*(2), 022126. doi:10.1103/PhysRevE.87.022126

Lee, H., Simpson, G. V., Logothetis, N. K., & Rainer, G. (2005). Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron, 45*(1), 147-156. doi:10.1016/j.neuron.2004.12.025

Lisman, J. (2005). The theta/gamma discrete phase code occuring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus, 15*(7), 913-922. doi:10.1002/hipo.20121

Litwin-Kumar, A., & Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat Neurosci, 15*(11), 1498-1505. doi:10.1038/nn.3220

London, M., Roth, A., Beeren, L., Hausser, M., & Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature, 466*(7302), 123-127. doi:10.1038/nature09086

Mackay, W. A. (1997). Synchronized neuronal oscillations and their role in motor processes. *Trends Cogn Sci, 1*(5), 176-183. doi:10.1016/S1364-6613(97)01059-0

Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science, 268*(5216), 1503-1506.

Margrie, T. W., & Schaefer, A. T. (2003). Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J Physiol, 546*(Pt 2), 363-374.

Masland, R. H. (2001). Neuronal diversity in the retina. *Curr Opin Neurobiol, 11*(4), 431-436.

Masuda, N., & Doiron, B. (2007). Gamma oscillations of spiking neural populations enhance signal discrimination. *PLoS Comput Biol, 3*(11), e236. doi:10.1371/journal.pcbi.0030236

Matsuoka, K. (1985). Sustained oscillations generated by mutually inhibiting neurons with adaptation. *Biol Cybern, 52*(6), 367-376.

Mazarakis, N. K., Cybulska-Klosowicz, A., Grote, H., Pang, T., Van Dellen, A., Kossut, M., . . . Hannan, A. J. (2005). Deficits in experience-dependent cortical plasticity and sensory-discrimination learning in presymptomatic Huntington's disease mice. *J Neurosci, 25*(12), 3059-3066. doi:10.1523/JNEUROSCI.4320-04.2005

Mehta, M. R., Lee, A. K., & Wilson, M. A. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature, 417*(6890), 741-746. doi:10.1038/nature00807

Meldrum, B. S. (2000). Glutamate as a neurotransmitter in the brain: review of physiology and pathology. *J Nutr, 130*(4S Suppl), 1007S-1015S.

Mensi, S., Naud, R., Pozzorini, C., Avermann, M., Petersen, C. C., & Gerstner, W. (2012). Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms. *J Neurophysiol, 107*(6), 1756-1775. doi:10.1152/jn.00408.2011

Meyers, E. M., Freedman, D. J., Kreiman, G., Miller, E. K., & Poggio, T. (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol, 100*(3), 1407-1419. doi:10.1152/jn.90248.2008

Miura, K., Mainen, Z. F., & Uchida, N. (2012). Odor representations in olfactory cortex: distributed rate coding and decorrelated population activity. *Neuron, 74*(6), 1087-1098. doi:10.1016/j.neuron.2012.04.021

Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., & Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr Biol, 18*(5), 375-380. doi:10.1016/j.cub.2008.02.023

Monto, S., Palva, S., Voipio, J., & Palva, J. M. (2008). Very slow EEG fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans. *J Neurosci, 28*(33), 8268-8272. doi:10.1523/JNEUROSCI.1910-08.2008

Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., & Pouget, A. (2014). Information-limiting correlations. *Nat Neurosci, 17*(10), 1410-1417. doi:10.1038/nn.3807

Naud, R., Bathellier, B., & Gerstner, W. (2014). Spike-timing prediction in cortical neurons with active dendrites. *Front Comput Neurosci, 8*, 90. doi:10.3389/fncom.2014.00090

Naud, R., & Gerstner, W. (2012). Coding and decoding with adapting neurons: a population approach to the peri-stimulus time histogram. *PLoS Comput Biol, 8*(10), e1002711. doi:10.1371/journal.pcbi.1002711

Naud, R., Marcille, N., Clopath, C., & Gerstner, W. (2008). Firing patterns in the adaptive exponential integrate-and-fire model. *Biol Cybern, 99*(4-5), 335-347. doi:10.1007/s00422-008-0264-7

Nienborg, H., & Cumming, B. (2010). Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Curr Opin Neurobiol, 20*(3), 376-381. doi:10.1016/j.conb.2010.05.002

O'Keefe, J., & Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus, 15*(7), 853-866. doi:10.1002/hipo.20115

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res, 34*(1), 171-175.

Okun, M., Steinmetz, N., Cossell, L., Iacaruso, M. F., Ko, H., Bartho, P., . . . Harris, K. D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature, 521*(7553), 511-515. doi:10.1038/nature14273

Palva, J. M., Zhigalov, A., Hirvonen, J., Korhonen, O., Linkenkaer-Hansen, K., & Palva, S. (2013). Neuronal long-range temporal correlations and avalanche dynamics are correlated with behavioral scaling laws. *Proc Natl Acad Sci U S A, 110*(9), 3585-3590. doi:10.1073/pnas.1216855110

Peng, W., & Tang, D. (2016). Pain Related Cortical Oscillations: Methodological Advances and Potential Applications. *Front Comput Neurosci, 10*, 9. doi:10.3389/fncom.2016.00009

Petermann, T., Thiagarajan, T. C., Lebedev, M. A., Nicolelis, M. A., Chialvo, D. R., & Plenz, D. (2009). Spontaneous cortical activity in awake monkeys composed of

neuronal avalanches. *Proc Natl Acad Sci U S A, 106*(37), 15921-15926. doi:10.1073/pnas.0904089106

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., & Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature, 454*(7207), 995-999. doi:10.1038/nature07140

Pine, J. (2006). A History of MEA Development. In M. Taketani & M. Baudry (Eds.), *Advances in Network Electrophysiology*. Boston, MA: Springer.

Plenz, D. (2012). Neuronal avalanches and coherence potentials. *Eur. Phys. J. Spec. Top., 205*(1).

Plenz, D. (2013). Viewpoint: The Critical Brain. *Physics, 6*(47).

Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nat Neurosci, 7*(12), 1345-1352. doi:10.1038/nn1352

Quian Quiroga, R., & Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nat Rev Neurosci, 10*(3), 173-185. doi:10.1038/nrn2578

Quiroga, R. Q., Nadasdy, Z., & Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput, 16*(8), 1661-1687. doi:10.1162/089976604774201631

Ranade, S., Hangya, B., & Kepecs, A. (2013). Multiple modes of phase locking between sniffing and whisking during active exploration. *J Neurosci, 33*(19), 8250-8256. doi:10.1523/JNEUROSCI.3874-12.2013

Remme, M. W., Lengyel, M., & Gutkin, B. S. (2009). The role of ongoing dendritic oscillations in single-neuron dynamics. *PLoS Comput Biol, 5*(9), e1000493.

Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., & Harris, K. D. (2010). The asynchronous state in cortical circuits. *Science, 327*(5965), 587-590. doi:10.1126/science.1179850

Roux, F., & Uhlhaas, P. J. (2014). Working memory and neural oscillations: alpha-gamma versus theta-gamma codes for distinct WM information? *Trends Cogn Sci, 18*(1), 16-25. doi:10.1016/j.tics.2013.10.010

Sakmann, B., & Neher, E. (1984). Patch clamp techniques for studying ionic channels in excitable membranes. *Annu Rev Physiol, 46*, 455-472. doi:10.1146/annurev.ph.46.030184.002323

Salinas, E., & Abbott, L. F. (1994). Vector reconstruction from firing rates. *J Comput Neurosci, 1*(1-2), 89-107.

Schaefer, A. T., Angelo, K., Spors, H., & Margrie, T. W. (2006). Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. *PLoS Biol, 4*(6), e163. doi:10.1371/journal.pbio.0040163

Schroeder, C. E., & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci, 32*(1), 9-18. doi:10.1016/j.tins.2008.09.012

Schwartz-Bloom, R. D., & Sah, R. (2001). gamma-Aminobutyric acid(A) neurotransmission and cerebral ischemia. *J Neurochem, 77*(2), 353-371.

Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci, 16*(4), 1486-1510.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci, 18*(10), 3870-3896.

Shaukat, A., & Thivierge, J. P. (2016). Statistical Evaluation of Waveform Collapse Reveals Scale-Free Properties of Neuronal Avalanches. *Front Comput Neurosci, 10*, 29. doi:10.3389/fncom.2016.00029

Sherman, A., & Rinzel, J. (1992). Rhythmogenic effects of weak electrotonic coupling in neuronal models. *Proc Natl Acad Sci U S A, 89*(6), 2471-2474.

Shew, W. L., & Plenz, D. (2013). The functional benefits of criticality in the cortex. *Neuroscientist, 19*(1), 88-100. doi:10.1177/1073858412445487

Shew, W. L., Yang, H., Yu, S., Roy, R., & Plenz, D. (2011). Information capacity and transmission are maximized in balanced cortical networks with neuronal avalanches. *J Neurosci, 31*(1), 55-63. doi:10.1523/JNEUROSCI.4637-10.2011

Silva, L. R., Amitai, Y., & Connors, B. W. (1991). Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science, 251*(4992), 432-435.

Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron, 24*(1), 49-65, 111-125.

Skinner, F. K. (2017). Toward Theoretical and Experimental Synergies in Neuroscience: A Personal View. *Can J Neurol Sci*, 1-2. doi:10.1017/cjn.2017.215

Softky, W. R., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci, 13*(1), 334-350.

Song, K., Meng, M., Chen, L., Zhou, K., & Luo, H. (2014). Behavioral oscillations in attention: rhythmic alpha pulses mediated through theta band. *J Neurosci, 34*(14), 4837-4844. doi:10.1523/JNEUROSCI.4856-13.2014

Srinivas, K. V., Jain, R., Saurav, S., & Sikdar, S. K. (2007). Small-world network topology of hippocampal neuronal network is lost, in an in vitro glutamate injury model of epilepsy. *Eur J Neurosci, 25*(11), 3276-3286. doi:EJN5559 [pii]

10.1111/j.1460-9568.2007.05559.x

Staudigl, T., & Hanslmayr, S. (2013). Theta oscillations at encoding mediate the context-dependent nature of human episodic memory. *Curr Biol, 23*(12), 1101-1106. doi:10.1016/j.cub.2013.04.074

Steriade, M. (2000). Corticothalamic resonance, states of vigilance and mentation. *Neuroscience, 101*(2), 243-276.

Sussillo, D., & Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron, 63*(4), 544-557. doi:10.1016/j.neuron.2009.07.018

Tagliazucchi, E., Balenzuela, P., Fraiman, D., & Chialvo, D. R. (2012). Criticality in large-scale brain FMRI dynamics unveiled by a novel point process analysis. *Front Physiol, 3*, 15. doi:10.3389/fphys.2012.00015

Tajima, S., Mita, T., Bakkum, D. J., Takahashi, H., & Toyoizumi, T. (2017). Locally embedded presages of global network bursts. *Proc Natl Acad Sci U S A, 114*(36), 9517-9522. doi:10.1073/pnas.1705981114

Tauskela, J. S., Aylsworth, A., Hewitt, M., Brunette, E., & Mealing, G. A. (2012). Preconditioning induces tolerance by suppressing glutamate release in neuron culture ischemia models. *J Neurochem, 122*(2), 470-481. doi:10.1111/j.1471-4159.2012.07791.x

Tauskela, J. S., Fang, H., Hewitt, M., Brunette, E., Ahuja, T., Thivierge, J. P., . . . Mealing, G. A. (2008). Elevated synaptic activity preconditions neurons against an in vitro model of ischemia. *J Biol Chem, 283*(50), 34667-34676. doi:10.1074/jbc.M805624200

Thivierge, J. P. (2014). Scale-free and economical features of functional connectivity in neuronal networks. *Phys Rev E Stat Nonlin Soft Matter Phys, 90*(2), 022721.

Thivierge, J. P., & Cisek, P. (2008). Nonperiodic synchronization in heterogeneous networks of spiking neurons. *J Neurosci, 28*(32), 7968-7978. doi:10.1523/JNEUROSCI.0870-08.2008

Thivierge, J. P., & Cisek, P. (2011). Spiking neurons that keep the rhythm. *J Comput Neurosci, 30*(3), 589-605. doi:10.1007/s10827-010-0280-1

Thomas, C. A., Jr., Springer, P. A., Loeb, G. E., Berwald-Netter, Y., & Okun, L. M. (1972). A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exp Cell Res, 74*(1), 61-66.

Tiesinga, P., Fellous, J. M., & Sejnowski, T. J. (2008). Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci, 9*(2), 97-107. doi:10.1038/nrn2315

Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res, 23*(8), 775-785.

Touboul, J., & Destexhe, A. (2010). Can power-law scaling and neuronal avalanches arise from stochastic dynamics? *PLoS One, 5*(2), e8982. doi:10.1371/journal.pone.0008982

Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science, 286*(5446), 1943-1946.

Turesson, H. K., Logothetis, N. K., & Hoffman, K. L. (2012). Category-selective phase coding in the superior temporal sulcus. *Proc Natl Acad Sci U S A, 109*(47), 19438-19443. doi:10.1073/pnas.1217012109

van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science, 274*(5293), 1724-1726.

van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput, 10*(6), 1321-1371.

Velandia-Romero, M. L., Castellanos, J. E., & Martinez-Gutierrez, M. (2013). In vivo differential susceptibility of sensory neurons to rabies virus infection. *J Neurovirol*. doi:10.1007/s13365-013-0179-5

Vincent, K., Tauskela, J. S., Mealing, G. A., & Thivierge, J. P. (2013). Altered network communication following a neuroprotective drug treatment. *PLoS One, 8*(1), e54478. doi:10.1371/journal.pone.0054478

Vincent, K., Tauskela, J. S., & Thivierge, J. P. (2012). Extracting functionally feedforward networks from a population of spiking neurons. *Front Comput Neurosci, 6*, 86. doi:10.3389/fncom.2012.00086

Vogels, T. P., & Abbott, L. F. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat Neurosci, 12*(4), 483-491. doi:10.1038/nn.2276

Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., & Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science, 334*(6062), 1569-1573. doi:10.1126/science.1211095

Wagenaar, D. A., Pine, J., & Potter, S. M. (2006). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci, 7*, 11. doi:10.1186/1471-2202-7-11

Wang, H. P., Spencer, D., Fellous, J. M., & Sejnowski, T. J. (2010). Synchrony of thalamocortical inputs maximizes cortical reliability. *Science, 328*(5974), 106-109. doi:10.1126/science.1183108

Wang, X.-J., & Rinzel, J. (1992). Alternating and synchronous rhythms in reciprocally inhibitory model neurons. *Neural Comput, 4*(1), 84-97.

Wang, X. J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiol Rev, 90*(3), 1195-1268. doi:10.1152/physrev.00035.2008

Wang, X. J., Golomb, D., & Rinzel, J. (1995). Emergent spindle oscillations and intermittent burst firing in a thalamic model: specific neuronal mechanisms. *Proc Natl Acad Sci U S A, 92*(12), 5577-5581.

Wang, Y., Iliescu, B. F., Ma, J., Josic, K., & Dragoi, V. (2011). Adaptive changes in neuronal synchronization in macaque V4. *J Neurosci, 31*(37), 13204-13213. doi:10.1523/JNEUROSCI.6227-10.2011

Werner, G., & Mountcastle, V. B. (1963). The Variability of Central Neural Activity in a Sensory System, and Its Implications for the Central Reflection of Sensory Events. *J Neurophysiol, 26*, 958-977.

Williams, T. L. (1992). Phase coupling by synaptic spread in chains of coupled neuronal oscillators. *Science, 258*(5082), 662-665.

Wong, F. K., Bercsenyi, K., Sreenivasan, V., Portalés, A., Fernández-Otero, M., & Marín, O. (2018). Pyramidal cell regulation of interneuron survival sculpts cortical networks. *Nature, 557*(7707), 668-673. doi:10.1038/s41586-018-0139-6

Yu, S., Ribeiro, T. L., Meisel, C., Chou, S., Mitz, A., Saunders, R., & Plenz, D. (2017). Maintained avalanche dynamics during task-induced changes of neuronal activity in nonhuman primates. *Elife, 6*. doi:10.7554/eLife.27119

Zapperi, S., Baekgaard Lauritsen, K., & Stanley, H. E. (1995). Self-organized branching processes: Mean-field theory for avalanches. *Phys Rev Lett, 75*(22), 4071-4074. doi:10.1103/PhysRevLett.75.4071

Zbinden, C. (2010). *Leader neurons in living neural networks and in leaky integrate and fire neuron models.*, University of Geneve. (4183)

Zbinden, C. (2011). Leader neurons in leaky integrate and fire neural network simulations. *J Comput Neurosci, 31*(2), 285-304. doi:10.1007/s10827-010-0308-6

Zhang, S. J., Ye, J., Couey, J. J., Witter, M., Moser, E. I., & Moser, M. B. (2014). Functional connectivity of the entorhinal-hippocampal space circuit. *Philos Trans R Soc Lond B Biol Sci, 369*(1635), 20120516. doi:10.1098/rstb.2012.0516

Zhang, S. J., Ye, J., Miao, C., Tsao, A., Cerniauskas, I., Ledergerber, D., . . . Moser, E. I. (2013). Optogenetic dissection of entorhinal-hippocampal functional connectivity. *Science, 340*(6128), 1232627. doi:10.1126/science.1232627