

2016

Point process modeling as a framework to dissociate intrinsic and extrinsic components in neural systems

<https://hdl.handle.net/2144/19065>

"Downloaded from OpenBU. Boston University's institutional repository."

BOSTON UNIVERSITY
SCHOOL OF MEDICINE

Dissertation

**POINT PROCESS MODELING AS A FRAMEWORK TO DISSOCIATE
INTRINSIC AND EXTRINSIC COMPONENTS IN NEURAL SYSTEMS**

by

GRANT MICHAEL FIDDYMENT

A.B., University of Georgia, 2009
B.S., University of Georgia, 2009
M.A., University of Georgia, 2011

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2016

Approved by

First Reader

Mark Kramer, Ph.D.
Associate Professor, Mathematical Neuroscience

Second Reader

Uri Eden, Ph.D.
Associate Professor, Statistics

Third Reader

Jason Ritt, Ph.D.
Assistant Professor, Biomedical Engineering

DEDICATION

Dedicated to Hadas Peles, whose boundless love and support have made this endeavor possible,

And to the memory of Gordon Fiddymont, whose passion for numbers and the written word lives on.

ACKNOWLEDGMENTS

First and foremost, I want to extend my profound gratitude to Mark Kramer and Uri Eden, who have been endlessly knowledgeable and supportive mentors throughout this project. Our lab meetings were a highlight of my graduate experience – both for the collegial atmosphere and the depth of feedback and practical knowledge regularly shared. In particular, I must thank Manu Martinet, Kyle Lepage, Mikio Aoi, Wei Tang, Xinyi Deng, Emily Stephen, and Stefania Sokolowski for their many helpful discussions.

Thanks to Nancy Kopell and Jason Ritt for guidance. Their keen insights and healthy skepticism challenged me to think critically and greatly enhanced the quality of this work.

I would also like to thank Sydney Cash, Omar Ahmed, Jason Naftulin and the epilepsy lab at Massachusetts General Hospital for sharing the data at the core of this thesis and helping in its analysis and interpretation. Our hack-a-thons were another highlight to be remembered.

Thanks to Nicholas James and Joan Martinez for sharing experimental data sets used in Chapters III and IV; to Jason Sherfey, Austin Soplata, Caroline Moore-Kochlacs, and Dave Stanley for sharing expertise in computing and technology; and to the NIH and Epilepsy Foundation for funding that enabled Chapters II and III.

I am extremely grateful to the Allen Institute for Brain Science for hosting me at their summer workshop, where many of the ideas for Chapter IV took shape. Thanks especially to Adrienne Fairhall and Christof Koch for organizing the event, and to all the

faculty and attendees – particularly Lydia Ng, Michael Buice, Anton Arkhipov, Jeremy Freeman, Nicholas Cain, Kenneth Lattimer, Aaron Oldre, Josh Merel, Taylor Newton, and Gioia Di Francheschi. The spirited atmosphere and thoughtful pursuit of cooperative, innovative science were deeply inspiring.

Finally, I must acknowledge Shelley Russek and Sandi Grasso for their tireless work in cultivating the neuroscience community at Boston University. I can imagine no better place to work over the last several years. It will always feel like home.

**POINT PROCESS MODELING AS A FRAMEWORK TO DISSOCIATE
INTRINSIC AND EXTRINSIC COMPONENTS IN NEURAL SYSTEMS**

GRANT MICHAEL FIDDYMENT

Boston University School of Medicine, 2016

Major Professor: Mark Kramer, Ph.D., Associate Professor, Mathematical Neuroscience

ABSTRACT

Understanding the factors shaping neuronal spiking is a central problem in neuroscience. Neurons may have complicated sensitivity and, often, are embedded in dynamic networks whose ongoing activity may influence their likelihood of spiking. One approach to characterizing neuronal spiking is the point process generalized linear model (GLM), which decomposes spike probability into explicit factors. This model represents a higher level of abstraction than biophysical models, such as Hodgkin-Huxley, but benefits from principled approaches for estimation and validation.

Here we address how to infer factors affecting neuronal spiking in different types of neural systems. We first extend the point process GLM, most commonly used to analyze single neurons, to model population-level voltage discharges recorded during human seizures. Both GLMs and descriptive measures reveal rhythmic bursting and directional wave propagation. However, we show that GLM estimates account for covariance between these features in a way that pairwise measures do not. Failure to account for this covariance leads to confounded results. We interpret the GLM results to speculate the mechanisms of seizure and suggest new therapies.

The second chapter highlights flexibility of the GLM. We use this single framework to analyze enhancement, a statistical phenomenon, in three distinct systems. Here we define the enhancement score, a simple measure of shared information between spike factors in a GLM. We demonstrate how to estimate the score, including confidence intervals, on a simulated network. In real networks, we find that enhancement occurs prominently during human seizure, while redundancy tends to occur in mouse auditory networks. We discuss implications for physiology, particularly during seizure.

In the third part of this thesis, we apply point process modeling to spike trains recorded from single units *in vitro* under external stimulation. We re-parameterize models in a low-dimensional and physically interpretable way; namely, we represent their effects in principal component space. We show that this approach successfully separates the neurons observed *in vitro* into different classes consistent with their gene expression profiles.

Taken together, this work contributes a statistical framework for analyzing neuronal spike trains and demonstrates how it can be applied to create new insights into clinical and experimental data sets.

TABLE OF CONTENTS

TITLE PAGE	i
COPYRIGHT	ii
APPROVAL PAGE	iii
DEDICATION	iv
ACKNOWLEDGMENTS	v
ABSTRACT	vii
TABLE OF CONTENTS	ix
LIST OF TABLES	xii
LIST OF FIGURES	xiii
LIST OF ABBREVIATIONS	xv
CHAPTER I. INTRODUCTION	1
1. Issues in Spike Train Analysis	2
2. Generalized Linear Models in Neuroscience	3
2.1 Point process GLM	3
2.2 Applications of point process models in neuroscience	9
3. Summary of Dissertation: Motivation and Approach	10
3.1 Chapter II. Two categories of ictal discharges propagate with different spatiotemporal dynamics during human seizure	11

3.2 Chapter III. Point process modeling reveals a unique type of enhancement during human seizures.....	12
3.3 Chapter IV. Point process modeling reveals a hierarchy of functional cell types based on self-history dependence	14
CHAPTER II. TWO CATEGORIES OF ICTAL DISCHARGES PROPAGATE WITH DIFFERENT SPATIOTEMPORAL DYNAMICS DURING HUMAN SEIZURE	17
1. Introduction.....	17
2. Methods.....	20
2.1. Patients and recordings	20
2.2. Ictal discharge identification and clustering	22
2.3. Calculation of speed.....	26
2.4. Point process analysis	27
2.5. Directional analysis.....	31
3. Results.....	34
4. Discussion.....	52
CHAPTER III. POINT PROCESS MODELING REVEALS A UNIQUE TYPE OF ENHANCEMENT DURING HUMAN SEIZURES	59
1. Introduction.....	59
2. Methods.....	60
2.1. Data	61
2.2. Point process modeling	64
3. Results.....	71

4. Discussion.....	86
CHAPTER IV. POINT PROCESS MODELING REVEALS A HIERARCHY OF FUNCTIONAL CELL TYPES BASED ON SELF-HISTORY DEPENDENCE.....	
1. Introduction.....	92
2. Methods.....	93
2.1. Medial entorhinal cortex data	94
2.2. Allen Institute Cell Types data	95
2.3. Point process model	98
3. Results.....	101
4. Discussion	108
CHAPTER V. CONCLUSION.....	
1. Innovation and Impact	111
1.1. A rigorous and patient-specific approach to seizure analysis	111
1.2. A generalizable method for measuring enhancement	112
1.3. A method for clustering neurons according to their function	112
2. Future Directions	113
2.1. Understanding and extending enhancement	113
2.2. Clustering by functional profile	116
BIBLIOGRAPHY	118
CURRICULUM VITAE.....	136

LIST OF TABLES

Table II.1. Patients with implanted MEA.	21
Table II.2. Total counts of identified ictal discharges, large amplitude discharges, and small amplitude discharges.	23
Table II.3. Features of spatial waves for LADs and SADs.	26
Table II.4. Summary of distributions of LAD directionality.	47
Table II.5. Summary of differences in directionality between large and small amplitude discharges.	51
Table III.1. Differences between estimated and true parameters for individual and joint models.	78
Table IV.1. Cre driver lines in the Allen Cell Types Database.	97
Table IV.2. Spiking profile clusters correlate with genetic cell type in MEC interneurons.	103

LIST OF FIGURES

Figure II.1. Ictal discharges propagate in waves with multi-faceted structure.	25
Figure II.2. Overview of analysis approach for estimation of spatiotemporal structure from point process data.	33
Figure II.3. Correlation analyses reveal multiple time scales and directionality in ID dynamics but suffer from confounding effects.	39
Figure II.4. Modeling analyses reveal multiple time scales and directionality in ID dynamics without the confounds of descriptive analyses.	43
Figure II.5. Modeling analysis of small amplitude discharges reveals effects that depend on both SADs and LADs.	49
Figure III.1. Enhancement is measured by analyzing the model hierarchy.	70
Figure III.2. Enhancement depends on the type of system feedback.	73
Figure III.3. Confidence intervals for the enhancement score can be reliably bootstrapped with data.	81
Figure III.4. Enhancement is prevalent in ictal discharges during human seizure.	84
Figure III.5. Mouse cortical networks show redundancy and independence between spike trains during passive listening.	85
Figure IV.1. MEC interneurons' spiking profiles can be clustered into groups based on their estimated intrinsic effects.	102
Figure IV.2. Allen Database neurons' spiking profiles can be clustered into groups based on their estimated intrinsic effects.	105

Figure IV.3. Spiking profile clusters correlate with genetic cell type from the Allen Cell

Types Database 107

LIST OF ABBREVIATIONS

ACF / CCF	Auto / cross-correlation function
AC	Auditory cortex
CIF	Conditional intensity function
FS	Fast spiking
GAD	Glutamic acid decarboxylase
GLM	Generalized linear model
ID	Ictal discharge
IID	Inter-ictal discharge
ISI	Inter-spike interval
KS	Kolmogorov-Smirnov
LAD	Large amplitude discharge
LIF	Leaky integrate-and-fire
LFP	Local field potential
MEA	Micro-electrode array
MEC	Medial entorhinal cortex
PCA	Principal component analysis
PFC	Prefrontal cortex
PV	Parvalbumin
PID	Pre-ictal discharge
SAD	Small amplitude discharge

CHAPTER I. INTRODUCTION

Action potentials, or spikes, are thought to be the primary means of communication in the brain. Characterizing ongoing spike patterns and the factors that affect them is a fundamental problem spanning both cellular and systems neuroscience. At the cellular level, spikes convey information about individual neurons, such as the types of ion channels present, how a neuron adapts to changing inputs, and the cell's preferred stimulus. At the network level, spikes shed light on coordination among cells, providing insight into which neurons are connected and how information flows through a network.

Computational theories of the brain seek to address both how neurons spike (e.g. the biophysical or dynamical mechanisms) and why spike patterns have particular structure (e.g. evoked by a stimulus or ongoing correlations). Exploration of these complex issues benefits from two complementary modeling approaches. The first approach is to build detailed mathematical models with hundreds or thousands of equations and variables (Traub et al., 2005a; Markram et al., 2015; Hawrylycz et al., 2016). Many neuroscientists prefer such models because the model terms represent familiar – and directly controllable – quantities. However, this high-dimensional representation comes at a cost. Analyzing detailed models typically involves the numerical integration of a large number of coupled differential equations with many unknown parameters; a complete understanding of the model's behavior can be time consuming, and the best procedure to constrain unknown model parameters is not known.

The complementary approach is to formulate statistical models with a small number of equations and variables (Okatan et al., 2005). Statistical models trade precision for conciseness; as a result, statistical models are relatively easy to simulate and analyze, and principled methods exist to constrain these models directly from data.

This dissertation deals primarily with the second approach: building small, analytically tractable models of neural spiking. We apply a specific modeling framework, the point process generalized linear model (GLM), to five different sets of spike trains. These spike trains consist of single- and multi-unit action potentials, simulated data, and a non-standard type of spike: a population-wide voltage discharge that occurs during human seizure, or “ictal spike.” With only minor modifications, the same GLM framework successfully characterizes all five sets of spike data. In this chapter, we first discuss general approaches and issues to model-building in neuroscience. We then briefly review concepts related to the point process GLM and survey its use in neuroscience. Finally, we summarize the results presented in Chapters II-IV, the three implementations of the point process GLM that comprise this dissertation. We conclude in Chapter V with a summary of the research and proposals for continued work.

1. Issues in Spike Train Analysis

Neuroscientists probe the brain by collecting experimental data and explaining it with mathematical models. Modern technologies like spatially-dense electrodes (Blanche

2005; Scholvin et al. 2016), calcium imaging (Kerr and Denk 2008; Hawrylycz et al. 2016), and optogenetics (Stanley 2013) provide high-resolution ways to record and control neural circuits. They also generate large data sets, which demand powerful and efficient tools for spike train analysis (Brown et al. 2004; Horwitz 2016). Point process theory, a statistical framework for characterizing event data such as spike trains, offers a language with which to develop these tools.

2. Generalized Linear Models in Neuroscience

In this section we briefly motivate and define concepts related to the generalized linear model (GLM) framework for spike train analysis. We then review previous uses of the GLM in neuroscience, highlighting several practical advantages. Although the review focuses on the GLM's role in analyzing neural coding, point process theory can be applied in other contexts, such as inferring spike times from calcium imaging (Vogelstein et al. 2009; Pnevmatikakis et al. 2013).

2.1 Point process GLM

2.1.1. Point processes. Spikes are stereotyped events isolated in time. A spike train, then, resembles a set of points plotted along the continuous time axis. In statistics such data are known as temporal *point processes*. Likewise, events isolated in space are known as spatial point processes. For a review of statistical point process theory, including topics beyond the present scope, see (Cox and Isham 1980; Daley and Vere-Jones 2002).

There are several ways to represent a point process: as a sequence of event times, waiting times, or counts. All three are equivalent, so we can flexibly move between them if one form lends itself to a particular application. Typically, spike trains are represented as event times during the recording stage and later converted to count sequences for analysis; this is the approach we use here. So-called “binless” representations have been used in the analysis of single neurons (Victor 2002) but become analytically intractable for more complex cases. Here we briefly describe the process for translating event times to counts, introducing notation that will be used along the way.

Suppose a spike train recorded over the time interval $[0, T]$ consists of N spikes at times s_1, \dots, s_N and let the function $n(t)$ give the total number of spikes observed up to time t . We first divide time into smaller windows of duration Δ , giving time points $t_k = k\Delta$ for $k = 0, \dots, T/\Delta$. Now in discrete time, we define the sequence $dn_k = n(t_k) - n(t_{k-1})$ $k \geq 1$, which counts the number of spikes in each time interval $[t_{k-1}, t_k]$. If the bin size Δ is chosen to be sufficiently small, the count sequence dn_k will be binary: either 1 if a spike occurs in bin k or 0 otherwise.

2.1.2. Conditional intensity function. We analyze a point process through the conditional intensity function (CIF), which describes its history-dependent structure. As the name implies, the CIF is the conditional probability defined by the formula,

$$\lambda(t | H_t) = \lim_{\Delta \rightarrow 0} \frac{\Pr[n(t+\Delta) - n(t) = 1 | H_t]}{\Delta} \text{(Equation 1.1)}$$

where H_t is the set of covariates that influence spiking, all observed up to time t (Daley and Vere-Jones, 2002; Truccolo et al., 2005). The goal of model selection, addressed in Chapters II and III, is to identify and interpret the set of covariates H_t in light of the observed system. In neuroscience applications, H_t typically involves a combination of (1) the previous history of spiking for the cell being modeled; (2) the previous history of spiking for additional cells; (3) inputs to the neuron due to external stimuli, and/or (4) oscillations in the extracellular voltage. In the subsequent chapters, we consider models with history-dependence involving factors (1) and (2), which mimic the biophysical features of intrinsic membrane channels and synaptic input, respectively. Frequently we refer to (1) as intrinsic effects and (2) as extrinsic effects.

2.1.3. Generalized linear model. The point process GLM, as opposed to other point process models, assumes a particular form for the conditional intensity function.

Specifically, it assumes the intensity can be written in the form,

$$\lambda_k = f(X\beta) \text{ (Equation 1.2).}$$

The function f , called the inverse link, transforms a classical linear model ($X\beta$) into the distribution of the response variable (in this case, λ_k). In the case where X is a stimulus and β a stimulus filter, or tuning curve, Equation 1.2 has also been called a Linear-Nonlinear Poisson (LNP) model (Simoncelli et al. 2004). A common choice of link, used throughout this dissertation, is the exponential link,

$$\lambda_k = \exp(X\beta)$$

or, equivalently,

$$\log \lambda_k = X\beta \text{ (Equation 1.3).}$$

For comparisons of different choices of link function, see (Paninski et al., 2004b; Ahrens et al., 2008; McFarland et al., 2013).

We estimate the model parameters β by maximum-likelihood. That is, we fix dn_k to be an observed spike train and optimize the likelihood function

$$L = \exp \left(\sum_{k=1}^K dn_k \log(\lambda_k \Delta) - \lambda_k \Delta \right) + o(\Delta^N)$$

or, equivalently, the log-likelihood function

$$\log L = \sum_{k=1}^K dn_k \log(\lambda_k \Delta) - \lambda_k \Delta + o(\Delta^N) \text{ (Equation 1.4),}$$

where $o(\Delta^N)$, higher-order terms introduced from approximating the limit in Equation 1.1, are negligible (see Truccolo et al. 2005). The log-likelihood is convex as a function of β and therefore has no local maxima. As a result, max-likelihood solutions exist and can be estimated rapidly through gradient descent (McCullagh and Nelder 1989).

One way to mitigate overfitting, called regularization, is to add a penalty term to the log-likelihood before optimizing that will favor models with fewer parameters. For uses of regularization in fitting a point process GLM see (Gerwinn et al., 2010; Chen et al. 2009; Kelly et al. 2010).

2.1.4. Goodness of fit. An important advantage of statistical models over biophysically motivated mathematical models is that principled methods exist to quantitatively assess

the statistical model's goodness of fit. Several methods have been proposed to test goodness of fit for point process models. Perhaps the most well-known method is based on the time rescaling theorem (Brown et al. 2002), which rescales the waiting times and compares them with an exponential distribution. However, researchers have noted several issues with approaches based on time rescaling. One is that certain requisite assumptions break down for discrete data with high firing rates and short interspike-intervals (Haslinger et al., 2010). To deal with this, other goodness of fit tests based on thinning and complementing have been proposed (Gerhard and Gerstner 2010). An additional issue is that time rescaling assesses residual structure only within a single spike train. As a result, models could ignore important cross-population structure but still pass a goodness of fit test. Gerhard et al. (2011) address this issue by developing a multivariate version of the test that can be applied to an entire network.

2.1.5. Relationship between point process GLM and the integrate-and-fire model

neuron. Developed before Hodgkin and Huxley revealed the specific mechanisms of the action potential (Abbott 1999), the leaky integrate-and-fire (LIF) model is the most basic phenomenological model of a spiking neuron. The LIF model is defined as follows:

$$\tau \dot{V} = -V + I(t); \text{ if } V(t) \geq 1, V(t + dt) = 0.$$

In this model, the differential equation describes how the cell integrates time-varying current $I(t)$ while it is opposed by a voltage “leak” $-V$ that drives the voltage towards zero. If the input current exceeds the leak, the voltage eventually reaches one, and is then

reset to 0; at this time of reset, we say that the LIF model generates a “spike”.

Additionally, several authors have extended the LIF model to incorporate broader spiking dynamics such as bursting and afterhyperpolarization (AHP) currents (Smith and Sherman 2002; Casti et al. 2002; Gerstner et al. 2014).

The LIF model and its extensions can all be completely specified using an appropriate GLM. Stevens and Zador (1996), for example, derive the mapping from a noisy LIF model to an equivalent Poisson GLM. From here, capturing additional LIF model dynamics with a GLM becomes relatively straightforward. For example, membrane currents map directly onto intrinsic effects terms in the GLM (Jolivet 2004; Jolivet and Gerstner 2004; Gerstner et al. 2014), while synaptic inputs to an LIF neuron map onto GLM extrinsic effects (Latimer et al. 2014).

2.1.6. Applications of the point process modeling framework in other fields.

Although here we focus on models of neural spiking, the point process modeling framework can be used to analyze other kinds of events as well. Indeed, point process models have been used to model diverse phenomena including military activity (Zammit-Mangion et al. 2012), social network behavior (Zadeh and Sharda 2015), ambulance demand (Zhou et al. 2015b), and basketball shot quality (Franks et al. 2015). In the next section we briefly summarize some applications of the GLM framework to problems in neuroscience.

2.2. Applications of point process models in neuroscience

GLMs have been applied to study how neural systems encode stimuli as well as to decode their spike activity (Paninski et al., 2007). In encoding analyses, the GLM is used to describe how neurons – both single cells and networks – spike in response to stimuli. For example, Brown et al. (2001) estimated place fields of neurons in the rat hippocampus in vivo. To do so, they developed an adaptive estimation procedure, later extended by Eden et al. (2004) and Ergun et al. (2007), that successfully tracked receptive fields evolving over a span of minutes. Similarly, for single neurons the GLM has been used to characterize receptive fields in the retina (Pillow 2005), tuning curves related to movement direction and velocity in motor cortex (Paninski et al. 2004a; Truccolo et al. 2005), spectrotemporal receptive fields in the auditory midbrain (Calabrese et al. 2011), movement-induced dynamics in subcortical structures related to Parkinson's disease (Sarma et al. 2008; Deng et al., 2013), and task-related changes in spiking in macaque LIP (Latimer et al. 2015) and SEF (Ventura et al. 2002). The GLM has also been applied to model network activity. In particular, it has been used to describe sensory encoding by networks in the macaque retina (Pillow et al. 2008; Vidne et al. 2012), cat motor cortex (Chen et al. 2009), and macaque V1 (Gerhard et al. 2011; Kelly and Kass 2012), and, in small cases, validated against simulated networks (Chornoboy et al., 1988; Okatan et al., 2005; Kim et al. 2011).

A dual application of the GLM is to decode the signal giving rise to a particular spike train. For example, place cell spiking has been used to decode animal location and

future navigational decisions (Huang et al. 2009; Deng et al. 2015). Similarly, spiking in monkey LIP has been used to predict a monkey's decision of where to saccade (Park et al. 2014), and multi-unit activity from motor cortex to decode monkeys' hand movements (Lawhern et al. 2010).

In addition, the GLM framework has been developed to overcome known confounds present in the descriptive analysis of spike train data. For example, descriptive measures of spike-field coherence depend on firing rate (Lepage et al., 2011); this confound can be addressed through the development of a GLM that relates spiking to the inferred phase of a field (Lepage et al. 2013; Zhou et al. 2015a). Measures of correlation between spike trains also suffer from confounds (Moore et al. 1970; Cohen and Kohn 2011), as we discuss in detail in Chapter II.

3. Summary of Dissertation: Motivation and Approach

This dissertation utilizes a single framework, namely the point process GLM, to analyze spike trains in different *in vivo*, *in vitro*, and *in silico* scenarios. The three chapters describe these applications (or case studies) in modeling with the GLM. Where appropriate, we extend current methods and discuss practical advantages related to the GLM approach. In the first chapter, we apply the GLM to analyze spatiotemporal patterns during seizure. The second chapter uses this modeling framework to estimate a quantity known as enhancement or synergy. The third chapter applies the GLM to cluster

neurons by functional profile. Taken together, this work demonstrates the power and flexibility of modeling spike trains with a point process GLM.

3.1. Chapter II. Two categories of ictal discharges propagate with different spatiotemporal dynamics during human seizure

Chapter II applies the GLM framework to ictal discharges (IDs), a type of population spike data recorded in microelectrode arrays (MEA). These population-level voltage fluctuations travel across the MEA with rich spatiotemporal structure. By developing an algorithm that automatically extracts IDs from voltage data, we were able to identify for each seizure hundreds of IDs per electrode. We first analyzed the slope and amplitude of the discharges and found that they naturally follow a bimodal distribution. This suggests there are two types of discharge, so we clustered the IDs and assigned each one to its most likely type. We call these two types “large amplitude discharges” (LADs) and “small amplitude discharges” (SADs).

To explore the spatiotemporal structure of LADs and SADs, we utilized two methods: descriptive measures (namely auto- and cross-correlation) and the point process GLM. We find that both approaches reveal signatures of rhythmic bursting and directional wave propagation. The direction of wave propagation, whether estimated by correlation or GLM, was highly consistent across multiple seizures per patient. Correlations, however, are confounded and become spatially homogeneous at long lags, whereas GLM estimates show spatial specificity at long lags. We use the GLM

framework to show that the discharges evolve temporally as bursts, which become more rhythmic approaching seizure termination. Although both LADs and SADs propagate spatially as waves over the cortical surface, the large amplitude discharges propagate with more spatial organization than the smaller amplitude discharges. These spatiotemporal features remain consistent for each patient's seizures. Such characterizations provide insight into the different types of spatiotemporal dynamics displayed by ictal discharges during human seizure at the sub-millimeter spatial scale, and provide clues to the mechanisms of human seizure and possible targets for improved surgical therapy.

3.2. Chapter III. Point process modeling reveals a unique type of enhancement during human seizures

In analyzing the GLMs discussed in Chapter II, we observed evidence of enhancement, a curious and little-reported statistical phenomenon. Typically, we expect that if two correlated variables are used to predict a modeled variable, then the two variables will be partially redundant in their explanation. In fact, some statistics textbooks state – incorrectly – that the variability explained by a second model covariate is always less than the information that same covariate provides alone. However, this is not necessarily true. There are a limited number of cases in social sciences, fewer still in neuroscience, where the converse occurs. In these situations, information behaves super-linearly: two variables provide more information in a joint model than they provide in

total individually. In the statistical literature, this phenomenon has been variously called enhancement, synergy, and suppression. For clarity, we call this phenomenon enhancement. In Chapter III, we plainly define and analyze enhancement in terms of GLM.

We first define an enhancement score for two covariates that registers negative in cases of redundancy (i.e., where two variables provide less information in a joint model than they provide in total individually) and positive in cases of enhancement. This score is based on the GLM deviance, a well-known measure that is straightforward to compute. We then propose a data bootstrapping procedure by which confidence bounds for the score can be estimated. Using a simulated two-cell network, we validate this data bootstrapping procedure against parameter bootstrapping, or simulating more data. Over a range of network configurations leading to both redundancy and enhancement, we find that confidence intervals for data bootstrapping are highly consistent with parameter bootstrapping. This suggests that bootstrapping from a single block of data - how the technique is applied in practice - gives similar confidence intervals as a very large data sample, or the theoretical distribution of scores.

We then proceed to estimate the enhancement score with confidence for two neural data sets. For the ictal discharges introduced in Chapter II, we find a consistent result for the entire population of patients and seizures: strong evidence of a positive score, consistent with enhancement between intrinsic and extrinsic effects. We continue and apply the same technique to analyze multi-unit spiking recorded in mouse auditory

and pre-frontal cortex. Assessing the enhancement score in these data, we find the opposite result: scores are predominantly negative, indicating redundancy between intrinsic and extrinsic effects.

We describe in this chapter a new method for evaluating enhancement using GLMs. Although enhancement is well-described in the classical linear modeling literature, much less is known for generalized linear models. As a result, the classical case offers some insight into enhancements; nevertheless, the full implications of enhancement – for seizure in particular and neural systems in general – remains an active research area to be explored. We conclude by interpreting these results in light of the underlying neural networks, particularly for seizure.

3.3. Chapter IV. Point process modeling reveals a hierarchy of functional cell types based on self-history dependence

Identifying the various types of neurons is a critical problem in neuroscience. Increasingly, cells can be biologically classified using genetic assays; however the functional consequences of neuronal biology remains incompletely understood. Some overlap between the schemes has been observed. For instance, genetically-identified excitatory principal neurons are often modeled as regular-spiking (RS) and genetically-identified parvalbumin-positive (PV+) interneurons as fast-spiking (FS) cells. In this chapter, we describe a procedure to extend such labeling by clustering neurons into

functional cell types using the point process GLM. Our approach complements biological classification schemes by describing neurons purely in terms of their spike outputs.

We apply the clustering procedure to two sets of spike trains collected from neurons *in vitro*. The first data set consists of PV+ and non-PV+ inhibitory interneurons in the mouse medial entorhinal cortex (MEC). The second data set, from the publicly available Allen Cell Types Database, consists of ten cell types including both excitatory and inhibitory neurons. After fitting GLMs to the cells' spike trains, we dimensionally-reduce the model estimates in various ways, including principal component analysis (PCA) and k-means clustering, and compare their performance. We find that the range of GLM intrinsic effects is effectively low-dimensional; that is, for both data sets the intrinsic effects cluster into a small number of types. We find notable overlap between functional types assigned from the GLM estimates and neurons' genetic lineage. In the MEC interneurons, we find that clusters that reliably separate PV+ and non-PV+ cells. In the Allen Cell Types data, we find that the two largest clusters correspond closely to excitatory and inhibitory cells. Furthermore, the clustering also hints at sub-types within and across the classical divisions. These techniques provide a powerful method to establish taxonomy of functional cell types across cortex that complements our understanding of biological cell types.

3.4. Chapter V: Conclusion

We conclude by summarizing the main research results, discussing the implications for the broader scientific community, and proposing future directions of research.

CHAPTER II. TWO CATEGORIES OF ICTAL DISCHARGES PROPAGATE WITH DIFFERENT SPATIOTEMPORAL DYNAMICS DURING HUMAN SEIZURE

1. Introduction

Epilepsy, one of the most common neurological problems, is a devastating disease, impacting over five million people in the United States alone. Although epilepsy is an ancient disease, observed for thousands of years (Foote Smith and Bayne 1991; Ozer 1991), many aspects of this disease remain poorly understood (Lado and Moshé 2008; Frei et al. 2010; Jiruska et al. 2013; Krook-Magnuson and Soltesz 2015). To completely understand and treat epilepsy remains an active research challenge that benefits from numerous approaches including experimental (Wagner et al. 2015; Paz et al. 2013), computational (Traub et al., 2005; Lytton 2008; Destexhe 1998; Fröhlich et al., 2010), and clinical (Rummel et al. 2013; Blumenfeld et al. 2009).

The most prominent characteristic of epilepsy is the repeated, spontaneous occurrence of seizures. Although the electrographic components of a seizure are varied, one of the most common features consists of brief (5-20 ms), sharp (at least 10 standard deviations / sec) changes in voltage. Voltage discharges in epilepsy can be divided into three types: interictal, pre-ictal, and ictal; and can manifest over large volumes, even brain wide (Stufflebeam et al. 2011; Lüttjohann et al., 2014; Sabolek et al. 2012). Interictal discharges - which occur between seizures - are useful in diagnosing epilepsy,

although the relationships between interictal discharges and seizures are not completely understood (Staley, Hellier, and Dudek 2005). Interictal discharges have been associated with the development of cognitive deficits and memory impairment (Kleen et al. 2013; Gelinas et al. 2016), and proposed as useful for targeting the epileptogenic zone (Alarcon et al. 1997). During an interictal discharge, it is thought that local cortical neurons generate action potentials due to a synchronous, paroxysmal membrane depolarization (Prince and Connors 1986; Dichter and Spencer 1969), although the mechanisms by which interictal discharges initiate and spread is not yet completely defined (Sabolek et al. 2012). Pre-ictal discharges - which occur immediately before seizure onset - have been found in many brain areas (Bartolomei et al. 2004; Huberfeld et al. 2011) and correlated with increased glial density (Spencer et al. 1999). A recent analysis of human tissue slices from patients with mesial temporal lobe epilepsy further distinguishes these two classes of discharges - pre-ictal and interictal - which differ in amplitude (pre-ictal is larger) and the cellular networks that support the discharges (Huberfeld et al. 2011). How these interictal and pre-ictal discharges organize and spread in the cortical network - and relate to ictal discharges - remains incompletely understood. Improving this understanding is the primary goal of this manuscript.

Recently, microelectrode arrays (MEAs) - implanted in the superficial layers of human cortex - have provided an unprecedented view of human brain activity at high spatial (< 0.5 mm) and temporal (e.g., 30,000 Hz) resolution. These data have revealed

fine-scale phenomena such as microseizures (Schevon et al. 2008), and insights into the activity of individual neurons (Truccolo et al. 2011) and multi-unit activity during seizure (Schevon et al. 2012). Although aspects of these data remain controversial (e.g., the isolation of individual units (Merricks et al. 2015); the appearance of an ictal wavefront (Wagner et al. 2015)) a common theme has begun to emerge: at the sub-millimeter spatial scale, human seizure activity exhibits a complex yet organized spatiotemporal structure.

In this manuscript, we analyze the spatiotemporal organization of ictal discharges observed at the sub-millimeter spatial scale in four patients. Consistent with recent *in vitro* observations, we show that these discharges naturally divide into two groups with distinct electrophysiological features. Analysis of the spatiotemporal dynamics of these two discharge groups, using both descriptive methods and statistical modeling, shows that – despite their different electrophysiological features – both groups exhibit similar, rhythmic temporal organization, which manifests as bursts consisting of doublet and triplet sequences of discharges approaching seizure termination. Consistent with this similarity in temporal organization, both discharge groups propagate in similar spatial directions. However, the spatial organization of this propagation is higher for the larger amplitude discharges, compared to the smaller amplitude discharges. These analyses reveal both similarities and distinctions between the spatiotemporal dynamics of the two groups of ictal discharges, provide insights into the possible mechanisms that support these discharges, and suggest refined targets for surgical intervention.

2. Methods

2.1. Patients and recordings

Four patients with medically intractable focal epilepsy underwent clinically indicated intracranial cortical recordings using grid electrodes for epilepsy monitoring (see Table 1). Clinical electrode implantation, positioning, duration of recordings and medication schedules were based purely on clinical need as judged by an independent team of physicians. The patients were implanted with a 10 x 10 (4 mm x 4 mm) microelectrode array (MEA; Blackrock Microsystems, Utah) in a neocortical area expected to be resected with high probability. This research probe consisted of 96 recording platinum-tipped silicon probes, with a length of either 1-mm (Patient D) or 1.5-mm (Patients A-C), corresponding to neocortical layer III as confirmed by histology after resection. Signals from the MEA were acquired continuously at 30 kHz per channel. The reference electrode was either subdural or epidural, chosen dynamically based on recording quality.

Seizure onset and end times were determined by an experienced encephalographer through inspection of clinical voltage recordings, referral to the clinical report of the invasive electroencephalogram and clinical manifestations recorded on video. The number of seizures varied across the patients. Owing to operational issues, not all of the seizures were recorded or provided data with a high signal-to noise ratio. We selected 11 seizures from the four patients. Seizure onsets were detected 2-3 cm away from the

research MEA, based on the clinical voltage electrodes. These recordings were therefore outside the seizure onset zone.

Table II.1. Patients with implanted MEA. TG: Temporal Gyrus.

Patient	#	Age: Onset	Age: Surgery	Sex	Hand	MEA Placement	Distance from seizure onset
A	3	10	32	M	L	Superior TG	3 cm
B	3	--	45	M	R	Middle TG	2 cm
C	2	14	25	F	L	Middle TG	3 cm
D	3	15	21	M	R	Middle TG	2 cm

These data have been previously used in other studies. Patients A and D correspond to LFP #1 and #2 in (Kramer et al. 2012). Patient D corresponds to Patient B in (Truccolo et al. 2011). Patients A-D correspond to patients (2,1,3,4) in study (Wagner et al. 2015); detailed clinical information for each patient may be found in reference (Wagner et al. 2015).

This research was approved by local Institutional Review Boards at Massachusetts General Hospital/Brigham and Women's Hospitals (Partners Human Research Committee) and at Boston University according to National Institutes of Health guidelines.

2.2. Ictal discharge identification and clustering

Ictal discharges (IDs) are defined as large amplitude, brief, seizure-related fluctuations in brain voltage. To identify IDs quickly and consistently, we implemented the following automated detection procedure. First, we excluded electrodes with excessive noise (average autocorrelation from 0-2 ms < 0.985). For most patients, there were 2-4 excluded channels; in one case (Patient D), there were 11 excluded channels. We note that we only analyzed the interior 8 x 8 sub-grid of the MEA, so that each electrode possessed neighbors in all directions. We then bandpass filtered the voltage data (1-20 Hz, linear-phase FIR filter, order 30, zero-phase filtering) and z-scored the result for each electrode over the entire duration of the seizure.

To extract the IDs, we detected extrema in the filtered and z-scored voltage activity. Conceptually, an ID occurs at a voltage extremum, where the change in voltage over time - or the slope - reaches zero; near the extremum, the slope changes sign between (large) negative and positive values. In practice, we computed the slope over non-overlapping 20 ms windows. We then found consecutive windows where the slope changed from a large negative value to a large positive values; we required the change in z-score to exceed 10 per second in each direction. The local minimum over this 40 ms interval was defined as the time of an ictal discharge. We found that this automatic identification process was consistent with visual inspection and identification of IDs by

an experienced electroencephalographer (Sydney S. Cash), although much faster (automated detection times of ~ 3 s per channel versus ~ 1 hour per channel manually).

A key challenge of this automated procedure is tuning the sensitivity of ID detections. In comparing window sizes, we found that slope is largely stable as the window size varies from 15-30 ms. Furthermore, we compared different choices of slope threshold (5-50 standard deviations per sec). We found that a threshold of 5 standard deviations per second resulted in many false positive detections, whereas most false positives vanished at a threshold of 10 standard deviations. Larger choices (> 30 standard deviations) reduced the number of true positives. We found that for some channels no choice of threshold produced a consistent detection of discharges; these channels were excluded from analysis (minimum of 4, maximum of 28, mean of 12.8 electrodes excluded). The average and standard deviation of the total number of electrodes analyzed for all patients and seizures is shown in Table II.2.

Table II.2. Total counts of identified ictal discharges (IDs), large amplitude discharges (LADs), and small amplitude discharges (SADs). Because electrode number and seizure duration varied between seizures (top two rows), we normalize counts of IDs, LADs, and SADs per electrode and per second. The mean (first column) and standard deviation (second column) are indicated for each measure.

Total counts (N=11 seizures)	Mean	Standard Deviation
Number of electrodes	45.18	14.18
Seizure duration (sec)	73.73	18.76

Total counts (N=11 seizures)	Mean	Standard Deviation
IDs per electrode	679.53	230.88
LADs per electrode	289.58	72.3
SADs per electrode	389.95	172.74
IDs per electrode per sec	8.92	2.3
LADs per electrode per sec	3.88	0.98
SADs per electrode per sec	5.04	1.59

We recorded each detected IDs' amplitude and left-hand slope. For each seizure, we clustered the set of IDs by estimating a multivariate Gaussian mixture model with two clusters, and assigned each ID to its most likely cluster. These two steps utilized the MATLAB functions *fitgmdist.m* and *cluster.m*, respectively (Mathworks, Natick, MA). We found that the clustering algorithm failed to converge for some patients and seizure when using a choice of 3 or 4 clusters. Using a choice of two clusters, the clustering algorithm converged for all eleven seizures, suggesting this is a natural number of ID types for these data. We label each discharge as “small amplitude discharge” (SAD) or “large amplitude discharge” (LAD), which denotes one feature of its associated cluster (the amplitude).




Figure II.1. Ictal discharges propagate in waves with multi-faceted structure.

(A) Example seizure with two categories of ictal discharges marked (black and green o's). Six time points are highlighted for further visualization below. (B) Detailed visualization of the brain voltage at the times shown in (A). Each colored grid illustrates the voltage across the microelectrode array at a moment in time, with IDs indicated as circles. Panels show voltages in three channels over a four second interval surrounding the indicated moment. (C) Example distribution of ictal discharge shapes. Shape is quantified by discharge slope (left) and amplitude (right). Empirical distributions (blue) of the shape show bimodal composition. Estimating the underlying mixture of the distributions (black and green) permits classification of each discharge as large or small amplitude.

2.3. Calculation of speed

We approximated the propagation speed of IDs in two steps. First, we identified spatial waves: rapid, spatially-continuous sequences of IDs. We defined a spatial wave as multiple discharges that (1) contain no pause greater than 2 ms; (2) propagate between spatially adjacent electrodes (namely, from one electrode to its eight nearest neighbors); (3) include at least 10 discharges total; and (4) persist for at least 3 ms from first to final discharge. Note that the third criterion is intentionally chosen to be low, leading to a high discovery rate; false positives are reduced by the fourth criterion. A high sensitivity is important for SAD waves especially, since these discharges tend to be smaller in amplitude. Finally, to avoid assigning a discharge to multiple waves, we combined any sequences with intersecting IDs. Example spatial waves are shown in Figure II.1B.

For each spatial wave, we estimated its speed as d/t where d is the straight-line distance between the first and last electrode in the wave, and t the time elapsed from first to last discharge. Table II.3 characterizes the spatial waves using several statistics, including the speed.

Table II.3. Features of spatial waves for LADs and SADs. The mean (first column) and standard deviation (second column) are indicated for each measure and both discharge types.

Totals (N=11 seizures)	Mean	Standard Deviation
<i>LADs</i>		

Totals (N=11 seizures)	Mean	Standard Deviation
Number of spatial waves per second	4.9	1.3
Spatial wave speed (mm / s)	278	131
Spatial wave duration (ms)	8.6	4.3
Electrodes per spatial wave	39	26
<i>SADs</i>		
Number of spatial waves per second	3.8	2.7
Spatial wave speed (mm / sec)	311	165
Spatial wave duration (ms)	6.0	2.3
Electrodes per spatial wave	30.6	21.0

2.4. Point process analysis

After identifying the discharges and summarizing their basic features, we analyze these point process sequences in two ways (Truccolo et al. 2005): (1) By computing auto- and cross-correlations across the channels, and (2) By estimating statistical models of within- and between-channel interactions. Here we define the mathematical notation and formulas used in the analyses applied here.

2.4.1. Correlation analysis. The sample correlation between channels i and j at lag K is,

$$\hat{\rho}_{i,j}(K) = \frac{1}{T} \sum_{k=1}^{T-K} (dn_i(k) - \bar{dn}_i)(dn_j(k+K) - \bar{dn}_j)$$

where T is the total duration of the recording, $dn_i(k)$ is the number of spikes at time index k of channel i , and \bar{dn}_i is the average firing rate at channel i . When $i = j$, this equation defines the autocorrelation; otherwise, this equation defines the cross-correlation. Peaks in the cross-correlation suggest physical relationships (Nowak et al. 1999; Brockwell and Davis 2002; Nowak and Bullier 2000). In the case of IDs, peaks in cross-correlation may be interpreted as excitatory functional connectivity between recording sites (Moore, 1970; Ostojic et al., 2009). Troughs, similarly, may be interpreted as inhibitory functional connectivity.

2.4.2. Point process generalized linear model. Like correlation analysis, the point process model quantifies the interactions between binary time series. While correlations summarize pairwise interactions, modeling can flexibly estimate higher-order interactions; auxiliary variables, which can confound correlations, are easily incorporated into a model. A point process model is defined by its conditional intensity function (CIF). The CIF is a vector that reflects the probability of events. Integrating the CIF over time gives the expected number of events over an interval. Following (Truccolo et al. 2005), we use the conditional intensity model:

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} dn_{i,k-\kappa} + \sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} dn_{c,k-\kappa}$$

Here $\exp \beta_0$ is the baseline discharge rate; $\sum_{\kappa=1}^Q \beta_{\kappa} dn_{i,k-\kappa}$ is the self-history dependence, i.e. how the recent past of channel i affects its current discharge odds; and $\sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} dn_{c,k-\kappa}$ is the ensemble-history dependence, which depends on the history of electrodes in C_i , channel i 's four nearest neighbors (see Figure II.2B). Throughout we label the self-history dependence “intrinsic effects” and the ensemble-history dependence “extrinsic effects.” These terms emphasize that intrinsic effects are factors from the same channel i , while extrinsic effects are due to external factors outside of channel i . Note that, in these sums, the intrinsic effects persist for Q time steps, the extrinsic effects for R time steps.

We utilize splines to parameterize the model's effects, which reduces the number of parameters to be estimated and enforces smoothness. The expression for the conditional intensity model becomes:

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i) + \sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} G_{\kappa}(dn_c)$$

The effects $\sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i)$ and $\sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} G_{\kappa}(dn_c)$ are now sums over spline basis functions. Weights β_{κ} and $\gamma_{c,\kappa}$ are then estimated from the data. Because the splines can

represent complicated functions with a small set of basis functions, the total number of parameters is substantially reduced. Specifically, here we reduce intrinsic effects from 10,000 to 13 parameters, and we reduced extrinsic effects from 800 to 16. In total there are 78 ($1+13+4*16$) parameters in the model.

Both correlations and models were estimated over 30 s moving windows (29 s overlap). For models we estimated parameters via an iteratively reweighted least squares algorithm (i.e. Newton-Raphson method), implemented in MATLAB as *glmfit.m* (McCullagh and Nelder 1989). All models were tested for goodness-of-fit using the procedure described in (Brown et al. 2002). Briefly, for each model we computed a Kolmogorov-Smirnov (KS) statistic. This KS statistic was tested against computed confidence bounds, and models with significantly high KS statistics (i.e. poor scores) were rejected. For the data analyzed here, only a small percentage of models were rejected: 2.5% and 1.2% of LAD and SAD models, respectively. To further mitigate uncertainty, for all visualizations of model effects (Figure II.3A-C, Figure II.4A-C, Figure II.5A-E) we also mask effects that are not significant at a level of $\alpha=0.05$. That is, for each model we first compute confidence bounds associated with the effects curves. We then identify lags where the confidence intervals fail to exclude zero and set the model effects at those lags equal to zero. Hence all remaining effects are verified as statistically significant, which emphasizes the models' most salient features.

2.5. Directional analysis

For each seizure, we computed cross-correlations and extrinsic effects estimates between every MEA electrode and its four nearest neighbors. These measures tended to vary between neighbors, consistent with the notion that IDs propagate in preferred directions. To further investigate the directionality, we implemented a two-step approach: (1) First, we defined the composite direction: a vector that summarizes the preferred direction for a set of cross-correlations or spatial effects. (2) Then we grouped all composite directions for each patient and summarized their circular distributions.

2.5.1. Composite direction. If a set of cross-correlations or spatial effects shows an imbalance across neighbors, it indicates discharge propagation is more likely in certain directions than others (see example in Figure II.2B). We therefore developed the composite direction: a simple measure to summarize directional preference in a set of cross-correlations or extrinsic effects. For a set of cross-correlations or extrinsic effects at each of a channel's four neighbors (in the cardinal directions, $c = \{N, S, E, W\}$), we average over lags spanning 1-10 ms and then use these four quantities to compute a vector sum. The resulting vector points in the direction IDs are most likely to propagate according to the cross-correlation or extrinsic effects.

2.5.2. Circular statistics. Having computed composite directions for all electrodes and time intervals in a given patient, we analyze their distribution. Assuming composite directions follow a von Mises distribution, we compute max-likelihood estimates,

including confidence intervals, for the mean direction θ_0 and concentration parameter κ . Confidence intervals for the mean direction appear as red arcs in Figure II.3D, II.4D, and II.5F. Confidence intervals for the concentration parameters are listed in Table II.4 and Table II.5. We also perform a Rayleigh test on the composite directions, which determines whether a set of angles is uniformly distributed across the circle. The alternative is that angles cluster about a preferred direction. Estimation and inference are implemented via the *CircStat* toolbox for MATLAB and verified using the R package '*circular*' (Berens, 2009).

In addition to approximating the propagation directions, we compared the directions for different measures and types of discharge. Specifically, we computed, for each channel and time window, (1) The difference in propagation angle for LADs estimated by cross-correlation versus by modeling, and (2) The difference in propagation angle for LADs versus SADs, both estimated using modeling. Similar angular directions computed using (1) different measures or (2) different discharge types produce angular differences near zero, while unrelated angular directions produce angular differences different from zero. We analyzed the angular differences in the same way as overall direction: the set of angular differences for each patient were used to estimate parameters for a von Mises distribution. Mean and concentration parameter estimates, including confidence intervals, appear in Table II.4 and Table II.5.

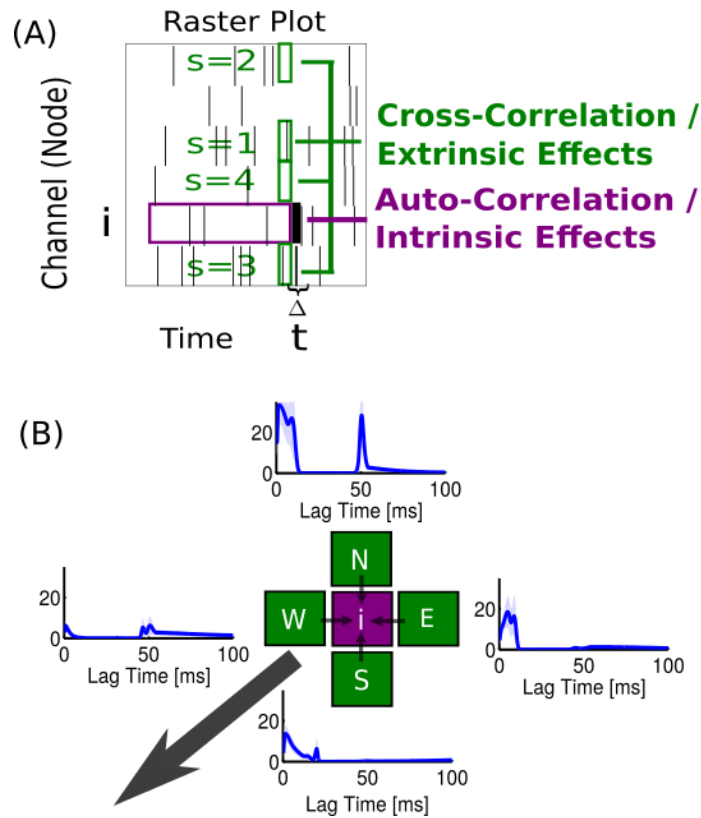


Figure II.2. Overview of analysis approach for estimation of spatiotemporal structure from point process data.

(A) Sequences of IDs in space and time are analyzed for within-channel (intrinsic or local circuit, purple) and between-channel (extrinsic or ensemble, green) dynamics using descriptive methods (auto- and cross-correlation) and statistical modeling. (B) Example of directional propagation. In the statistical model, IDs in four neighboring channels (green) modulate the discharge probability at channel i (purple). Each neighbor has a unique modulation profile (blue). In this example, IDs from the north and east neighbors are more influential than the south and west neighbors. The large arrow indicates the most likely direction of discharge propagation.

3. Results

We analyze local field potentials (LFP) recorded during seizures by microelectrode arrays (MEAs) implanted in four patients with focal epilepsy. All eleven seizures are spike-and-wave seizures that evolve with a similar electrographic pattern: fast low voltage activity at seizure onset followed by a transition to rhythmic bursting (Perucca, Dubeau, and Gotman 2014). We show an example of a representative seizure in Figure II.1. Visual inspection reveals that rapid voltage fluctuations (marked by green and black o's in Figure II.1A,B) occur throughout the seizure. These events are a characteristic signature (i.e., the “spike”) of the spike-and-wave seizure. We refer to these rapid voltage fluctuations as ictal discharges (IDs). We note that the IDs of interest here occur during seizure, and are temporally distinct from pre-ictal (Huberfeld et al. 2011) or interictal (Staley, Hellier, and Dudek 2005) discharges.

At a finer timescale, visual inspection reveals two properties of the IDs. First, the temporal features of the IDs (e.g., their rate of appearance or interval between IDs) appear to evolve during seizure. In some cases, the rate appears to decrease and become more rhythmic as seizures evolve (examples in Figure II.1A,B). Second, IDs appear organized in space. Visualization of the LFP at fixed moments in time reveals that IDs tend to appear over extended spatial intervals of the MEA, not at isolated spatial locations or randomly across the MEA (examples in Figure II.1B).

Visual inspection of the example LFP trace also suggests two different categories of IDs; namely, smaller amplitude IDs with smoother temporal profiles (green circles in Figure II.1A,B), and larger amplitude IDs with steeper temporal profiles (black circles in Figure II.1A,B). This distinction appears throughout the seizure, and becomes especially clear late in seizure, during which smaller amplitude IDs tend to follow larger amplitude IDs. To explore this distinction in more detail, we determine the amplitude and left-hand slope (i.e., steepness) associated with each ID. Analysis of the distributions of these quantities reveals multiple peaks, consistent with the notion of different ID categories (example for one patient and seizure in Figure II.1C). For each patient and seizure, we estimate two underlying distributions using a Gaussian mixture model (see Methods). From the estimated distributions, we assign each ID to its most likely group: either “large amplitude discharge” (LAD) or “small amplitude discharge” (SAD). Here we use the electrographic convention that discharges are negative deflections, thus units of amplitude and slope are in negative standard deviations. We note that inclusion of a third cluster in the analysis does not alter the results or provide additional meaningful categorization (not shown). We therefore conclude that two types of ictal discharges appear, distinguished by two measured features: voltage amplitude and slope. We note that the automated procedure to identify IDs allows for rapid and consistent cultivation of ID data from multiple electrodes and seizures, thereby permitting a high-throughput

analysis. Because we observe hundreds of IDs per seizure (see Table II.2), such automated analysis is necessary.

After extracting individual IDs, we next identified sequences of IDs that occurred across at least ten spatially-adjacent electrodes within a brief time interval (see Methods). We refer to such sequences as spatial waves. Visual inspection suggests that spatial waves are prominent and appear to propagate in a specific direction (examples in Figure II.1B). In Table II.3 we summarize features of the spatial waves computed for the LADs and SADs over the MEA. Overall, the spatial waves of LADs and SADs show comparable speeds, durations, and spatial extents; however, we note that SAD spatial waves are slightly faster, shorter, and involve fewer electrodes. Additionally, we observe fewer total SAD spatial waves. These results are consistent with the notion that both LADs and SADs possess similar spatial wave organization, although this organization is weaker for SADs than LADs.

To analyze further these IDs, we employed two analysis approaches: descriptive statistics and statistical modeling. Both approaches characterize the spatiotemporal structure of IDs by assessing the frequency of events (1) within channels and (2) across channels, as illustrated in Figure II.2A. We refer to the former (purple in Figure II.2A) as “intrinsic” effects, which represent the autocorrelation or self-history-dependence of the IDs within a channel. We refer to the latter (green in Figure II.2A) as “extrinsic” effects,

which represent the cross-correlation or ensemble-history dependence of the IDs between channels (Truccolo et al. 2005).

Visual inspection of the spatiotemporal organization of the IDs suggests that these discharges tend to propagate in preferred directions (examples in Figure II.1B; see [[Movie 1 and Movie 2]]). To assess this organization and compare the waves' directions, we developed a measure of directionality that can be estimated from the descriptive statistics and statistical models. We illustrate the application of this measure in Figure II.2B. Here, for a chosen interval of time (30 s) and a chosen electrode, the spatial effects are estimated (e.g., using the cross-correlations) between the chosen electrode and each of its four neighbors. These spatial effects (blue traces in Figure II.2B) are then combined to compute a vector sum (see Methods). The resulting vector indicates the direction of influence between the chosen electrode and its neighbors. In this example, chosen from the data shown in Figure II.1, visual inspection suggests IDs propagate southwest across this electrode at this time (see Figure II.1B and [[Movie 1 and Movie 2]]). Because the north and east neighbors' spatial effects dominate the south and west in Figure II.2B, the estimated directionality (black arrow) points southwest, as expected. In what follows, we examine for each patient and seizure how the directionality evolves across the MEA as the seizure progresses.

Having defined two categories of IDs, and introduced the analysis approach, we now examine the spatiotemporal evolution of the IDs in eleven seizures. We begin with a

descriptive statistical analysis of the large amplitude discharges (LADs, Figure II.3). The descriptive statistical analysis consists of two measures: (1) the autocorrelations, which measure the structure of LADs within a channel, and (2) the cross-correlations, which measure the structure of LADs between channels. We show in Figure II.3 these correlations for a single seizure (Figure II.3A-B) and across all eleven seizures (Figure II.3C-D). Correlations are computed over a moving window (size 30 s, horizontal axis) for each electrode and then averaged across electrodes. Horizontal bands of red (blue) indicate likely (unlikely) inter-discharge intervals that persist in time.




Figure II.3. Correlation analyses reveal multiple time scales and directionality in ID dynamics but suffer from confounding effects.

(A,B) Example (A) autocorrelations and (B) cross-correlations (vertical axis is lag time, correlation value in color and bar) over a moving window (horizontal axis) of LADs for a single seizure. In (B), each of the four neighboring electrodes (white arrows) are shown together with the autocorrelation over the same lag range. (C) Average windowed autocorrelation for all patients and seizures. (D) Rose plots indicating directionality values for each patient. Lines (red) indicate estimated mean direction with 95% confidence intervals for each patient.

The example in Figure II.3A shows two prominent features of rhythmic activity.

First, an interval of negative correlation (blue band) occurs at lags 0-50 ms throughout

seizure. This feature suggests a period in which subsequent LADs are unlikely immediately following a LAD. We may interpret this effect as an ID “refractory period”, although we note that this refractory period refers to the LADs, and is not related to the refractory period of an individual neuron. Second, intervals of positive correlation (red bands) occur at lags of 50-100 ms, near 500 ms, and near 1000 ms. These features suggest an increase of rhythmic bursting (i.e., intervals of rapid LADs separated by 0.5 s) approaching seizure termination. We note that this transition to bursting behavior is consistent with the example progression shown in Figure II.1B. The smallest time scale, 50-100 ms, is the intra-burst interval (i.e., the time between IDs within a burst), and the second time scale, 500 ms, is the inter-burst interval (i.e., the time between bursts). The 1000 ms time scale results from three sequential bursts separated by ~500 ms delays. These features, consistent with bursting activity, also appear in the population average results (Figure II.3C).

To assess the spatial structure of the LADs, we computed the cross-correlations between electrodes in the MEA. The example in Figure II.3B shows the average cross-correlations, here averaged across all electrodes (lags on vertical axis) over a moving window (horizontal axis), with correlation values indicated in color. We note in this example a large positive correlation at short lags (less than 10 ms) whose strength varies in the four directions throughout the seizure. This feature is consistent with the rapid propagation of LADs between neighboring electrodes; an electrode is more likely to

produce a LAD when one of its neighbors produces a LAD in the recent past. We also note that, in this example, the strength of the correlation varies depending upon the direction. Here the IDs in a given electrode tend to be more correlated with their neighbor to the east. This spatial organization is consistent with a planar wave that propagates from right to left across the MEA.

To assess the directionality of this propagation, we combined the cross-correlations in all four directions to compute a composite direction vector (see Methods and Figure II.2B). We show in Figure II.3D histograms and mean values of the angles of these direction vectors, grouped by patient. We find for all patients that the null hypothesis of the Rayleigh test is rejected ($p < 1e-6$). We conclude that, using the cross-correlation measure, LADs do not propagate uniformly but instead travel in a preferred direction that varies by patient.

We note that this correlation analysis also reveals an important confound: prominent features in the autocorrelation impact the cross correlation. This effect is clearly illustrated in Figure II.3B. We note that the strong cross correlation at lags 50-100 ms, apparent in each direction, matches the strong auto-correlations observed at the same lags (compare the top-right panel in Figure II.3B to the other panels in Figure II.3B). Although correlation structure at lags beyond 10 ms may identify additional spatial relationships in the data, these relationships are confounded by the rhythmic ID activity at

each electrode. As we show below, a statistical modeling approach mitigates this confounding effect.

Having described the spatiotemporal evolution of LADs using correlations, we now characterize their structure with a statistical model (more specifically, a point process generalized linear model; see Methods). Figure II.4 illustrates a set of models fit over the same LAD data as the correlations shown in Figure II.3. Here within-channel dynamics are characterized by “intrinsic effects” rather than autocorrelations, and between-channel dynamics, earlier measured by cross-correlations, are now characterized by “extrinsic effects.”

Model estimates for a single seizure are shown in Figure II.4A-B. Here the intrinsic and extrinsic effects are estimated simultaneously over windows that span the duration of the seizure (horizontal axis). Color indicates how LAD probability fluctuates (percent change) at specific lag times post-discharge (vertical axis). Plots are averages over the electrode array of all effects that reached statistical significance (i.e., $p < 0.05$, see Methods). Green indicates a value of 1, and suggests ID probability does not change significantly. Red (blue) indicates significantly enhanced (depressed) ID probability.




Figure II.4. Modeling analyses reveal multiple time scales and directionality in ID dynamics without the confounds of descriptive analyses.

(A,B) Example (A) intrinsic effects and (B) extrinsic effects (vertical axis is lag time, effect value in color bar) over a moving window (horizontal axis) of LADs for a single seizure. In (B), each of the four neighboring electrodes (white arrows) are shown together with the intrinsic effects over the same lag range (C) Average windowed intrinsic effects for all patients and seizures. (D) Rose plots indicating relative values of extrinsic effects in each direction. Lines (red) indicate estimated mean direction with 95% confidence intervals for each patient. All probability images are masked to show only statistically significant effects ($p < 0.05$).

Similar to the autocorrelation results, variations in the intrinsic effects (Figure II.4A) indicate the evolution of rhythmic activity during seizure. In this example, we observe a decreased probability of LADs occurring 100-400 ms after a previous LAD

(i.e., the probability of inter-discharge intervals between 100-400 ms is reduced, blue in Figure II.4A). We note that this reduction becomes more pronounced approaching seizure termination. In addition, we observe an increase in the probability of inter-discharge intervals near 50 ms and between 500-750 ms; the latter increase becomes more pronounced approaching seizure termination (orange in Figure II.4A). Together these results suggest that LADs tend to exhibit bursts, consisting of 2-3 LADs each separated by 50 ms, with 500-750 ms between burst onsets. The properties of these bursts evolve during seizure, such that the interval between bursts becomes more pronounced (blue region 100-400 ms). We also note that the 50 ms peak becomes reduced near seizure termination, suggesting the number of LADs per burst eventually decreases to one, consistent with visual inspection (Figure II.1B), as sequences of LADs become replaced by sequences of LADs and SADs. These signatures of rhythmic activity appear consistently across the population of patients and seizures (Figure II.4C).

An example of the extrinsic effect estimates for this patient and seizure are shown in Figure II.4B. These estimates, like the cross-correlation, measure between-channel structure. Unlike the two types of correlation, which are estimated separately and dominated by common signals (see Figure II.3B), the intrinsic and extrinsic effects are fit simultaneously in the statistical modeling framework. The example in Figure II.4B shows the extrinsic effects averaged across all electrodes (lags on vertical axis) over a moving window (horizontal axis), with color indicating the relative change in LAD probability.

We find, similar to the cross-correlation results, an increase in LAD probability at short lags (less than 10 ms), which persists in different neighbors throughout the seizure; again, this feature is consistent with the rapid propagation of LADs between neighboring electrodes. Unlike the correlation results, which are dominated by a common increase between 20-100 ms (Figure II.3B), we find here much more variability between directions in the extrinsic effects at intermediate lags (~20-100 ms). This variability was hidden in the cross correlation analysis due to the confounding effect of rhythmic ID activity at each electrode. The statistical model - which simultaneously estimates the intrinsic and extrinsic effects - mitigates this confound. The model results reveal less organized spatial influences at intermediate lags (~20-100 ms), compared to short lags.

We also assess whether the statistical models suggest directionality in LAD propagation. To do so we compute a measure of directionality similar to that used for cross-correlations, and examine the distributions of these directions. Namely, for each model we combine the spatial effects at short lags (< 10 ms) in all four directions and compute a composite direction vector (see Methods and Figure II.2B). The angles of the composite directions are summarized for all 11 seizures in Figure II.4D. The histograms indicate the distributions of angles for all channels and time windows, grouped by patient. Red arcs indicate the estimated mean direction with 95% confidence intervals.

To measure IDs' directionality - or lack thereof - deduced from the extrinsic effects, we again performed a Rayleigh test on the composite directions (see Methods).

For all four patients we rejected the null hypothesis that directions uniformly cover the circle ($p < 1e-6$). Additionally, we computed max-likelihood estimates of the concentration parameter for each patient's composite directions (Table II.4). For three of the four patients, the concentration parameters are near or above 1, which indicates their composite directions tend to align. We conclude from these results that LAD directionality at short timescales (less than 10 ms) concentrates in a specific direction and persists during seizure and across the MEA for three of the four patients.

Visual inspection of Figures II.3D and 4D suggests that the mean directionality is similar whether estimated by correlation (Figure II.3D) or statistical modeling (Figure II.4D). To investigate this further, for each electrode we compute the angular difference between the directionality estimated from the two approaches at each moment in time (see Methods). If the directionality estimates are similar for the two approaches, then we expect this difference to be near zero. Indeed, we find that the mean angular difference is small - less than 0.25 radians in magnitude - in three of the four patients (Table II.4, middle column). We note that, for the fourth patient (Patient D) the confidence intervals for the angular difference are large and contain 0. These results confirm our visual inspection; the directionalities estimated through the two approaches are qualitatively similar. However, we note that the angular differences for two patients (Patients A and B) have confidence intervals that do not include zero, and that possess large concentration

parameters (Table II.4, right column). These results suggest that, though small, significant angular discrepancies do exist between the two analysis approaches.

Table II.4. Summary of distributions of LAD directionality. Results in the first column are derived from the statistical modeling approach, while results in the second and third column compare the statistical modeling and descriptive approach. Listed are 95% confidence intervals for each estimate.

Patient	Concentration parameter of LAD directionality	Angular difference between LAD estimates from correlation and modeling	Concentration parameter of angular difference between LAD estimates from correlation and modeling
A	(0.97, 1.07)	(-0.17, -0.05)	(0.73, 0.82)
B	(2.42, 2.62)	(-0.24, -0.19)	(1.8, 1.93)
C	(0.91, 0.99)	(-0.01, 0.15)	(0.39, 0.47)
D	(0.11, 0.25)	(-1.67, 0.69)	(0.01, 0.14)

We conclude this analysis by applying the same modeling approach to analyze the spatiotemporal structure of small amplitude discharges (SADs). To start, we show in Figures II.5A, 5B, and 5E examples of the model estimates for a single seizure. In this case, there are two types of intrinsic effects (i.e., within-channel dynamics) that represent the influence of either a previous SAD (Figure II.5A) or a previous LAD (Figure II.5B) on the probability of a subsequent SAD. In this example, both types of intrinsic effects

suggest that IDs occur in rhythmic bursts. Visual inspection of Figure II.5B suggests that LADs consistently precede SADs by 50-100 ms, and that this effect becomes more prominent approaching seizure termination; notice the interval of increased probability (dark red) at these lags in Figure II.5B. We observe a similar - although weaker effect - of a preceding SAD in Figure II.5A. Combined these intrinsic effects are consistent with a three-discharge LAD-SAD-SAD burst, with intra-burst interval of 50-100 ms.

Approaching seizure termination, we observe an additional interval of increased SAD probability between lags 500-1000 ms following both a LAD and SAD. Notice that the effect of the preceding LAD is both stronger and earlier (i.e., at shorter lag) than the effect of the preceding SAD; these effects are again consistent with a LAD-SAD-SAD burst, with inter-burst interval of approximately 500 ms. This pattern of LAD/SAD dynamics remains evident in the population averages (Figure II.5C, II.5D). While the time between bursts is less consistent (i.e. lighter color) across patients, the tendency of SADs to follow LADs by 50-100 ms, particularly approaching seizure termination, remains strong.

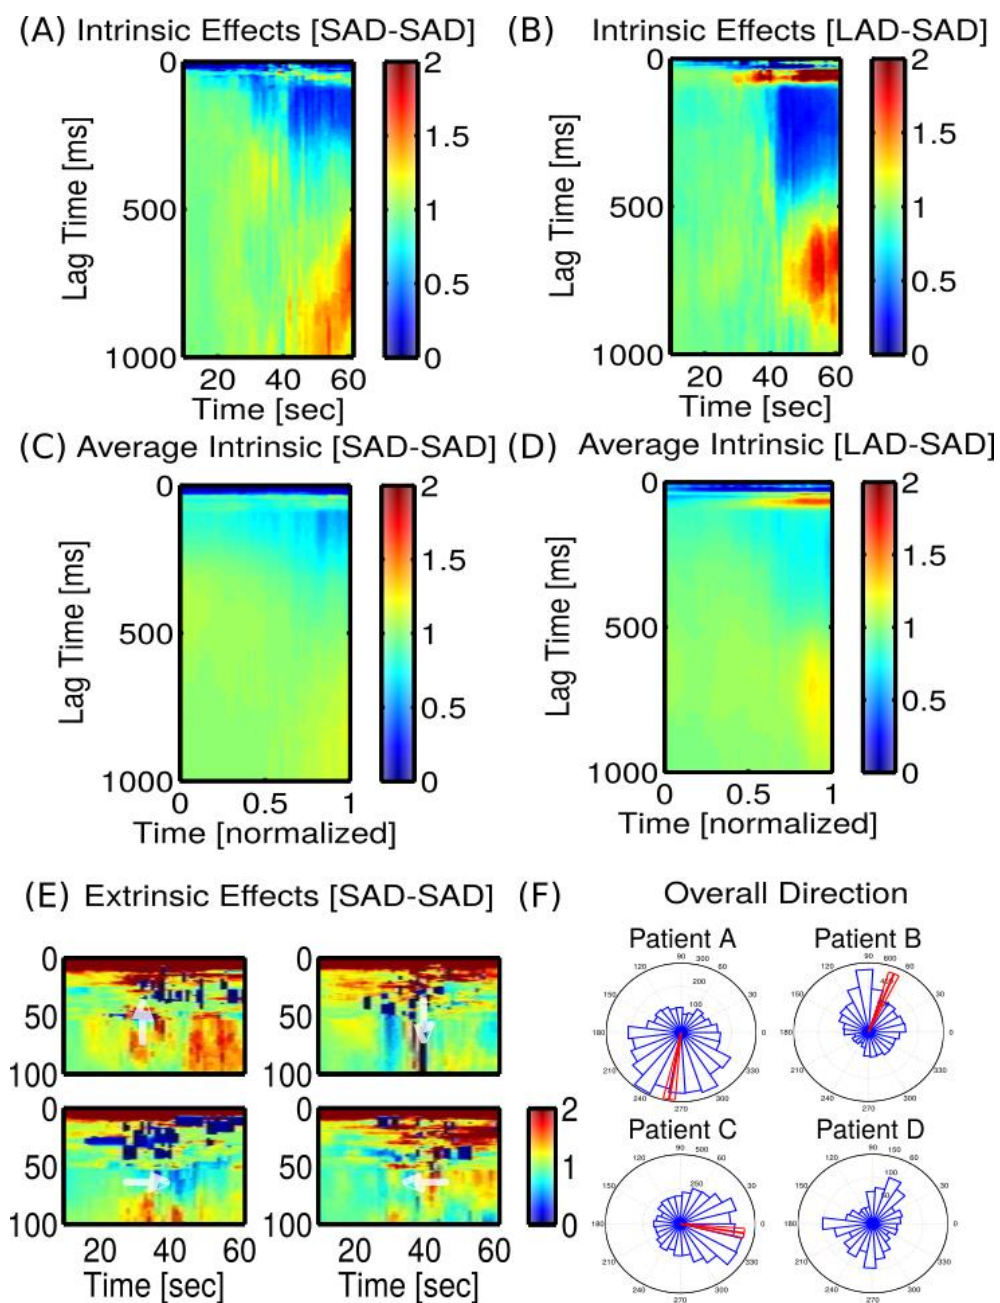


Figure II.5. Modeling analysis of small amplitude discharges (SADs) reveals effects that depend on both SADs and large amplitude discharges (LADs).

(A, B) Example intrinsic effects (lag on vertical axis, relative probability in color) over a moving window (horizontal axis) for (A) SAD-to-SADs and (B) LAD-to-SADs for a single seizure. (C, D) Average intrinsic

effects over a moving window for (C) SAD-to-SADs and (D) LAD-to-SADs for all eleven seizures. (E) Example extrinsic effects (lag on vertical axis, relative probability in color) for each of the four neighboring electrodes (white arrow) for SADs for a single seizure. (F) Rose plots indicating relative values of SAD-to-SAD extrinsic effects in each direction. Lines (red) indicate estimated mean direction with 95% confidence intervals for each patient. All probability images are masked to show only statistically significant effects ($p < 0.05$).

To examine the spatial influence of neighboring SADs, we show an example of the statistical model extrinsic effects in Figure II.5E. We note that, in this model, the extrinsic effects characterize the influence of SADs at four neighboring electrodes on the probability of a SAD at each center electrode (see Methods, and Figure II.2). In this example, we find the most prominent effect at short lags (1-10 ms), consistent with the rapid spatial propagation of SADs. At longer lags, from 10-100 ms, the spatial effects become less uniform and more difficult to characterize. To assess the directionality of SADs for all patients and seizure, we again compute a composite direction vector from each model's extrinsic effects (see Methods). Histograms of the direction vectors - computed for all time windows, MEA electrodes, and seizures - are shown for each patient in Figure II.5F. Like the population results for LADs (Figure II.4D), we find here that SADs tend to propagate in consistent directions for each patient.

Figure II.4D and Figure II.5F suggest that the mean angular directions of LADs and SADs are consistent. However, visual inspection of these figures reveals that the standard error of the angular mean (red arcs in Figures II.4D and 5F) tends to be larger for SADs compared to LADs, consistent with the directionality of SAD propagation being more varied. To test this hypothesis, we compared the concentration parameters for

the two sets of directions (see Methods). We find that the concentration parameters for SADs are significantly less than the concentration parameters of LADs in three of four patients (compare first and second columns of Table II.5). For the remaining case, Patient D, the confidence intervals for LAD and SAD concentration parameter overlap, but only slightly. The distribution of angular differences between LAD and SAD directions is near zero (third column of Table II.5) and tends to be concentrated (fourth column of Table II.5), although less concentrated than the LADs alone in three of four patients (compare the first and last columns of Table II.5). We conclude that SADs do exhibit spatial organization similar to LADs, but less of it, and that LADs and SADs tend to propagate in similar directions.

Table II.5. Summary of differences in directionality between large and small amplitude discharges. Estimates of the concentration parameters of LADs (first column) and SADs (second column), as well as their mean angular differences (third column) and the concentration parameter for these angular differences (fourth column). Each entry indicates 95% confidence intervals, see Methods. The first column of Table II.5 is identical to Table II.4, but repeated here for completeness.

Patient	Concentration parameter of composite direction (LAD model)	Concentration parameter of composite direction (SAD model)	Mean angle difference (LAD - SAD)	Concentration of angular difference (LAD - SAD)
A	(0.97, 1.07)	(0.47, 0.56)	(-0.27, -0.13)	(0.57, 0.66)
B	(2.42, 2.62)	(0.4, 0.47)	(0.17, 0.34)	(0.38, 0.45)

Patient	Concentration parameter of composite direction (LAD model)	Concentration parameter of composite direction (SAD model)	Mean angle difference (LAD - SAD)	Concentration of angular difference (LAD - SAD)
C	(0.91, 0.99)	(0.41, 0.48)	(-0.12, 0.06)	(0.36, 0.44)
D	(0.11, 0.25)	(0, 0.13)	(-0.07, 0.44)	(0.21, 0.35)

4. Discussion

In this chapter, we analyzed ictal discharges (IDs) from microelectrode array recordings performed in four patients during seizure. We showed that IDs can be separated into two distinct categories by their shape: large amplitude discharges (LADs) and small amplitude discharges (SADs). Correlation and modeling analysis of LADs and SADs showed that both exhibit rhythmic bursts and traveling waves, which become more prominent approaching seizure termination. We showed that cross correlation analyses are confounded by autocorrelation effects, and addressed this confound in a statistical modeling framework, in which rhythmic and spatial effects were simultaneously and separately estimated. Finally, testing for differences in the propagation direction of LADs and SADs revealed that both travel in similar directions, although SADs tended to do so with less consistency. Together these results provide the first patient-specific statistical

characterization of the spatiotemporal dynamics of ictal discharges at the sub-millimeter spatial scale.

Voltage discharges in human epilepsy have been extensively studied in human and animal models. Interictal discharges (IIDs) - which occur between seizures - are common in epilepsy, although the relationships between interictal discharges and seizures are not completely understood (Staley et al., 2005). Like the ictal discharges studied here, IIDs have been shown to propagate over cortex (Emerson et al. 1995), with substantial variability in spatial propagation between sequential IIDs (Sabolek et al. 2012). While the clinical significance of IIDs is unclear, it has been hypothesized that regions where the earliest discharges occur are located in the epileptogenic zone and behave as pacemakers that drive propagation; targeting resections to brain regions that lead sequences of IIDs correlates with successful surgical outcome (Alarcon et al. 1997). Here, the microelectrode arrays were not placed in the epileptogenic zone, so that the traveling waves of IDs - observed during seizure - propagated from another brain region. Unlike IIDs, the LADs analyzed here displayed consistent directions of propagation within a seizure, and between a patient's seizures. This suggests that the direction of LAD propagation may help target the brain region from which the earliest IDs emanate. When the direction of LAD propagation is consistent, following the propagation backwards through established neural pathways may help identify the source, which could serve as a

target for resection. When LAD propagation is unclear or originates from multiple regions, we expect surgical outcome will be worse (Alarcon et al. 1997).

Recent work suggests rhythmic trains of IDs follow passage of an ictal wavefront (Smith et al. 2016). These IDs are proposed to propagate as traveling waves away from the ictal wavefront, and are associated with transient bursts of multiunit firing and high gamma field activity (Smith et al. 2016). Similar to the results in (Smith et al. 2016), we find that IDs - both LADs and SADs - propagate across the MEA as traveling waves with a preferred direction for each patient, and that intervals between discharges or bursts of discharges increase approaching seizure termination (i.e., the activity “slows”). Unlike (Smith et al. 2016), we find that IDs tend to become more regular approaching seizure termination, and we do not find strong evidence for two, opposing directions of propagation. This difference may result from the distinction here between two ID categories; late in seizure, the LADs may represent traveling waves that originate from the ictal wavefront, while the SADs may originate more locally and therefore exhibit more heterogeneity not directly yoked to the ictal wavefront. Combining the LADs and SADs may then reduce the regularity of IDs approaching seizure termination. In this scenario, we expect the ictal wavefront migrates in a constant direction away from the MEA, so that IDs continue to propagate from the wavefront to the multielectrode array at the same orientation throughout the late seizure interval.

How seizures terminate remains unknown (Lado and Moshé 2008). Here we observed increased temporal organization of IDs approaching seizure termination. This is consistent with the observation of increased coupling approaching seizure termination, common in invasive brain voltage recordings (Kramer and Cash 2012; Schindler et al. 2007; Topolnik et al., 2003; Schindler et al. 2008; Schiff et al. 2005). More specifically, we observed the emergence of discharge patterns approaching seizure termination that commonly consisted of LAD-SAD doublets, and LAD-SAD-SAD triplets. The mechanisms of the LADs and SADs observed here are not understood. We speculate that long range synaptic connections support LAD propagation from a possibly distant source, while mechanisms that support SADs may sustain more local propagation. For example, an ID emanating from a distant cortical source may follow excitatory synaptic pathways across the cortex, eventually reaching the observed cortical patch and propagating across this patch as a LAD. This LAD is then followed by one or two SADs, whose propagation over the cortical patch is sculpted by local mechanisms (e.g., local connectivity structure or local interneuron networks (Sabolek et al. 2012)). In this scenario, the seizure terminates when the wave source (or perhaps the ictal wavefront [Smith et al. 2016]) dissipates and ceases emitting traveling waves of IDs.

These proposed mechanisms for LADs and SADs are consistent with a recent analysis of human tissue slices from patients with mesial temporal lobe epilepsy (Huberfeld et al. 2011). In that work, two classes of discharges were identified, pre-ictal

and interictal, which occurred contemporaneously (Huberfeld et al. 2011). The pre-ictal discharges possessed larger amplitudes, like the LADs studied here, depended on glutamatergic signaling and were preceded by pyramidal cell firing. The interictal discharges had smaller amplitudes, like the SADs studied here, depended on both glutamatergic and depolarizing GABAergic signaling, and were preceded by interneuron firing. In addition, GABAergic networks have been shown as major contributors to the variance in discharge propagation patterns (Sabolek et al. 2012), consistent with the increased variability of SAD propagation compared to LAD propagation shown here. Although the mechanisms of IDs remain unknown, this work is consistent with existing observations that LADs and SADs utilize different neural networks.

We applied two types of analyses - descriptive analysis and statistical modeling - to these data. Both approaches revealed ID dynamics with similar spatiotemporal features (e.g., time scale and direction of propagation). However, the cross-correlation exhibited a clear confound: the appearance of autocorrelation (i.e., self-history) effects. Analysis of coupling in point process data often exhibits these types of confounds. For example, the spike-field coherence depends on the firing rate, such that lower firing rates reduce the coherence (Lepage et al., 2011). Similarly, correlations between spike trains depend on many factors, including the firing rate, which can affect and bias estimates (Cohen and Kohn 2011; de la Rocha et al. 2007; Dorn and Ringach 2003). Procedures to mitigate these confounds include corrections to descriptive measures (Aoi et al. 2015; Cohen and

Kohn 2011) and the development of multivariate point process models (Granger 1969, Lepage et al. 2013). Point process models have been used to characterize action potentials from individual neurons in a variety of contexts (Truccolo et al. 2005; Huang et al. 2009; Pillow et al. 2008; Latimer et al. 2015; Gerstner). Here we developed a point process modeling framework to analyze cortical voltage events generated by neural populations. In doing so, we utilized a predictor unique to these spatiotemporal data: the past discharge activity of spatial neighbors. This type of ensemble predictor is rarely available for individual neurons, for which the identification of a neighbor depends on complex - and typically unobserved - synaptic connectivity.

The results presented in this chapter suggest three areas for future investigation. First, additional research may identify an optimal method of discharge identification that depends on additional features developed for individual patients and electrodes. Second, different categories of discharge - beyond the two identified here - may exist. Finally, we estimated point process models for a particular model class that includes the effects of self-history and ensemble-history. These models were effective in capturing the spatiotemporal patterns observed in the data; more than 97% of the models passed a standard goodness-of-fit test. However, further model development may result in more accurate or simpler (lower-dimensional) model formulations. This may provide additional insight into the spatiotemporal dynamics and physiology of IDs.

How ictal discharges initiate and spread over human cortex is not understood. In this chapter, we developed a quantitative analysis of the spatiotemporal dynamics of human IDs at the sub-millimeter scale. We showed that two categories of IDs exist with similar temporal organization, but different spatial organization: large amplitude discharges propagate with more spatial organization, while smaller amplitude discharges propagate with less spatial organization. Understanding the mechanisms that support - and disrupt - these different discharge dynamics may help target interventions tailored to each discharge type.

CHAPTER III. POINT PROCESS MODELING REVEALS A UNIQUE TYPE OF ENHANCEMENT DURING HUMAN SEIZURES

1. Introduction

Characterizing how a network's information is distributed among multiple spike trains is a key problem in neuroscience. Sensory systems, for example, have been described in terms of how well they retain information from stimuli (Rolls 2003; Quiñero and Panzeri 2009; Dimitrov et al., 2011). However, details of the encoding and how information flows among spike trains and beyond early sensory areas remain incompletely understood.

In analyzing the GLMs discussed in Chapter II, we observed evidence of enhancement, a curious and little-reported statistical phenomenon related to information theory. Typically, we expect that if two correlated variables are used to predict a modeled variable, then the two variables will be at least partially redundant in their explanation. In fact, some statistics textbooks state – incorrectly – that the variability explained by a second model covariate is always less than the information that same covariate provides alone (Currie and Korabinski 1984). However, this is not necessarily true. There are a limited number of cases in social sciences, fewer still in neuroscience, where the converse occurs (Gat and Tishby 1998; Schneidman et al., 2003; Narayanan et al., 2005). In these situations, information behaves super-linearly: two variables provide more information in a joint model than they provide in total individually. In the statistical

literature, this phenomenon has been variously called enhancement (Friedman and Wall 2005), synergy (Hamilton 1988), and suppression (Lynn 2003). Following Friedman and Wall (2005) we use “enhancement” to mean super-linear information between model factors and “suppression” referring to changes in model effect sizes.

In this chapter, we define an enhancement score for a set of GLMs similar to the mutual information computed in previous neuroscience applications. We demonstrate how to estimate the score, including confidence bounds, by computing it for simulated data. Our results show that the estimated confidence is reliable across several networks. Applying this technique to *in vivo* neural data, we find that enhancement is ubiquitous during human seizures, while redundancy tends to be the norm in mouse auditory networks. We conclude by discussing the physiological implications of enhancement and redundancy, particularly for seizure.

2. Methods

We analyze enhancement, or lack thereof, in three neural data sets: simulated spiking from a simple two-cell network, population-level discharges recorded during human seizure, and multi-unit activity recorded from mouse auditory and prefrontal cortex.

2.1. Data

Here we describe the procedures undertaken to simulate and record spike train data. All analysis and visualization was performed using custom MATLAB software (Mathworks; Natick, MA).

2.1.1. Simulated two-cell network. The first data set we analyzed was simulated from a two-cell spiking network, shown schematically in Figure III.2A. Cell X and Cell Y are described by two point process generalized linear models (GLM) where the quantities $\lambda_{X,k}$ and $\lambda_{Y,k}$ give the expected number of spikes in time bin k for the respective cells. These spike probabilities vary depending on intrinsic effects (self-history-dependence, filled circles in Figure III.2A) and extrinsic effects (ensemble-history-dependence, open circles in Figure II.2A). Such effects are similar to the ones used in previous functional network simulations (Okatan et al., 2005; Kim et al. 2011) and discussed in greater detail in Chapter II Methods 2.4.2.

Mathematically, these spiking models can be described by the conditional intensity functions (CIF)

$$\log \lambda_{X,k} = w_{X,0} - w_1 \phi(k - k_X^*) + w_2 \phi(k - k_Y^*)$$

$$\log \lambda_{Y,k} = w_{Y,0} - w_3 \phi(k - k_Y^*) + w_4 \phi(k - k_X^*)$$

The terms $w_{X,0}$ and $w_{Y,0}$ are baseline spike rates; they are set to $\log 0.02$ and $\log 0.01$, which correspond to rates of 20 and 10 Hz, respectively. Terms $w_1 \phi(k - k_X^*)$ and

$w_3\phi(k - k_Y^*)$ are intrinsic effects for Cells X and Y, respectively; terms $w_2\phi(k - k_Y^*)$ and $w_4\phi(k - k_X^*)$ are extrinsic effects describing the drive from Y to X and X to Y, respectively. k_X^* and k_Y^* represent the times that Cells X and Y last spiked. Each cell is affected by both spike times; the function ϕ describes the time course of the effects. For simplicity, we set all effects to have the same time course $\phi(\tau) = \exp(-\frac{\tau}{50})$, plotted in Figure III.2.A.

To explore the system's behavior, we fixed w_1 and w_3 equal to one and let w_2 and w_4 vary from -5 to 5. For every (w_2, w_4) pair, we generated 1000 spike trains for the network, each of duration 20 seconds. From each pair of spike trains, we then estimated GLM parameters and enhancement scores using the procedures outlined in Chapter II, Methods 2.4.2. Median enhancement scores across the 1000 trials are summarized in Figure III.2.B over the range of w_2 and w_4 values. We define the enhancement score in Section 2.2.4 after we have introduced the data and form of the statistical model.

2.1.2. Human seizure networks. We also considered data from microelectrode array (MEA) recordings of eleven human seizures. For each seizure, we considered a set of ictal discharges, rapid voltage fluctuations with large amplitude. These discharges are identical to those analyzed in detail in Chapter II. Patient history, recording protocols, and discharge definitions are described in detail in Methods II.2.1. For each seizure, we

selected a 30-second window near the seizure's temporal midpoint. We analyzed the point process data for this window across all MEA electrodes, fitting 388 models in total.

2.1.3. Mouse cortical networks. The third data set we considered was multi-unit activity recorded from mouse cortex. All procedures involving animals were approved by the Boston University Institutional Animal Care and Use Committee (IACUC). These data comprise a subset of the data analyzed in (James 2016).

Ten ChAT-ChR2 transgenic mice were used, each three to six months old at the time of recording. Mice were surgically implanted with a head plate custom head plate designed to allow access to the prefrontal cortex (PFC) and auditory cortex (AC). Upon recovering from surgery, mice were habituated to a sound attenuation chamber where recordings took place. For recording, two small craniotomies were performed over PFC and AC and linear electrode probes were inserted into both areas. The electrode arrays were (1) a 16-contact linear probe with 100 μm spacing between electrode contacts (Neuronexus, Ann Arbor, MI; model A1x16-10mm-100-177-A16) inserted into PFC and (2) a 32-channel probe with 4 shanks each 400 μm apart and with 8 contacts spaced by 100 μm (Neuronexus, Ann Arbor, MI; model A4x8-5mm-100-400-177-A32) inserted into AC. To record spikes, signals were digitized at 24,414 Hz and bandpass filtered between 300-5000 Hz. Spikes were identified by threshold crossings, which were manually set at the beginning of each recording session. To identify individual neurons,

spike snippets were analyzed using Offline Sorter (Plexon Inc.; Dallas, TX) and manually clustered in principal component space.

Neural spiking was recorded during passive presentation of an auditory stimulus. Specifically, the stimulus was a 500 millisecond burst of white noise generated with a RZ2 Bioamp processor and RP2.1 real time processor (Tucker Davis Technologies) and digitized at a frequency of 48,828 Hz. Stimuli were presented to each mouse 100 times per recording block with an inter-trial interval of 10 seconds. Interspersed within the auditory stimulus trials are additional trials with optogenetic stimulation, however we do not consider these trials here.

2.2. Point Process Modeling

The point process generalized linear model (GLM) is a powerful and flexible way to characterize neural spike data. In the present chapter, we measured mutual information between spike trains using the GLM to compute enhancement. Here we describe the specific GLMs estimated for the three neural data sets; we define the enhancement score and its calculation from the GLM in Section 2.2.4. Point process models are described in more detail in Chapter II and, for example, by Truccolo et al. (2005). In all of the subsections below, we assume that time is discretized into bins of equal size Δ so that the k -th time point is $t_k := k\Delta$ and spikes are represented as binary time series $dn_{i,k}$ equal to 1 if a spike occurs in electrode channel i in the time interval $[t_k, t_{k+1})$ and 0 otherwise.

For simulated spike trains and multi-unit activity from mouse cortex, we set $\Delta = 1 \text{ ms}$; for ictal discharges, we set $\Delta = 0.27 \text{ ms}$.

2.2.1. GLM for simulated spikes. To analyze the spikes produced in the simulated network, we estimated parameters for the same conditional intensity that used to generate the data:

$$\log \lambda_{y,k}^{\wedge} = w_{y,0}^{\wedge} - w_3^{\wedge} \phi(k - k_y^*) + w_4^{\wedge} \phi(k - k_x^*) \text{ (Equation 3.1)}$$

where $\lambda_{y,k}^{\wedge}$ is the estimated CIF of cell y at time bin k , $w_{y,0}^{\wedge}$ is the estimated log baseline spike rate, w_3^{\wedge} is the estimated weight of the self-history dependence (intrinsic effects), and w_4^{\wedge} is the estimated log weight of the ensemble-history dependence (extrinsic effects).

We note that we only estimate parameters for the simulated neuron Y , while we observe the spiking activity of simulated neurons X and Y .

In this way, there is no misspecification between the GLM and the model that produces the data. Such a scenario is unlikely for *in vivo* data, but utilized here for simplicity and to clearly illustrate the enhancement phenomenon.

2.2.2. GLM for ictal discharges. The model used for ictal discharges was the same GLM as described in Chapter II; we briefly review this model here. The spike probability in channel i is modeled as

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} dn_{i,k-\kappa} + \sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} dn_{c,k-\kappa}$$

For this model $\exp \beta_0$ is the baseline discharge rate; $\sum_{\kappa=1}^Q \beta_{\kappa} dn_{i,k-\kappa}$ is the self-history dependence, i.e. how the recent past of channel i affects its current discharge odds; and $\sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} dn_{c,k-\kappa}$ is the ensemble-history dependence, which depends on the history of electrodes in C_i , channel i 's four nearest neighbors (see Figure II.2B). Furthermore, we utilize splines to parameterize the model's effects, which reduces the number of parameters to be estimated and enforces smoothness. The expression for the conditional intensity model becomes:

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i) + \sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} G_{\kappa}(dn_c)$$

The effects $\sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i)$ and $\sum_{c \in C_i} \sum_{\kappa=1}^R \gamma_{c,\kappa} G_{\kappa}(dn_c)$ are now sums over spline basis functions. Weights β_{κ} and $\gamma_{c,\kappa}$ are then estimated from the data. Because the splines can represent complicated functions with a small set of basis functions, the total number of parameters is substantially reduced. Specifically, here we reduce intrinsic effects from 10,000 to 13 parameters, and we reduced extrinsic effects from 800 to 16. In total there are 78 ($1+13+4*16$) parameters in the model.

2.2.3. GLM for multi-unit spiking in mouse auditory networks. The GLM used for multi-unit spiking was similar to the one used for seizure, as defined in Section 2.2.3. The full model for channel i , before reducing dimensionality, was

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} dn_{i,k-\kappa} + \sum_{i=1}^2 \sum_{\kappa=1}^R \gamma_{i,\kappa} \sum_{c \in C_i} dn_{c,k-\kappa}$$

Here the intrinsic effects are the same as for the ID model. Extrinsic effects again sum the modulation from other spike trains; a notable difference, though, is that the extrinsic effects in this model sum spiking activity from *all* other neurons C_i across auditory cortex ($i=1$) and prefrontal cortex ($i=2$), not just a subset in adjacent electrodes.

Just as with the seizure GLM, we replace the full effects, which involve many parameters, with a spline-basis representation of the effects that reduces the dimensionality and imposes smoothness. The spline-based representation takes the form

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i) + \sum_{i=1}^2 \sum_{\kappa=1}^R \gamma_{i,\kappa} G_{\kappa}(\sum_{c \in C_i} dn_c).$$

This change reduces the intrinsic effects from 2,400 to 6 parameters and extrinsic effects from 720 parameters to 5. In total there are 17 ($1+6+2*5$) parameters in the model.

2.2.4. Enhancement score. To measure enhancement, we first estimated the models described above for each type of data. Model parameters were estimated via an iteratively reweighted least squares algorithm (i.e. Newton-Raphson method), implemented in MATLAB as *glmfit.m*. For an estimated model, deviance was computed in the standard way as the Kullback-Liebler divergence (difference in log-likelihood) between the estimated model and the saturated model where each spike is described by a dedicated

parameter (McCullagh and Nelder 1989). Deviances were compared across the model hierarchy as summarized schematically in Figure III.1.

Let Δ_0 be the deviance of the null model with only one parameter; Δ_1 be the deviance of a model with intrinsic effects; Δ_2 be the deviance of a model with extrinsic effects; and let Δ_3 be the deviance of a joint model with both intrinsic and extrinsic effects. Then the quantity $\Delta_{0x} = \Delta_0 - \Delta_x$ represents the information gain (i.e. reduction in deviance) associated with a particular model $x = \{1,2,3\}$. Other authors have measured enhancement, sometimes called synergy, to occur when the information gain from a pair of variables Δ_{03} exceeds the sum of information gains they yield individually $\Delta_{01} + \Delta_{02}$ (Gawne and Richmond 1993; Gat and Tishby 1998; Optican et al., 1991; Nakahara and Amari 2002; Schneidman et al., 2003; Quinn et al. 2011). We adopt a similar definition. However, rather than measuring enhancement as the absolute difference between these two quantities, as in other cases, we normalize by the total information gain of the joint model. Therefore we define for the model hierarchy the enhancement score E :

$$E = [\Delta_{03} - (\Delta_{01} + \Delta_{02})]/\Delta_{03} = 1 - (\Delta_{01} + \Delta_{02})/\Delta_{03}.$$

The advantage of a normalized enhancement is that it is easier to compare across multiple models. Because deviance is easily computed even for more complicated GLMs, this enhancement score readily generalizes beyond the example models studied here.

2.2.5. Confidence intervals for the enhancement score. To compute confidence intervals for the enhancement score using a single fixed set of spike train data, we implement a resampling technique (Hastie, Tibshirani, and Friedman 2001; Sarma et al. 2011). Consider the case of a single enhancement score where y is the target spike train being modeled and X is the design matrix of the joint model, both with T rows which correspond to the T observations. To bootstrap more scores from the data, we implemented the following procedure: (1) Resample the rows of y and X with replacement; (2) Fit a hierarchy of models from this new data sample; and (3) Compute the enhancement score. We repeated steps (1)-(3) from 100-1000 times per model, depending on the computational load required. For the seizure data, where the arrays are large and estimating a model hierarchy takes approximately 20 minutes, we performed 100 iterations. For other data, where estimating a model hierarchy takes a matter of seconds, we were able to perform 1000 iterations without issue. Confidence intervals were then taken to be the 5th and 95th percentiles of the bootstrapped score distribution.

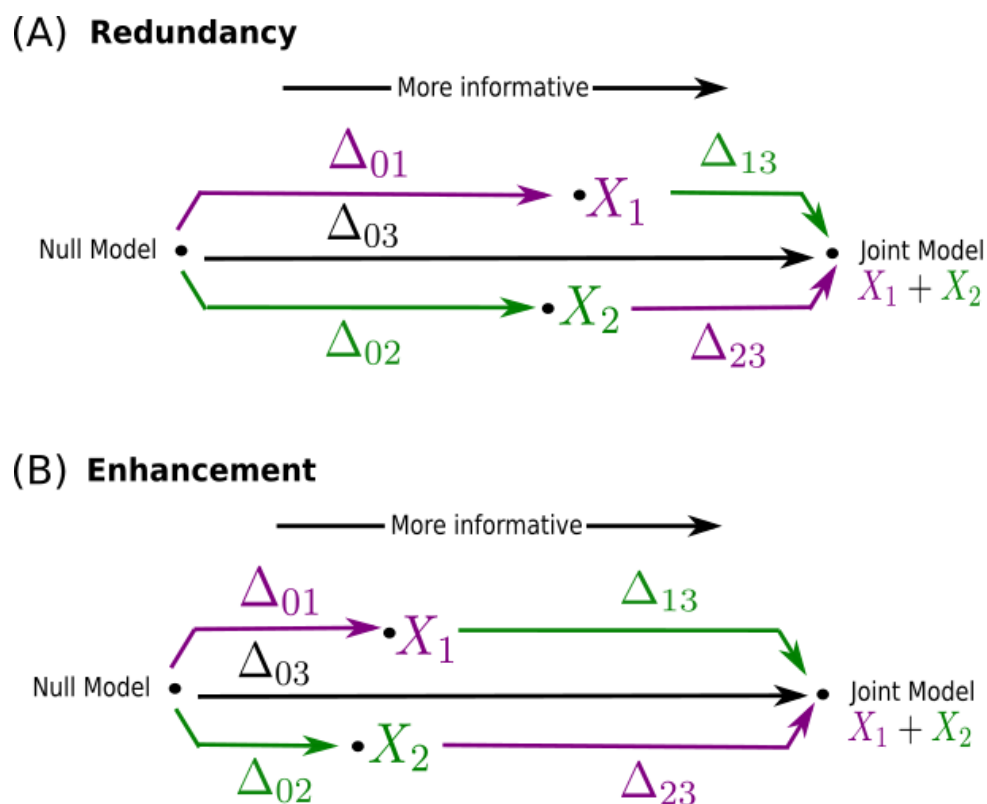


Figure III.1. Enhancement is measured by analyzing the model hierarchy.

(A) Schematic of model hierarchy with redundancy. In this figure, the models progress from least information (the null model, left) to most informative (the joint model, right). Feature X_1 or X_2 improves the null model considerably, but each provides subsequently less improvement in higher-dimensional models. These improvements are indicated schematically by the length of the horizontal line associated with each model; Δ_{01} is the information gain due to intrinsic effects alone, Δ_{02} is the information gain due to extrinsic effects alone, Δ_{03} is the information gain from intrinsic and extrinsic effects together. Similarly Δ_{13} and Δ_{12} represent the information gain extrinsic effects provide over intrinsic effects alone, and vice versa. In this figure, purple indicates a change in information where intrinsic effects are added, and green indicates a change in information where extrinsic effects are added. (B) Model hierarchy with enhancement. Feature X_1 or X_2 improves the null model modestly. However, both X_1 and X_2 improve one another's models more so than their individual improvements of the null model.

3. Results

We analyzed the enhancement, or lack thereof, in three neural data sets. The first was simulated spiking from a simple two-cell network; the second was population-level spikes during human seizures; and the third was multi-unit activity from mouse auditory and prefrontal cortex. In all three cases, enhancement between intrinsic and extrinsic effects was estimated using a score derived from the GLM framework. Additionally, confidence intervals for the scores were computed using a data bootstrapping procedure (see Methods). Varying simulation parameters and computing the enhancement score, we found that the score is useful for assessing the model identifiability and that confidence intervals accurately reflected changes in the system dynamics. Finally, we applied the same enhancement analysis procedure to two cases of *in vivo* neural spiking. The first was ictal discharges observed during human seizure, where we found significant evidence of enhancement between intrinsic and extrinsic effects. The second was multi-unit activity in mouse auditory networks (auditory and prefrontal cortex) during passive listening, where we observed redundancy between the effects and, in a few cases, independence.

The first data set consisted of simulated data from a simple two cell network. This network has only six degrees of freedom and a fully known structure, which allows us to verify the accuracy of our methods and build intuition. The network is diagrammed in Figure III.2.A. In these statistical models of a neuron, spikes from Cells X and Y are

drawn from an inhomogeneous Poisson distribution. Cell X (Y) spikes at a baseline rate of 20 Hz (10 Hz).

Additionally, the cells' spike probabilities fluctuate according to two factors: intrinsic and extrinsic effects. Intrinsic effects depend on each cell's own spike history. These are represented in Figure III.2.A as inhibitory self-synapses (filled circles) with strengths w_1 and w_3 . When Cell X spikes, its spike probability is multiplied by a factor $\exp(-w_1)$ one millisecond later. As shown in the time course of the simulated intrinsic effects (Figure III.2.A, blue curve at left), this depression in spike probability persists for up to one second but rapidly decays. By 100 milliseconds post-spike, intrinsic effects depress the spike probability by a factor $\exp(-0.5w_1)$. In our case, for example, Cell X has a baseline spike probability of 0.02 spikes per bin and w_1 equal to 1. When Cell X spikes, its spike probability reduces to approximately 0.007 spikes per bin at the next time step. But spike probability eventually rebounds to 0.012 per bin after 100 milliseconds and 0.018 after 200 milliseconds. Cell Y follows similar dynamics with strength defined by w_3 , which is also set equal to 1.




Figure III.2. Enhancement depends on the type of system feedback.

(A) Schematic of the simulated network. Cells X and Y have baseline spike rates of 20 and 10 Hz, respectively. Spike probabilities are decreased by intrinsic effects (filled circles) producing refractory periods with strengths w_1 and w_3 . Spike probabilities are also modulated by extrinsic effects (open circles) that mimic synaptic coupling. When the coupling strengths w_2 or w_4 is positive (negative), the drive excites (inhibits) spiking. Time courses of intrinsic and extrinsic effects are pictured; both rapidly decay after 200 ms but persist for up to one second. (B) Enhancement scores along two-dimensions of parameter space. Colors depict median enhancement ratios over 1000 simulations (see Methods). If X sends inhibitory feedback in response to drive from Y (top left), then intrinsic effects and extrinsic effects become redundant. On the other hand, if the drives between X and Y match in sign (diagonal), then effects show enhancement. (C) Comparison of spiking for the three networks labelled in (B). Network N1 exhibits feedforward inhibition: Y excites X, which inhibits Y. Network N2 exhibits recurrent excitation or positive feedback: X and Y both excite each other, leading to frequent overlap in spiking. Because this network spikes so frequently, the example raster is shorter (1 second long) than the other rasters (5 seconds long). Network N3 exhibits mutual inhibition: X and Y inhibit each other, so that they tend to spike at different times.

Extrinsic effects, depicted as filled circles in Figure III.2.A, mimic the effects of synaptic coupling. In the statistical model, these terms allow Cell X's spike history to drive Cell Y's spike probability with strength w_2 , and vice versa with strength w_4 . Negative (positive) weights lead to statistical coupling that mimics inhibitory (excitatory) synaptic drive between neurons. Here we varied w_2 and w_4 from -5 to 5, spanning a range of inhibitory and excitatory values. The resulting levels of enhancement are shown in Figure III.2.B. In this figure, color depicts the enhancement score; cool colors indicate redundancy, while warm colors indicate enhancement.

We find that the level of enhancement depends critically on the signs of w_2 and w_4 . This is evident from the warm/cool color patches in Figure III.2.B, which roughly divide the plane into four quadrants separated by the lines $w_2 = 0$ and $w_4 = 0$. However, unlike sign, the magnitudes of w_2 and w_4 have little influence on enhancement or

redundancy, as evidenced by the approximately constant enhancement score within each quadrant.

To further examine the simulation results depicted in Figure III.2.B, we divide the model state space into three representative networks N1-N3. These three networks (labelled in Figure III.2.B) have coupling terms with different signs; in N1 both coupling terms are negative, while in N2 and N3 the coupling terms have opposite sign. As a result, these combinations express different relationships of excitation and inhibition between cells. Their range of behaviors is illustrated by example spike trains in Figure III.2.C.

In the simulated example, we know the true values of the synaptic weights, and can examine the estimation procedure by comparing the estimated parameters to those true values. In Table III.1 we quantify the error in parameter estimation for networks N1, N2, and N3. We find that describing any of the three networks with only one type of effects (individual models) leads to bias in the parameter estimates (first two columns). This is expected, since a one-effect model incompletely describes the network generating the data. We will now interpret the various biases that occur in light of the network connectivity.

In the first type of network (N1), w_2 is positive and w_4 is negative. Classically, this would be described as feedforward inhibition or negative feedback network (Freund and Antal 1988; Li et al. 2014). In this configuration, Cell Y excites Cell X, which in turn

inhibits Cell Y. Example spiking rasters in Figure III.2.C illustrate this pattern: Each time Cell Y spikes, it triggers a burst of spikes in X, which immediately stop Y's spiking. This type of network leads to redundancy. Intuitively, the reason for this redundancy is that when Y excites X which inhibits Y, the net effect is that Y inhibits Y. This post-spike inhibition can be described equally well as self-inhibition (Y inhibits Y through intrinsic effects) or synaptic inhibition (X inhibits Y through extrinsic effects). Such an interpretation is supported by Table III.1, where the amount of inhibition is overestimated for individual models of intrinsic and extrinsic effects (row N1, first and second column). To see this, note that the parameter difference between the estimated and true values of w_3 in this case is 3.166, meaning that the amount of self-inhibition is overestimated in the individual model. Similarly, the difference between estimated and true values of w_4 is -0.244, indicating that the extrinsic effects are estimated to be even more inhibitory than they actually are.

In network N2, the synaptic weights are both positive. This could be described as recurrent excitation or positive feedback, where either cell can initiate a prolonged sequence of mutual spiking in the other. Such connectivity leads to intervals of rapid spiking. This is illustrated in Figure III.2.C. We note that, in this figure, the scale bar indicates 200 ms, compared to 1000 ms in the other example rasters. In this scenario, the spiking is much denser, as expected for this type of network configuration. Network N2 also exhibits enhancement. Intuitively, the reason for this enhancement is that because

they covary but modulate spiking in opposite directions, intrinsic and extrinsic effects suppress one another's effects in an individual model. In a joint model, these covarying effects can be properly separated, as suggested by the third and fourth column of Table III.1.

Finally, in network N3, w_2 and w_4 are both negative. This is the most complicated case, which can be interpreted physically as a pair of mutually-coupled inhibitory cells. The negative synaptic weights cause each cell to inhibit the other, leading to indirect increases in spike probability. For example, if Cell Y spikes, Cell X is inhibited and becomes less likely to inhibit Cell Y; thus Y's future spike probability is increased relative to not spiking. We note that this configuration differs from the network in N2, where w_2 and w_4 are both positive. To see this, compare the example spike rasters in Figure III.2.C; in N1 the cells frequently spike together, whereas for N3 the cells tend to spike out-of-phase with one another. Like mutual excitation, mutual inhibition also leads to enhancement. Intuitively, the reason for this enhancement is similar to network N2. In a model without extrinsic effects, the inhibitory influence of X onto Y is misattributed to the self-history effects of Y. This misattribution suppresses the true intrinsic effects in Y; by including the positive impact of the extrinsic effect from X (due to release of inhibition from X on Y) in the modeled intrinsic effects of Y, we inaccurately estimate the inhibitory effect of the Y's intrinsic effects. We interpret this misattribution as contributing "negative information" to our ability to predict the next

spike of Y. In a joint model that includes both intrinsic and extrinsic effects, this negative information is reduced. By modeling the extrinsic effects term from X, the estimated intrinsic effects in Y becomes more accurate – this term is no longer suppressed by the a net excitatory extrinsic effect. Therefore, including the extrinsic effects in the joint model increases the amount of correct information provided by the intrinsic effects term. This is the essence of enhancement; adding a new term to a model increases the information of other terms.

We note that self-inhibition is overestimated in the individual model (Table III.1, first column).

Table III.1. Differences between estimated and true parameters for individual and joint models.

Differences are measured as Estimated Parameter – True Parameter, so positive (negative) values represent an estimate that is larger (smaller) than expected. In interpreting the values, note the signs of w_3 and w_4 in Equation 3.1: w_3 is fixed to be -1, so its values represent the amount of self-inhibition. For instance, a positive (negative) w_3 in the table indicates the self-inhibition is over-estimated (under-estimated) by the model (see first column). The values of w_4 vary in sign depending on the network, so interpretation of the differences in the table requires care. When w_4 is negative in both the network and the table (N1, second column), inhibition is over-estimated; when w_4 varies in sign between network and table, the extrinsic effects are under-estimated (N2 and N3, second column). In these cases, estimation errors in the joint model (third and fourth column) approach zero, suggesting this network model becomes identifiable when correctly specified.

Parameter Difference,	$w_3,$	$w_4,$	$w_3,$ Joint	$w_4,$ Joint
Estimated – True	Individual	Individual	Model	Model
	Model	Model		
<i>Network N1. $w_2 + / w_4 -$</i>				
Mean	3.166	-0.244	0.599	0.019

Parameter Difference, Estimated – True	$w_3,$ Individual Model	$w_4,$ Individual Model	$w_3,$ Joint Model	$w_4,$ Joint Model
Standard Deviation	2.050	0.574	1.813	0.635
<i>Network N2. $w_2 + / w_4 +$</i>				
Mean	-0.860	-0.942	-0.02	-0.0965
Standard Deviation	0.180	0.295	0.177	0.341
<i>Network N3. $w_2 - / w_4 -$</i>				
Mean	-0.964	0.395	0.096	-0.034
Standard Deviation	0.428	0.510	0.526	0.605

In this chapter, we describe not only how to compute a model's enhancement score, but also a method to estimate confidence bounds. This estimate of variability is important, for example, when interpreting an enhancement score, and determining whether it lies significantly above or below 0. To estimate the confidence bounds, we developed a bootstrapping procedure which proceeds as follows: Suppose the original model had a design matrix X and spike train y . First, we resample rows from X and y with replacement 100-1000 times. We choose the number of resamples depending on the

size of the data and the computational load; for simulated networks and mouse multi-unit activity we compute 1000 resamples, for seizure data we compute 100 resamples.

We then compute the enhancement score for each resampled data set. After repeating for all resamples, we take as confidence intervals the 5th and 95th percentiles of resampled enhancement scores. We show examples of the bootstrapped confidence intervals (shaded regions) for the simulated data in Figures III.3, and for the *in vivo* data in Figures III.4.A and III.5.A-B.

To test the reliability of this bootstrapping procedure for estimating confidence intervals of the enhancement score, we compared the variability estimated by this data bootstrapping procedure with that from parameter bootstrapping, which requires simulating new data and calculating the enhancement score. Parameter bootstrapping is only possible in cases where the true model is known, such as in the simulated network implemented here. Figure III.3 shows confidence intervals computed by data bootstrapping (blue) and parameter bootstrapping (black) for a range of w_4 values with w_2 held fixed at 3. Thick lines are the median enhancement scores over 1000 independent 20-second samples of the network. We note that the thick blue line corresponds to a horizontal slice across the image of Figure III.2.B at $w_2 = 3$. Across this range of w_4 values, the score transitions from negative (redundancy) to positive (enhancement). Throughout this transition, the two bootstrapping methods give similar confidence intervals, as indicated by the overlap in Figure III.3. While this does not

show that data bootstrapping works for all networks, it confirms that the procedure estimates reasonable confidence intervals in the case in which the true model is known.




Figure III.3. Confidence intervals for the enhancement score can be reliably bootstrapped with data.

Comparison of 95% confidence intervals (shaded) for the enhancement score under two bootstrapping techniques (see Methods). The confidence intervals are in approximate agreement whether they were measured from 1000 independently-generated samples of simulated data (parameter bootstrap, blue) or by resampling 100 times from a single pair of spike trains (data bootstrap, gray).

With the intuition gained from analyzing enhancement in the simulated data, we now compute the enhancement between intrinsic and extrinsic effects from two neural data sets recorded *in vivo*. The first data set consists of ictal discharges (IDs): large-amplitude population spikes, recorded during human seizures. These IDs were described in detail in Chapter II. Figure III.4.A shows the distribution of enhancement scores,

including confidence intervals (shaded), for a total of 377 MEA electrodes from eleven seizures. We find that only a minority of electrodes show redundancy or independence between effects. Instead, the majority of electrodes show significant evidence of enhancement.

A challenge of interpreting enhancement is that it may manifest through different sources, as illustrated in the examples shown in Figure III.2.B. To further explore the source of enhancement during seizure, we compared the effects estimates for individual models where IDs are described purely by intrinsic or extrinsic effects (i.e. rhythmic activity or synaptic coupling alone) versus joint models where both effects are included. We showed in Table III.1 that, in simulation, the un-modeled effects biased parameter estimates in both the positive and negative direction. Therefore, we expect in the human seizure data that the intrinsic and extrinsic effects could change in either direction. We observed that intrinsic effects indeed changed in both directions when estimated under a joint model. A consistent finding across all eleven seizures, however, is that extrinsic effects increased in the joint model (Figure III.4, blue) as compared to the individual models (Figure III.4.B, red).

Interpreting this result in terms of network dynamics, we note that the extrinsic coupling during seizure is quite strong (i.e., neighboring brain regions become more influential). Without negative feedback such as a refractory period, such positive feedback could presumably trigger a runaway spike cascade. We speculate that intrinsic

effects, by mitigating instability associated with strong coupling, enhance extrinsic effects by allowing them to be large enough to reflect the spatial features of the data. In this way, inclusion of the intrinsic effects in a joint model unmask features in the extrinsic effects misallocated in an individual model.

The third data set for which we analyzed enhancement was neural spiking recorded in mouse cortex. The details of this experimental preparation and recording may be found in (James 2016). Briefly, extracellular spikes were recorded simultaneously in prefrontal cortex and auditory cortex during passive presentation of an auditory stimulus. Individual neurons were then isolated using standard spike sorting techniques during 39 different experimental blocks, 297 cells in total. Here, we compute the enhancement score for the data aggregated from all experimental blocks. For each identified cell, we fit model parameters and compute the enhancement; Figure III.5.A illustrates the resulting enhancement scores for all blocks and cells, resulting in 297 models in total. As in the other examples, shaded regions correspond to 95% confidence intervals estimated using the data bootstrapping procedure (see Methods). Observing the enhancement scores in Figure III.5.A, we find that the majority of confidence intervals are below zero, indicating redundancy in the model effects. For a subset of cells, the enhancement scores reach zero, indicating independence. After accounting for the number of comparisons performed, we do not find significant evidence for enhancement in these data. In other words, the extrinsic effects terms for auditory cortex and prefrontal cortex population

spiking describe similar spiking features already captured in the intrinsic effects (self-history) terms. These results suggest that the recorded spike trains do not influence one another directly as much as they influence themselves. This has important consequences for any correlation-based analysis of the multi-unit activity: namely we expect to find confounds in correlation-based analysis. As shown in Chapter II, such confounds occur unless the self-history effects and ensemble history effects are jointly modeled. We hypothesize that these confounds are related to activity induced by the shared stimulus. This hypothesis is consistent with the idea that external input can strongly drive primary sensory areas such as auditory cortex.




Figure III.4. Enhancement is prevalent in ictal discharges during human seizure.

(A) Histogram of enhancement scores for statistical models of ictal discharges (IDs). Each data point, i.e. model, represents one electrode over a 30-second interval during the middle of seizure. Shaded are 95% confidence intervals obtained through a data bootstrap procedure (see Methods). In this case, most models

show enhancement between intrinsic and extrinsic effects; virtually none show redundancy. (B) Changes in intrinsic effects (top) and extrinsic effects (bottom) when the effects are estimated individually versus jointly. In this case, intrinsic effects and extrinsic effects generally increase. This consistently holds for extrinsic effects in the seizure data, while in some cases intrinsic effects decrease.




Figure III.5. Mouse cortical networks show redundancy and independence between spike trains during passive listening.

(A) Distribution of enhancement ratios for models of multi-unit activity in mouse cortex during passive listening. Each data point, i.e. model, represents a sorted cell during a recording block of approximately 100 trials. Confidence intervals bootstrapped with 1000 resamples (see Methods). Most models show redundancy between intrinsic and extrinsic effects (negative score). (B) Zooming in on models with the greatest scores, we see that for a subset of models the intrinsic and extrinsic effects are independent (zero score). (C) Effects estimates across the model hierarchy for one of the redundant cells in (A). In the case shown, intrinsic effects and extrinsic effects both decrease at all lags. While this holds for virtually all extrinsic effects models estimated for the auditory data, some intrinsic effects profiles increase under the joint model. (D) Effects estimates across the model hierarchy for one of the independent cells in (B). As the term implies, estimates are unchanged whether the effects are estimated separately or together.

To further explore the lack of enhancement in these data, we compared estimates for individual models, where the effects are estimated separately, versus joint models, where effects are estimated simultaneously. Examples of these different effects for cells which exhibit redundancy and independence are shown in panels C and D of Figure III.5, respectively. In the case of redundancy (Figure III.5.C) we find that both intrinsic and extrinsic effects decrease when estimated together. This is consistent with the notion of redundancy, in which effect sizes grow weaker as variables share explanatory power. In the case of independence (Figure III.5.D) we find effect sizes remain unchanged. This is also consistent with independence: the effects explain different aspects of spiking and thus are insensitive to one another. Again, we speculate that redundancy occurs because intrinsic and extrinsic effects both indirectly capture features of the stimulus-response.

4. Discussion

We estimated the degree of enhancement in three neural data sets. Despite the diverse origins of the data, the generalized linear modeling (GLM) framework offers a

simple and unified way to analyze each data type. From the GLM deviance we defined an enhancement score, which characterizes how well two factors explain spiking together versus apart; if the enhancement score is positive (negative), the joint information is greater (less) than the sum of individual information. We also presented a bootstrap method to estimate confidence bounds for the score. Because of the flexibility of the GLM framework, these techniques can be extended to analyze a variety of informational interactions among the factors affecting neural spiking.

Motivated by Shannon's seminal work on information theory (1948), there has been much interest in quantifying information flow in the brain (MacKay and McCulloch, 1952; Quian Quiroga and Panzeri, 2009; Rolls and Treves, 2011). This approach has been mostly limited to studying sensory encoding, with researchers measuring how much information spike trains from different subsets of sensory neurons convey about a stimulus (Warland et al., 1997; Rolls, 2003; Pillow and Simoncelli, 2006). This stage of processing could be summarized as information flow from stimulus to neuron; however, the subsequent information processing from neuron to neuron is less well-characterized. Indeed, measuring information flow for even simple neural network models remains an open problem (Gat and Tishby, 1998).

One challenge in the analysis of information flow in neural systems is that the measures for comparing information content can be complicated or specific to a particular application (Dimitrov et al., 2011). Moreover, developing an interpretation for

the functional meaning of enhancement remains a challenge. One possible interpretation is that systems with enhancement are likely to be functionally integrated and functionally specialized networks (Balduzzi et al., 2008), while systems with redundancy are likely to be strongly modular.

Here the method proposed for studying enhancement, or information gain due to multiple factors affecting spiking, is quite flexible and could be extended to investigate information in theoretical models as well as *in vitro* and *in vivo* neural systems. Moreover, the enhancement score we define to compare information is easily estimated from the GLM theory, and confidence intervals can be rapidly computed with as few as 100 iterations of a bootstrapping procedure.

In this chapter, we discuss tools for computing enhancement and provide three preliminary analyses. In general, we find that enhancement results from the unmasking of misallocated spike-influencing factors that suppress one another due to covariance and opposing signs of modulation. However, understanding this phenomenon in a deeper way will require continued investigation.

For our simulations, we designed a model system with as little complexity as possible to facilitate understanding. As a result, several aspects of the system remain to be explored. The first is the number of cells; we simulated a network of only two cells with all-to-all connectivity. We note that, using the point process GLM, we are able to generate simulated data rapidly, so larger networks with more elaborate connectivity

could easily be considered. Understanding how enhancement manifests in larger networks is a rich area for exploration, as previously network size has been shown to play a role in the observation of enhancement among stimulus information (Narayanan et al., 2005). Furthermore, neuroscience experiments necessarily sample a limited number of cells. Simulating larger networks where only a subset of cells is observed, similar to (Pillow and Latham, 2007; Gerhard et al., 2011; Kim et al. 2011), would be useful for assessing sensitivity and specificity of the enhancement score with respect to hidden influences.

Another important aspect that impacts measures of enhancement is the time course of intrinsic and extrinsic effects in the GLM. We chose intrinsic and extrinsic effects to be single exponential functions for simulation. However, in the observed cases of enhancement, namely the seizure data, intrinsic effects have multiple peaks. How incorporating these multiple time scales would affect the enhancement with extrinsic effects remains unclear. Previous results suggest intrinsic effects can strongly influence the propagation of spiking activity through a network (Timofeev et al., 2004; Schneidman et al., 2006; Tripathy et al., 2013; Padmanabhan and Urban, 2010; Gjorgjieva et al., 2014; Han et al., 2015; Schneidman, 2016). Elucidating the relationship between time scales and intrinsic/extrinsic enhancement may shed more light on this phenomenon.

A related and unanswered question is how enhancement relates to seizure. Intrinsic/extrinsic enhancement scores vary across the microelectrode array (MEA), but it

is unclear why. Earlier work has found that, for some patients, inter-electrode redundancy corresponds closely with patches of local synchrony and the seizure onset zone. Conversely, informational enhancement was associated with the surrounding area (Erramuzpe et al., 2015). Redundancy and enhancement also were observed most prominently in high and low frequency bands, respectively. Here we have analyzed a different type of enhancement in a more spatially-restricted recording: namely, microelectrodes as opposed to electrocorticography (ECoG). We observed no redundancy, perhaps because the recording is restricted to a small area outside the seizure onset zone. We did, however, observe that enhancement was linked to low frequency activity: Intrinsic effects often became unmasked at long lag times (~1500 ms, 0.7 Hz) when estimated in a full model, though not always. More work will be necessary to link observations of enhancement at the two scales. One possibility is that enhancement reflects combined changes in excitation and inhibition (Boido et al., 2014).

If this type of enhancement relates to the mechanisms underlying seizure, then the scores could ultimately inform treatment strategies. For example, recording sites with the highest enhancement scores may be interpreted to represent functional hubs because they signify areas where both intrinsic and extrinsic effects are prevalent. Therefore, these sites could serve as prime targets to disrupt seizure through stimulation or drug delivery.

Another direction for future work would examine other types of enhancement. Here we have defined and analyzed intrinsic/extrinsic enhancement, which is computed

from a GLM with intrinsic and extrinsic effects. By building other GLMs, we can compute enhancement between other factors. For example, one interesting quantity might be oscillatory/extrinsic enhancement, which compares how spikes and oscillatory fields jointly influence a spike train. Currently such influence is estimated separately, for instance by spike-spike or spike-field coherence, but measuring enhancement with a GLM could offer a more comprehensive picture (Lepage et al., 2013). Similarly, with an appropriate choice of GLM we could also compute extrinsic/extrinsic enhancement. This measure would compare the effect of network coupling from multiple sub-networks and would relate most directly to the earlier uses of information theory in neuroscience.

CHAPTER IV. POINT PROCESS MODELING REVEALS A HIERARCHY OF FUNCTIONAL NEURON TYPES BASED ON INTRINSIC EFFECTS

1. Introduction

Varying in morphology, genetic expression, spike patterns, and many other factors, cells in the brain exhibit considerable diversity (Kawaguchi & Kondo, 2002; Moore et al., 2010). A number of characteristic cell types have been described, particularly in early sensory systems like the retina (Seung & Sümbül, 2014); however, classification schemes vary and are often limited in their scope (Gonchar, 2008; Rudy et al., 2011). Building a complete and consistent classification of the brain's cell types remains an open problem in the field of neuroscience, and is a critical component of the NIH BRAIN Initiative (Bargmann et al., 2014). Such a scheme is important because the distribution of neuron types constrains how and where computations can take place in the brain.

A comprehensive list of cell types, such as the one proposed by the BRAIN Initiative, has two related tasks; namely, it must characterize both functional and biological aspects of the neuron. Most descriptions of function involve electrophysiological measures such as the width of an action potential and the ratio of the first interspike interval to the second (Toledo-Rodriguez et al., 2004). More recently, molecular assays have made possible exploration of the genetic diversity of neurons and the classification of neuron on this basis (Wang et al., 2009).

The relationship between biological and functional types, however, has been incompletely mapped. Mensi et al., (2012) characterized and compared function for different neuron types in terms of statistical models. They showed that differences in model estimates between cell types were large enough to enable automatic classification of neurons into three clusters that overlapped with familiar biological classes.

In this chapter, we apply the point process GLM to generate functional profiles describing the spike outputs of single neurons. We analyze spike data from two neural populations: a set of inhibitory interneurons in the mouse medial entorhinal cortex (MEC), and a set of ten biological cell types from the Allen Cell Types Database. We find that, even though the records are short (3-5 seconds), the intrinsic effects profiles successfully cluster both cell populations; in other words, the full range of effects curves from the different neurons can be reduced to a small number of representative types. Like Mensi et al. (2012), our clusters overlap with biological cell types for both the MEC and Allen Cell Types data. These results represent an initial proof of concept and motivate the application of the proposed methods to a larger array of neural spike train data.

2. Methods

Using statistical models, we analyze the spike trains of various types of neuron in response to applied current stimulation. Here we describe the protocols for stimulating cells, recording spikes, estimating models, and clustering model effects. All analysis and

visualization was performed using custom MATLAB software (Mathworks; Natick, MA).

2.1. Medial Entorhinal Cortex Interneuron Data

First we analyze spike trains recorded from interneurons in the mouse medial entorhinal cortex (MEC). The cells we consider are a subset of the ones recorded and analyzed by Martinez (2015). We briefly repeat relevant methods here and invite the interested reader to consult the original work for more details. Brain slices were harvested from the superficial layers of medial entorhinal cortex (MEC) of 18-35 day old transgenic mice. Two transgenic strains were used: cre-dependent GAD2-IRES-tdTomato transgenic mice (Taniguchi et al., 2011, The Jackson Laboratories, strain 010802), which labeled glutamic acid decarboxylase 2 gene (GAD2) expressing cells and thus facilitated targeting of GABAergic cortical interneurons; and PV-tdTomato transgenic mice (Hippenmeyer et al., 2005, The Jackson Laboratories, strain 008069), which labeled all parvalbumin (PV) expressing cells and thus facilitated targeting of the specific PV+ genotype in inhibitory interneurons.

All electrophysiological protocols were conducted in current clamp and were performed within 30 minutes of breaking the cell membrane to engage the whole cell patch clamp recording. For each cell, a bias current was applied in current clamp to polarize the cell to -70 mV. A series of one-second current pulses (with a four second rest between pulses) were injected to determine the frequency-current (F-I) relationship

of the cell. These current pulses ranged from -100 pA to up to 1500 pA, depending on what current amplitude was required to reach a firing rate plateau, and were introduced in 20 nA increments. For point process modeling, we considered only the 5 trials that received the most stimulation for each cell. From the original 115 interneurons analyzed in (Martinez, 2015), we also excluded any neuron that spiked fewer than 20 times during stimulation, leaving 90 cells to analyze.

2.2. Allen Institute Cell Types Data

The Allen Institute for Brain Science shares a number of free data sets and tools with the neuroscience community. Among them is the Allen Cell Types Database, a collection of multimodal data (electrophysiology, morphology, gene expression) from single neurons in the mouse visual system (Hawrylycz et al., 2016). Cells in the database come from a variety of transgenic mice, each of which facilitates targeting different cell types. These data can be viewed and downloaded online at <http://celltypes.brain-map.org>. Also available online are technical whitepapers that describe the protocols for cell labeling, recording, and reconstruction. Here we repeat the relevant methods related to electrophysiology (Allen Institute for Brain Science, 2016) and invite the interested reader to consult the whitepapers for more details.

Brain slices are prepared from P45 to P70, male and female transgenic mice (either an interneuron or layer specific Cre driver line crossed to an Ai14 tdTomato

reporter line). In total, the Allen Cell Types Database consists of cells from sixteen transgenic lines. Cells considered here – ones for which a 'Noise 1' stimulus trial was recorded (see below) – come from ten of the sixteen lines, described in Table IV.1.

Slices (350 μm) are sectioned using a vibrating microtome: each slice is imaged to aid in brain region identification and registration to the Allen Mouse Common Coordinate Framework (CCF). Whole cell current clamp recordings are made from identified, tdTomato-positive neurons or nearby tdTomato-negative neurons for a subset of the experimental data. Electrophysiology data were reported with metadata detailing experimental conditions such as electrode resistance, tight seal resistance, and series resistance, as well as more granular details such as bath temperature and amplifier settings on a sweep by sweep basis.

Neurons in the Allen Cell Types Database are subjected to a variety of stimulation patterns. Stimulation waveforms are designed to: 1) interrogate intrinsic membrane properties that contribute to the input/output function of neurons, 2) understand aspects of neural response properties *in vivo*, and 3) construct and test computational models of varying complexity emulating the neural response to stereotyped stimuli. In this chapter, we restrict our focus to cells responding to a portion of the 'Noise 1' stimulus, a waveform that addresses design aim (2) above.

The Noise 1 stimulus is pink noise with a coefficient of variation (CV) equal to 0.2, chosen to resemble *in vivo* data. An entire trial in the Database consists of three, 3 s

noise epochs each superimposed on square pulses at 0.75, 1, and 1.5 times the cell rheobase. Recovery intervals between stimuli within a trial are 5 s. Here we consider only the third 3 s block, or the time interval 18-21 seconds within each trial. This segment has a number of desirable properties: At 3 seconds, the time interval is longer than other stimulus patterns such as square pulses. The stimulus here is also strong enough to generate a number of spikes and noisy enough to avoid perfectly regular ISIs. These are important considerations because spikes trains that are too sparse or regular can introduce numerical issues during statistical model estimation (Gerhard et al., 2013)

Table IV.1. Cre driver lines in the Allen Cell Types Database. We analyzed spike trains from ten different neuron types in the Allen Cell Types Database, described below. See the Allen Institute whitepapers for more information.

Cre Driver line	Originating lab	Public repository stock	Cell type expression	Layers of expression
Gad2-IRES-Cre	Z. Josh Huang	The Jackson Laboratory (010802)	Inhibitory	1-6
Htr3a-Cre_NO152	Nathaniel Heintz and Charles Gerfen	MMRRC (036680)	Inhibitory	1-6
Nr5a1-Cre	Bradford Lowell	The Jackson Laboratory (006364)	Excitatory	4
Ntsr1-Cre_GN220	Nathaniel Heintz and Charles Gerfen	MMRRC (036680)	Excitatory	6

Cre Driver line	Originating lab	Public repository stock	Cell type expression	Layers of expression
Pvalb-IRES-Cre	Silvia Arber	The Jackson Laboratory (008069)	Inhibitory	2-6
Rbp4-Cre_KL100	Nathaniel Heintz and Charles Gerfen	MMRRC (031125)	Excitatory	5
Rorb-IRES2-Cre	Allen Institute for Brain Science	The Jackson Laboratory (023526)	Excitatory	4-5
Scnn1a-Tg2-Cre	Allen Institute for Brain Science	The Jackson Laboratory (009112)	Excitatory	4
Scnn1a-Tg3-Cre	Allen Institute for Brain Science	The Jackson Laboratory (009613)	Excitatory	4-5
Sst-IRES-Cre	Z. Josh Huang	The Jackson Laboratory (013044)	Inhibitory	2-6

2.3. Point Process Model

For each cell, we fit a point process GLM with model terms representing the cells' baseline firing rate and intrinsic effects (i.e., self-history-dependence) (Truccolo et al., 2005). Because the cells are recorded individually, with all synaptic input blocked, we include no extrinsic effects. The overall model is similar to those implemented in

Chapters II and III. See Chapter II: Methods for a more detailed discussion of point process GLM and intrinsic effects. We briefly described the model form here. We discretize time into bins of equal size Δ so that the k -th time point is $t_k := k\Delta$ and spikes are represented as binary time series $dn_{i,k}$ equal to 1 if a spike occurs in cell i in the time interval $[t_k, t_{k+1})$ and 0 otherwise. As in other cases of single-unit activity analyzed in Chapter III, we set $\Delta = 1 \text{ ms}$.

The conditional intensity model for cell i is defined as

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} dn_{i,k-\kappa}.$$

where the intrinsic effects extend up to 100 milliseconds for the MEC data set and up to 500 milliseconds for the Allen Cell Types data.

2.3.1. Dimensional reduction via spline functions. To impose smoothness and avoid over-fitting, we replaced the full effects above with

$$\log \lambda_k = \beta_0 + \sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i),$$

where $\sum_{\kappa=1}^Q \beta_{\kappa} B_{\kappa}(dn_i)$ is a representation of the effects in terms of spline basis functions. For the MEC data, we place control points at lags 1, 5, 20, and 100 milliseconds. For the Allen Cell Types data, we place control points at 1, 20, 100, 300, and 500 milliseconds. Splines reduce the number of parameters in each respective case from 101 to 7 and from 501 to 8.

2.3.2. Smoothing of intrinsic effects for MEC data. Under the type of constant, high-amplitude stimulation delivered to MEC interneurons, many cells will spike periodically. In these cases, the intrinsic effects perfectly predict the absence of spikes at a number of lag times (namely, every lag time less than the period of spiking). As a result, intrinsic effects at these lags diverge to negative infinity (Gerhard et al., 2013). This can also introduce numerical issues at lags longer than the period. Specifically, we found in the MEC data that most inter-spike intervals were less than 10-20 ms and, correspondingly, that most intrinsic effects estimates were strongest at these short lags. However, the intrinsic effects often exhibited multiple peaks corresponding to different aspects of the activity. Typically, we observed a peak at the period of spiking, and at multiples of this period. To capture the peaks at shorter lags, we smoothed the intrinsic effects estimated from the MEC data with an exponentially decaying function before clustering. For each lag of τ milliseconds, we multiplied the intrinsic effects by $\exp(-0.06\tau)$. This processing step was only applied to the MEC data.

2.3.3. K-means clustering. After estimating the intrinsic effects for all cells within the two sets of experimental data, we clustered the (exponentiated) effects using a k-means algorithm. The procedure is an expectation-maximization algorithm that iteratively (1) computes cluster means, and (2) assigns each data point to the cluster with the nearest

mean according to Euclidean distance (James et al., 2013). This procedure is implemented in MATLAB as *kmeans.m*.

3. Results

We first fit GLMs to describe the spike-response profiles of a variety of stimulated neurons. These neurons come from two different *in vitro* data sets: one a collection of MEC interneurons of PV+ and non-PV+ type stimulated with five square 1-second pulses, the other a set of cells from the Allen Cell Types Database stimulated with a noisy 3-second pulse. We then clustered the resulting spike-response profiles and compared the genetic cell types of the cells within each cluster. We found that this procedure produced clusters consistent with the known divisions between PV+/non-PV+ and excitatory/inhibitory cells. We also found in both data sets unexpected divisions that suggest a further refinement of these divisions and merits further investigation.

Intrinsic effects profiles estimated for the 90 MEC cells are shown in Figure IV.1. The diversity of interneurons in the nervous systems is well-documented, as they are known to have widely-varying biological and firing properties (Hajos, 2004; Moore et al., 2010). We therefore hypothesized that there would be multiple types of response profiles, i.e. intrinsic effects, in the GLM estimates. Panels (A) and (B) of Figure IV.1 illustrate the effects before and after clustering, respectively. Effect size is depicted in color, with warm (cool) colors indicating that spikes are likely (unlikely) after a specific lag time.

The blue streaks near lag 0, for instance, reflect the refractory period of each cell. At longer lags when the refractory period ends, warm streaks indicate a rebound-excitation-like phenomenon in certain cells, while green streaks indicate arrhythmic behavior. In Figure IV.1 (A) visual inspection reveals that these phenomena, particularly refractory period length and degree of rebound excitation, vary among the cells; however, it is unclear how many types of responses are present. Separating the effects into four clusters, as shown in Figure IV.1 (B) (k-means clustering; see Materials & Methods), helps clarify to what extent responses fall into particular types.




Figure IV.1. MEC interneurons' spiking profiles can be clustered into groups based on their estimated intrinsic effects.

(A) Intrinsic effects estimates in a GLM for MEC neurons undergoing *in vitro* stimulation. Most response profiles show a combination of excitability at short (5-10 ms) and/or moderate (20-40 ms) lags (i.e. 100-200 Hz and 25-50 Hz rhythmicity, respectively). (B) The same intrinsic effects profiles, clustered by similarity with a K-means algorithm. In both figures color indicates the factor by which spike probability changes at each lag. Warm (cool) colors suggest the likelihood of spiking is increased (decreased) relative to baseline.

When applying the k-means algorithm to the MEC data, we assumed there were four clusters. One approach to finding an optimal value of k is to identify the inflection point of the within-cluster distance (i.e. the value at which the rate of decrease is maximal) (James et al., 2013). Based on a preliminary comparison (i.e. running the algorithm one time), we visually identified such an inflection point at $k=4$. Because previous analysis of these MEC interneurons also identified four distinct types among the cells (Martinez, 2015), we therefore took $k=4$ throughout. However, the issue of how to cluster the functional profiles remains an open question, as we discuss below.

The relationship between PV+ genetic status and cluster membership is shown in Table IV.2. Of the four clusters identified by intrinsic effects, two correspond closely with PV+ cells (Clusters 1 and 2) and the other two clusters with non-PV+ cells (Clusters 3 and 4). Indeed, this separation holds for 80 of the 90 cells (89%), with only 2-3 exceptions per cluster.

Table IV.2. Spiking profile clusters correlate with genetic cell type in MEC interneurons. Clusters are assigned using a K-means algorithm (see Methods). Clusters 1 and 2 correspond closely with PV+ cells, while clusters 3 and 4 consist of mostly non-PV+ cells.

Cluster Number	PV+ Interneurons	Non-PV+ Interneurons
1	15	3
2	14	3
3	2	38

Cluster Number	PV+ Interneurons	Non-PV+ Interneurons
4	2	13
Total	33	57

In addition to the MEC interneurons we analyzed a collection of spike trains from the Allen Institute Cell Types Database. These spike data were recorded in 171 cells comprising ten different cell types. Intrinsic effects estimates for all 171 cells appear in Figure IV.2. As in Figure IV.1, we show in Figure IV.2 (A) the estimated intrinsic effects before clustering and in Figure IV.2 (B) the same estimated intrinsic effects after clustering. It is evident from panel Figure IV.2 (A) that short (~10 ms), moderate (30-70 ms) and long (100-200 ms) time scales figure prominently in the intrinsic effects. Figure IV.2 (B) confirms that most cells exhibit some combination of these three time scales, as the clusters have visually similar intrinsic effects profiles.




Figure IV.2. Allen Database neurons' spiking profiles can be clustered into groups based on their estimated intrinsic effects.

(A) Intrinsic effects estimates in the GLM for neurons in the Allen Cell Types Database recorded *in vitro*. Most response profiles show a combination of excitability at short (~10 ms), moderate (30-70 ms), and or long (100-200 ms) lags (i.e. 100, 20, and 10 Hz rhythmicity, respectively). (B) The same intrinsic effects profiles, clustered by similarity with a K-means algorithm with $k=11$. In both figures color indicates the factor by which spike probability changes at each lag. Warm (cool) colors suggest the likelihood of spiking is increased (decreased) relative to baseline.

We again compared the quality of clustering versus the number of clusters. In this case, we found an inflection point at $k=11$. We therefore set this as the number of clusters when applying the k-means algorithm to the Allen Cell Type data. Visual inspection of Figure IV.2 (B) reveals that three clusters (Cluster 6, Cluster 10, and Cluster 11) have only 1-4 members. These clusters suggest that several outliers – with intrinsic effects that differ from the other cell types – occur in the data. Excluding these outliers, we then find eight effective clusters for the intrinsic effects of these data.

After clustering the estimated intrinsic effects for the Allen Cell Type data, we compared the Cre driver lines represented within each functional cluster. The results are

summarized in Figure IV.3. Here rows represent the different cell types, with the six excitatory and four inhibitory types separated by a horizontal line, and the columns represent clusters estimated from the algorithm. Numbers indicate the cell counts within each cluster; unlabeled entries have a cell count of zero. Colors additionally show the relative distribution of each row, i.e. how much of each cell type's total count falls in a particular cluster.

The multiple red blocks and bands in Figure IV.3 suggest strong overlap between functional clusters and genetic cell types. In particular, Clusters 7-9 contain roughly half of all inhibitory cells (29/56, 52%) and Clusters 2-5 three-quarters of excitatory cells (86/115, 75%). We note that Cluster 8, whose intrinsic effects are characterized by fast time scales, contains 100% of the PV+ neurons. These results suggest that the estimated intrinsic effects capture functional properties of the cells' spiking activity that align with the differences in their genetics.

We also observe additional structure in the clustered functional profiles beyond the separation between excitatory and inhibitory cell type. To see this, we focus on the 41 cells whose classification violates expectation (i.e., from the 171 cells we exclude the 115 that follow the expected excitatory/inhibitory types [Clusters 2-5 and 7-9] and the 15 cells in "noisy" Clusters 1, 6, 10, and 11). Among this sub-collection of 41 cells, we observe two notable trends. First, 7 of 10 *Ntsr1*-Cre cells fall into the presumptive "inhibitory" clusters (Clusters 7-9). Second, 14 of 29 *Sst*-IRES-Cre cells are in excitatory

clusters. Therefore while many cells have a stereotypical response based on being excitatory or inhibitory, there are still exceptions. These results suggest that cells from Sst-IRES-Cre and Ntsr1-Cre lines in particular may have multiple sub-types. Such an idea is supported by results in neurobiology, where multiple genetic subtypes of somatostatin-positive interneurons have been observed (Gonchar, 2008). Functional profiling may help clarify which biological differences are superficial and which lead to changes in the cell's spike-response.




Figure IV.3. Spiking profile clusters correlate with genetic cell type from the Allen Cell Types Database.

Clustering results (horizontal axis) versus cell type (vertical axis). Numbers indicate the number of cells per cluster, while colors show the fraction of each cell type contained in a particular cluster. Unlabeled entries are zero. A horizontal line separates excitatory (top) versus inhibitory (bottom) cell types. Red blocks and bands indicate strong overlap between cluster and cell type. Clusters 2-5 and 7-9 correspond well with excitatory and inhibitory cells, respectively.

4. Discussion

In this chapter, we computed a set of functional response profiles for neurons of different cell types. These neurons came from two data sets. The first is a collection of 90 PV+ and non-PV+ interneurons in the mouse MEC. The second is a set of 171 cells from the Allen Institute Cell Types Database comprising ten different cell types. We found a diverse set of responses, as illustrated in Figures IV.1 and IV.2. Clustering these functional responses, we concluded that they strongly correlate with the biologically identified cell types in both data sets, as shown in Table IV.2 and Figure IV.3.

These results represent further proof of concept that the GLM intrinsic effects can be successfully used as a feature space to cluster neurons. Mensi et al. (2012) used a similar technique to analyze 27 cells and discussed its potential utility in a high-throughput setting. By analyzing nearly 200 cells of ten different types, we provide an example of the increased scale accessible to this statistical modeling approach. These results inspire confidence that high-throughput screening of a cells functional profile, as described in (Mensi et al., 2012), may be achievable.

While our results show proof of concept, they do not yet suggest a definitive method for classifying functional cell types. Indeed, several issues remain to be resolved.

The first is which clustering algorithm to use. As more data become available, the performance and assumptions of different clustering algorithms should be systematically compared. Here we applied k-means clustering, which assumes clusters are spherical and have similar numbers of cells. Another approach would be to use a hierarchical clustering method (James et al., 2013). The advantage of this approach would be that it assumes neither an *a priori* number of clusters, nor similar cluster sizes, which allows for greater flexibility in cluster discovery.

The second issue to be resolved is how to establish the number of clusters present in the data. Here we approach this by comparing the inter-areal cluster distance and identifying where it has an inflection point. Related approaches identify an inflection point based on the total variance explained (James et al., 2013), using a silhouette method where inter- and intra-cluster distances are compared (Rousseeuw, 1987), or by computing a bootstrapped gap statistic (Tibshirani et al., 2001).

Finally, we note that the clustering of the estimated intrinsic effects revealed distinctions beyond the most basic identification of cells as either excitatory or inhibitory. A cell's functional profile (i.e., how it spikes in response to stimuli) may reveal subtle distinctions that further sub-classify the cells within each group. Whether these distinctions are more relevant, less relevant, or complementary to distinctions based on genetics remains unknown. A more complete understanding of the identified sub-

classifications – and their potential relevance for behavior and disease – remains an important topic for further research.

CHAPTER V. CONCLUSION

1. Impact and Innovation

This dissertation concerns the problem of inferring history-dependent structure in neural spike trains. The point process GLM provides a flexible and rigorous framework for this type of analysis. We adapted this framework to apply in a variety of neural data sets in order to address clinical and experimental questions. Here we summarize the contributions of this work to the field and discuss future directions of research.

1.1. A rigorous and patient-specific approach to seizure analysis

Before our work on ictal discharges (IDs) in Chapter II, the point process GLM had been exclusively applied in neuroscience to characterize action potentials generated by individual neurons. IDs are field effects and physiologically distinct from action potentials, yet in Chapter II we show that IDs and action potentials can be analyzed using the same GLM framework. This is useful for at least two reasons. First, the point process GLM provides a rigorous procedure to estimate important features (e.g., history dependence) from observed data. Second, the point process GLM provides a bridge between the observed data and features of biophysically detailed models (e.g., time scales). This combination of properties becomes quite useful in the clinical domain: By using the point process GLM we construct patient-specific and seizure-specific models of activity, which then facilitates comparison and simulation, and development of hypotheses with biophysical interpretation.

1.2 A generalizable method for measuring enhancement

Enhancement has only been measured in a handful of neural systems. Previously, measures of enhancement have focused on how information *about a stimulus* is distributed across multiple spike trains. Transmitting information about external stimuli, however, is only one facet of the neural code. Another important challenge is to understand how information manifests at higher levels of abstraction not directly related to the stimulus, and how information is passed between brain areas. In Chapter III we developed a tool to measure this information flexibly and robustly. We created this tool using the point process GLM framework, and through comparison of increasingly complex statistical models. We showed in simulation that the proposed measure accurately detected cases of redundancy and enhancement, and developed assess the significance of these detections. We also demonstrated that enhancement occurs in scenarios beyond stimulus coding, for instance between intrinsic and extrinsic effects during human seizure.

1.3 A method for clustering neurons according to their function

In Chapter IV, we demonstrated how the GLM can be used to classify neuronal cell types. The functional description of cells offered by the GLM complements biological descriptions provided by genetic assays. Whereas earlier functional descriptions have relied on a series of ad-hoc descriptive measures (e.g. action potential width), our description is based on a larger statistical modeling framework that is easily

extended, quickly estimated, and has established techniques to assess goodness-of-fit. We found that intrinsic effects profiles reliably distinguished a number of cell types, most notably PV+ versus non-PV+ cells in the MEC and a large fraction of excitatory versus inhibitory neurons in the Allen Cell Types Database. Furthermore, this analysis of functional properties hints at sub-structure in both data sets: namely, potential divisions within cell types as well as exotic cell types composed of unexpected combinations.

2. Future Directions

2.1. Understanding and extending enhancement

In Chapter III, we used the GLM to compute informational enhancement between intrinsic and extrinsic effects. While classical enhancement has been interpreted in terms of sparseness of a neural code, the type of enhancement detected here is more challenging to interpret. To aid in our interpretation, we simulated data from a simple two-cell network that exhibits both redundancy and enhancement. These simulated data offered some insights – mostly notably a straightforward procedure to produce redundancy and enhancement between intrinsic and extrinsic effects.

However, understanding enhancement in a deeper way and how it relates to neural systems will require continued simulations. For our simulations, we designed a system with as little complexity as possible. As a result, several aspects of the system remain to be explored. First is the number of cells; we simulated a network of only two cells with

all-to-all connectivity. We note that, using the point process GLM, we are able to generate simulated data rapidly, so larger networks with more elaborate connectivity could easily be considered. Understanding how enhancement manifests in larger networks is a rich area for exploration, as previously network size has been shown to play a role in the observation of enhancement among stimulus information (Narayanan et al., 2005). Furthermore, neuroscience experiments necessarily sample a limited number of cells. Simulating larger networks where only a subset of cells is observed, similar to (Pillow and Latham, 2007; Gerhard et al., 2011; Kim et al. 2011), would be useful for assessing sensitivity and specificity of the enhancement score with respect to hidden influences.

Another important aspect that impacts measures of enhancement is the time course of intrinsic and extrinsic effects in the GLM. We chose intrinsic and extrinsic effects to be single exponential functions for simulation. However, in the observed cases of enhancement, namely the seizure data, intrinsic effects have multiple peaks. How incorporating these multiple time scales would affect the enhancement with extrinsic effects remains unclear. Previous results suggest intrinsic effects can strongly influence the propagation of spiking activity through a network (Timofeev et al., 2004; Tripathy et al. 2013; Padmanabhan and Urban 2010; Gjorgjieva et al. 2014; Han et al. 2015). Elucidating the relationship between time scales and intrinsic/extrinsic enhancement may shed more light on this phenomenon.

A related and unanswered question is how enhancement relates to seizure.

Intrinsic/extrinsic enhancement scores vary across the microelectrode array (MEA), but it is unclear why. If this type of enhancement relates to the mechanisms underlying seizure, then the scores could ultimately inform treatment strategies. For example, recording sites with the highest enhancement scores may be interpreted to represent functional hubs. Therefore, these sites could serve as prime targets to disrupt seizure through stimulation or drug delivery.

Another direction for future work would examine other types of enhancement.

Here we have defined and analyzed intrinsic/extrinsic enhancement, which is computed from a GLM with intrinsic and extrinsic effects. By building other GLMs, we can compute enhancement between other factors. For example, one interesting quantity might be oscillatory/extrinsic enhancement, which compares how spikes and oscillatory fields jointly influence a spike train. Currently such influence is estimated separately, for instance by spike-spike or spike-field coherence, but measuring enhancement with a GLM could offer a more comprehensive picture (Lepage et al. 2013). Similarly, with an appropriate choice of GLM we could also compute extrinsic/extrinsic enhancement. This measure would compare the effect of network coupling from multiple sub-networks and would relate most directly to the earlier uses of information theory in neuroscience.

2.2. Clustering by functional profile

In Chapter IV, we used the GLM to cluster cells by their functional profiles. These results demonstrate that the intrinsic effects are useful as a feature space for clustering. Our results do not, however, provide a definitive method for clustering. Indeed, many details remain to be explored. One key issue is the selection of clustering algorithm; here we chose k-means, which is practical for its speed and intuitive simplicity. However, its assumptions – namely that clusters have elliptic shape and uniform size – are likely too unrealistic in large-scale neural data. The latter is especially problematic, since it seems possible that certain cell types – and perhaps even the most interesting ones – may be relatively uncommon compared to others. Future work should compare clustering algorithms in a systematic way and, importantly, evaluate their underlying assumptions. Furthermore, such work should consider how pre-processing affects clustering results. We observed “noisy” clusters – proposed to consist of outliers – in the Allen Cell Types data and found perfect predictors, which complicate numerical estimation, in the MEC data. Both issues merit careful pre-processing, which could be done prior to fitting models or afterwards. The goodness-of-fit tests available for point process models could prove useful in evaluating which model effects to discard before the clustering stage.

Like enhancement, future work on clustering benefits from the flexibility of the GLM framework. A natural approach to extend the feature space for clustering is to

extend the type of GLM being estimated. Here we considered a GLM with intrinsic effects, but extrinsic effects (relating to network connectivity), oscillatory effects (relating to LFP dynamics), and stimulus effects (particularly for sensory systems) could all prove useful in classifying cells. Indeed, because the cells they record are stimulated with known patterns designed to mimic synaptic input, Mensi et al. (2012) include extrinsic effects estimates in their feature space. These effects are not appropriate for the data sets analyzed in this thesis, in which the cells are recorded independently. We therefore exclude extrinsic effects here. Future work could design stimulation and/or recording protocols that probe network connectivity across different cell types. This description would further enrich the functional profiles we have compiled in Chapter IV.

BIBLIOGRAPHY

- Abbott, Larry F. 1999. "Lapicque's Introduction of the Integrate-and-Fire Model Neuron (1907)." *Brain Research Bulletin* 50 (5): 303–4.
- Ahrens, Misha B., Liam Paninski, and Maneesh Sahani. 2008. "Inferring Input Nonlinearities in Neural Encoding Models." *Network: Computation in Neural Systems* 19 (1): 35–67. doi:10.1080/09548980701813936.
- Alarcon, G., J. J. Garcia Seoane, C. D. Binnie, M. C. Martin Miguel, J. Juler, C. E. Polkey, R. D. Elwes, and J. M. Ortiz Blasco. 1997. "Origin and Propagation of Interictal Discharges in the Acute Electrocorticogram." *Brain* 120: 2259–82.
- Allen Institute for Brain Science. 2016. "Technical White Paper: Electrophysiology v.3." Allen Institute for Brain Science.
- Aoi, Mikio C., Kyle Q. Lepage, Mark A. Kramer, and Uri T. Eden. 2015. "Rate-Adjusted spike–LFP Coherence Comparisons from Spike-Train Statistics." *Journal of Neuroscience Methods* 240 (January): 141–53. doi:10.1016/j.jneumeth.2014.11.012.
- Balduzzi, David, and Giulio Tononi. 2008. "Integrated Information in Discrete Dynamical Systems: Motivation and Theoretical Framework." Edited by Olaf Sporns. *PLoS Computational Biology* 4 (6): e1000091. doi:10.1371/journal.pcbi.1000091.
- Bargmann, Cornelia I, William T. Newsome, David Anderson, Emery N. Brown, Karl Deisseroth, John P. Donoghue, Peter MacLeish, et al. 2014. "BRAIN 2025: A Scientific Vision."
- Bartolomei, F., F. Wendling, J. Régis, M. Gavaret, M. Guye, and P. Chauvel. 2004. "Pre-Ictal Synchronicity in Limbic Networks of Mesial Temporal Lobe Epilepsy." *Epilepsy Research* 61 (1-3): 89–104. doi:10.1016/j.eplepsyres.2004.06.006.
- Berens, Philipp. 2009. "CircStat: A MATLAB Toolbox for Circular Statistics." *J Stat Softw* 31 (10): 1–21.
- Blanche, T. J. 2005. "Polytrodes: High-Density Silicon Electrode Arrays for Large-Scale Multiunit Recording." *Journal of Neurophysiology* 93 (5): 2987–3000. doi:10.1152/jn.01023.2004.

- Blumenfeld, H., G. I. Varghese, M.J. Purcaro, J.E. Motelow, M. Enev, K. A. McNally, A.R. Levin, et al. 2008. "Cortical and Subcortical Networks in Human Secondarily Generalized Tonic-Clonic Seizures." *Brain* 132 (4): 999–1012. doi:10.1093/brain/awp028.
- Boido, Davide, Vadym Gnatkovsky, Laura Uva, Stefano Francione, and Marco de Curtis. 2014. "Simultaneous Enhancement of Excitation and Postburst Inhibition at the End of Focal Seizures: Focal Seizures." *Annals of Neurology* 76 (6): 826–36. doi:10.1002/ana.24193.
- Brockwell, Peter J., and Richard A. Davis. 2010. *Introduction to Time Series and Forecasting*. 2. ed., corr. at 8. printing. Springer Texts in Statistics. New York, NY: Springer.
- Brown, Emery N, Riccardo Barbieri, Valérie Ventura, Robert E Kass, and Loren M Frank. 2002. "The Time-Rescaling Theorem and Its Application to Neural Spike Train Data Analysis." *Neural Computation* 14 (2): 325–46. doi:10.1162/08997660252741149.
- Brown, Emery N, Robert E Kass, and Partha P Mitra. 2004. "Multiple Neural Spike Train Data Analysis: State-of-the-Art and Future Challenges." *Nature Neuroscience* 7 (5): 456–61. doi:10.1038/nn1228.
- Brown, E N, D P Nguyen, L M Frank, M a Wilson, and V Solo. 2001. "An Analysis of Neural Receptive Field Plasticity by Point Process Adaptive Filtering." *Proceedings of the National Academy of Sciences of the United States of America* 98 (21): 12261–66. doi:10.1073/pnas.201409398.
- Calabrese, Ana, Joseph W. Schumacher, David M. Schneider, Liam Paninski, and Sarah M. N. Woolley. 2011. "A Generalized Linear Model for Estimating Spectrotemporal Receptive Fields from Responses to Natural Sounds." Edited by M. Fabiana Kubke. *PLoS ONE* 6 (1): e16104. doi:10.1371/journal.pone.0016104.
- Casti, a R. R., a. Omurtag, a. Sornborger, E. Kaplan, B. Knight, J. Victor, and L. Sirovich. 2002. "A Population Study of Integrate-and-Fire-or-Burst Neurons." *Neural Computation* 14 (5): 957–86. doi:10.1162/089976602753633349.
- Chen, Zhe, David F. Putrino, De Ba, Soumya Ghosh, and Riccardo Barbieri. 2009. "A Regularized Point Process Generalized Linear Model for Assessing the

- Connectivity in the Cat Motor Cortex.” *Annual International Conference of the IEEE* 2009 (January): 5006–9. doi:10.1109/IEMBS.2009.5334610.
- Chornoboy, E S, L P Schramm, and a F Karr. 1988. “Maximum Likelihood Identification of Neural Point Process Systems.” *Biological Cybernetics* 59 (4-5): 265–75. doi:10.1007/BF00332915.
- Cohen, Marlene R, and Adam Kohn. 2011. “Measuring and Interpreting Neuronal Correlations.” *Nature Neuroscience* 14 (7): 811–19. doi:10.1038/nn.2842.
- Cox, David R, and Valerie Isham. 1980. *Point Processes*. Chapman and Hall.
- Currie, Iain, and Athol Korabinski. 1984. “Some Comments on Bivariate Regression.” *The Statistician* 33 (3): 283. doi:10.2307/2988232.
- Daley, Daryl J., and David Vere-Jones. 2002. *An Introduction to the Theory of Point Processes. Vol. 1: Elementary Theory and Methods*. 2. ed., corr. print. Probability and Its Applications. New York, NY: Springer.
- de la Rocha, Jaime, Brent Doiron, Eric Shea-Brown, Krešimir Josić, and Alex Reyes. 2007. “Correlation between Neural Spike Trains Increases with Firing Rate.” *Nature* 448 (7155): 802–6. doi:10.1038/nature06028.
- Deng, Xinyi, Emad N Eskandar, and Uri T Eden. 2013. “A Point Process Approach to Identifying and Tracking Transitions in Neural Spiking Dynamics in the Subthalamic Nucleus of Parkinson’s Patients.” *Chaos (Woodbury, N.Y.)* 23 (4): 046102. doi:10.1063/1.4818546.
- Deng, Xinyi, Daniel F. Liu, Kenneth Kay, Loren M. Frank, and Uri T. Eden. 2015. “Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter.” *Neural Computation* 27 (7): 1438–60. doi:10.1162/NECO_a_00744.
- Destexhe, a. 1998. “Spike-and-Wave Oscillations Based on the Properties of GABAB Receptors.” *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 18 (21): 9099–9111.
- Dichter, MARC, and W. ALDEN Spencer. 1969. “Penicillin-Induced Interictal Discharges from the Cat Hippocampus. II. Mechanisms Underlying Origin and Restriction.” *Journal of Neurophysiology* 32 (5): 663–87.

- Dorn, J. D., and D. L. Ringach. 2002. "Estimating Membrane Voltage Correlations From Extracellular Spike Trains." *Journal of Neurophysiology* 89 (4): 2271–78. doi:10.1152/jn.000889.2002.
- Eden, U T, L M Frank, R Barbieri, V Solo, and E N Brown. 2004. "Dynamic Analysis of Neural Encoding by Point Process Adaptive Filtering." *Neural Comput.* 16 (5): 971–98. doi:10.1162/089976604773135069.
- Emerson, Ronald G, Christine A Turner, Timothy A Pedley, Thaddeus S Walczak, and Mennato Forgione. 1995. "Propagation Patterns of Temporal Spikes." *Electroencephalography and Clinical Neurophysiology* 94 (5): 338–48.
- Ergun, Ayla, Riccardo Barbieri, Uri T. Eden, Matthew A. Wilson, and Emery N. Brown. 2007. "Construction of Point Process Adaptive Filter Algorithms for Neural Systems Using Sequential Monte Carlo Methods." *IEEE Transactions on Biomedical Engineering* 54 (3): 419–28. doi:10.1109/TBME.2006.888821.
- Erramuzpe, Asier, Guillermo J Ortega, Jesus Pastor, Rafael G de Sola, Daniele Marinazzo, Sebastiano Stramaglia, and Jesus M Cortes. 2015. "Identification of Redundant and Synergetic Circuits in Triplets of Electrophysiological Data." *Journal of Neural Engineering* 12 (6): 066007. doi:10.1088/1741-2560/12/6/066007.
- Foot Smith, Elizabeth, and Lydia Bayne. 1991. "Joan of Arc." *Epilepsia* 32 (6): 810–15.
- Franks, Alexander, Andrew Miller, Luke Bornn, and Kirk Goldsberry. 2015. "Characterizing the Spatial Structure of Defensive Skill in Professional Basketball." *The Annals of Applied Statistics* 9 (1): 94–121. doi:10.1214/14-AOAS799.
- Frei, Mark G., Hitten P. Zaveri, Susan Arthurs, Gregory K. Bergey, Christophe C. Jouny, Klaus Lehnertz, Jean Gotman, et al. 2010. "Controversies in Epilepsy: Debates Held during the Fourth International Workshop on Seizure Prediction." *Epilepsy & Behavior* 19 (1): 4–16. doi:10.1016/j.yebeh.2010.06.009.
- Friedman, Lynn, and Melanie Wall. 2005. "Graphical Views of Suppression and Multicollinearity in Multiple Linear Regression." *The American Statistician* 59 (2): 127–36.

- Frohlich, F., T. J. Sejnowski, and M. Bazhenov. 2010. "Network Bistability Mediates Spontaneous Transitions between Normal and Pathological Brain States." *Journal of Neuroscience* 30 (32): 10734–43. doi:10.1523/JNEUROSCI.1239-10.2010.
- Gat, Itay, and Naftali Tishby. 1998. "Synergy and Redundancy among Brain Cells of Behaving Monkeys." *Advances in Neural Information Processing Systems*.
- Gawne, Timothy J., and Barry J. Richmond. 1993. "How Independent Are the Messages Carried by Adjacent Inferior Temporal Cortical Neurons?" *The Journal of Neuroscience* 13 (7): 2758–71.
- Gelinas, Jennifer N, Dion Khodagholy, Thomas Thesen, Orrin Devinsky, and György Buzsáki. 2016. "Interictal Epileptiform Discharges Induce Hippocampal–cortical Coupling in Temporal Lobe Epilepsy." *Nature Medicine* 22 (6): 641–48. doi:10.1038/nm.4084.
- Gerhard, Felipe, and Wulfram Gerstner. 2010. "Rescaling, Thinning or Complementing? On Goodness-of-Fit Procedures for Point Process Models and Generalized Linear Models." *Advances in Neural Information Processing Systems*.
- Gerhard, Felipe, Tilman Kispersky, Gabrielle J. Gutierrez, Eve Marder, Mark Kramer, and Uri Eden. 2013. "Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone." Edited by Olaf Sporns. *PLoS Computational Biology* 9 (7): e1003138. doi:10.1371/journal.pcbi.1003138.
- Gerhard, Felipe, Gordon Pipa, Bruss Lima, Sergio Neuenschwander, and Wulfram Gerstner. 2011. "Extraction of Network Topology From Multi-Electrode Recordings: Is There a Small-World Effect?" *Frontiers in Computational Neuroscience* 5. doi:10.3389/fncom.2011.00004.
- Gerstner, Wulfram, Werner M. Kistler, Richard Naud, and Liam Paninski. 2014. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. 1. publ. Cambridge: Cambridge Univ. Press.
- Gerwinn, Sebastian, Jacob H Macke, and Matthias Bethge. 2010. "Bayesian Inference for Generalized Linear Models for Spiking Neurons." *Frontiers in Computational Neuroscience* 4. doi:10.3389/fncom.2010.00012.
- Gjorgjieva, Julijana, Rebecca A. Mease, William J. Moody, and Adrienne L. Fairhall. 2014. "Intrinsic Neuronal Properties Switch the Mode of Information

Transmission in Networks.” Edited by Wolfgang Einhäuser. *PLoS Computational Biology* 10 (12): e1003962. doi:10.1371/journal.pcbi.1003962.

- Gonchar, Yuri. 2008. “Multiple Distinct Subtypes of GABAergic Neurons in Mouse Visual Cortex Identified by Triple Immunostaining.” *Frontiers in Neuroanatomy* 1. doi:10.3389/neuro.05.003.2007.
- Granger, Clive WJ. 1969. “Investigating Causal Relations by Econometric Models and Cross-Spectral Methods.” *Econometrica* 137 (3): 424–38.
- Hajos, N. 2004. “Spike Timing of Distinct Types of GABAergic Interneuron during Hippocampal Gamma Oscillations In Vitro.” *Journal of Neuroscience* 24 (41): 9127–37. doi:10.1523/JNEUROSCI.2113-04.2004.
- Hamilton, David. 1988. “Sometimes $R^2 > r_{yx1}^2 + r_{yx2}^2$: Correlated Variables Are Not Always Redundant.” *American Statistical Association* 41 (2): 129–32.
- Han, Ruixue, Jiang Wang, Haitao Yu, Bin Deng, Xilei Wei, Yingmei Qin, and Haixu Wang. 2015. “Intrinsic Excitability State of Local Neuronal Population Modulates Signal Propagation in Feed-Forward Neural Networks.” *Chaos: An Interdisciplinary Journal of Nonlinear Science* 25 (4): 043108. doi:10.1063/1.4917014.
- Haslinger, Robert, Gordon Pipa, and Emery Brown. 2010. “Discrete Time Rescaling Theorem: Determining Goodness of Fit for Discrete Time Statistical Models of Neural Spiking.” *Neural Computation* 22 (10): 2477–2506.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2001. *The Elements of Statistical Learning*. Vol. 1. Springer series in statistics Springer, Berlin.
- Hawrylycz, Michael, Costas Anastassiou, Anton Arkhipov, Jim Berg, Michael Buice, Nicholas Cain, Nathan W. Gouwens, et al. 2016. “Inferring Cortical Function in the Mouse Visual System through Large-Scale Systems Neuroscience.” *Proceedings of the National Academy of Sciences* 113 (27): 7337–44. doi:10.1073/pnas.1512901113.
- Hippenmeyer, Simon, Eline Vrieseling, Markus Sigrist, Thomas Portmann, Celia Laengle, David R Ladle, and Silvia Arber. 2005. “A Developmental Switch in the Response of DRG Neurons to ETS Transcription Factor Signaling.” Edited by Joshua R. Sanes. *PLoS Biology* 3 (5): e159. doi:10.1371/journal.pbio.0030159.

- Horwitz, Rick. 2016. "Integrated, Multi-Scale, Spatial-temporal Cell Biology – A next Step in the Post Genomic Era." *Methods* 96 (March): 3–5. doi:10.1016/j.ymeth.2015.09.007.
- Huang, Yifei, Mark P. Brandon, Amy L. Griffin, Michael E. Hasselmo, and Uri T. Eden. 2009. "Decoding Movement Trajectories Through a T-Maze Using Point Process Filters Applied to Place Field Data from Rat Hippocampal Region CA1." *Neural Computation* 21 (12): 3305–34. doi:10.1162/neco.2009.10-08-893.
- Huberfeld, Gilles, Liset Menendez de la Prida, Johan Pallud, Ivan Cohen, Michel Le Van Quyen, Claude Adam, Stéphane Clemenceau, Michel Baulac, and Richard Miles. 2011. "Glutamatergic Pre-Ictal Discharges Emerge at the Transition to Seizure in Human Epilepsy." *Nature Neuroscience* 14 (5): 627–34. doi:10.1038/nn.2790.
- James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 103. Springer Texts in Statistics. New York, NY: Springer New York.
- James, Nicholas. 2016. "Cholinergic Modulation of Auditory and Prefrontal Cortical Interactions." Boston University.
- Jiruska, Premysl, Marco de Curtis, John G. R. Jefferys, Catherine A. Schevon, Steven J. Schiff, and Kaspar Schindler. 2013. "Synchronization and Desynchronization in Epilepsy: Controversies and Hypotheses: Synchronization in Epilepsy." *The Journal of Physiology* 591 (4): 787–97. doi:10.1113/jphysiol.2012.239590.
- Jolivet, R. 2004. "Generalized Integrate-and-Fire Models of Neuronal Activity Approximate Spike Trains of a Detailed Model to a High Degree of Accuracy." *Journal of Neurophysiology* 92 (2): 959–76. doi:10.1152/jn.00190.2004.
- Jolivet, Renaud, and Wulfram Gerstner. 2004. "Predicting Spike Times of a Detailed Conductance-Based Neuron Model Driven by Stochastic Spike Arrival." *Journal of Physiology-Paris* 98 (4-6): 442–51. doi:10.1016/j.jphysparis.2005.09.010.
- Kawaguchi, Yasuo, and Satoru Kondo. 2002. "Parvalbumin, Somatostatin and Cholecystikinin as Chemical Markers for Specific GABAergic Interneuron Types in the Rat Frontal Cortex." *Journal of Neurocytology* 31 (3-5): 277–87.

- Kelly, Ryan C., and Robert E. Kass. 2012. "A Framework for Evaluating Pairwise and Multiway Synchrony Among Stimulus-Driven Neurons." *Neural Computation* 24 (2009): 2007–32. doi:10.1162/NECO_a_00307.
- Kelly, Ryan C., Robert E. Kass, Matthew A. Smith, and Tai Sing Lee. 2010. "Accounting for Network Effects in Neuronal Responses Using L1 Regularized Point Process Models." *Advances in Neural Information Processing Systems* 23 (2): 1099–1107.
- Kerr, Jason N. D., and Winfried Denk. 2008. "Imaging in Vivo: Watching the Brain in Action." *Nature Reviews Neuroscience* 9 (3): 195–205. doi:10.1038/nrn2338.
- Kim, Sanggyun, David Putrino, Soumya Ghosh, and Emery N. Brown. 2011. "A Granger Causality Measure for Point Process Models of Ensemble Neural Spiking Activity." Edited by Karl J. Friston. *PLoS Computational Biology* 7 (3): e1001110. doi:10.1371/journal.pcbi.1001110.
- Kleen, Jonathan K., Rod C. Scott, Gregory L. Holmes, David W. Roberts, Melissa M. Rundle, Markus Testorf, Pierre-Pascal Lenck-Santini, and Barbara C. Jobst. 2013. "Hippocampal Interictal Epileptiform Activity Disrupts Cognition in Humans." *Neurology* 81 (1): 18–24.
- Kramer, M. A., and S. S. Cash. 2012. "Epilepsy as a Disorder of Cortical Network Organization." *The Neuroscientist* 18 (4): 360–72. doi:10.1177/1073858411422754.
- Krook-Magnuson, Esther, and Ivan Soltesz. 2015. "Beyond the Hammer and the Scalpel: Selective Circuit Control for the Epilepsies." *Nature Neuroscience* 18 (3): 331–38. doi:10.1038/nn.3943.
- Lado, Fred A., and Solomon L. Moshé. 2008. "How Do Seizures Stop?" *Epilepsia* 49 (10): 1651–64. doi:10.1111/j.1528-1167.2008.01669.x.
- Latimer, Kenneth W., E. J. Chichilnisky, Fred Rieke, and Jonathan W. Pillow. 2014. "Inferring Synaptic Conductances from Spike Trains with a Biophysically Inspired Point Process Model." In *Advances in Neural Information Processing Systems*, 954–62.
- Latimer, Kenneth W., Jacob L. Yates, Miriam LR Meister, Alexander C. Huk, and Jonathan W. Pillow. 2015. "Single-Trial Spike Trains in Parietal Cortex Reveal Discrete Steps during Decision-Making." *Science* 349 (6244): 184–87.

- Lawhern, Vernon, Wei Wu, Nicholas Hatsopoulos, and Liam Paninski. 2010. "Population Decoding of Motor Cortical Activity Using a Generalized Linear Model with Hidden States." *Journal of Neuroscience Methods* 189 (2): 267–80. doi:10.1016/j.jneumeth.2010.03.024.
- Lepage, Kyle Q., Georgia G. Gregoriou, Mark A. Kramer, Mikio Aoi, Stephen J. Gotts, Uri T. Eden, and Robert Desimone. 2013. "A Procedure for Testing across-Condition Rhythmic Spike-Field Association Change." *Journal of Neuroscience Methods* 213 (1): 43–62. doi:10.1016/j.jneumeth.2012.10.010.
- Lepage, Kyle Q., Mark a. Kramer, and Uri T. Eden. 2011. "The Dependence of Spike Field Coherence on Expected Intensity." *Neural Computation* 23 (9): 2209–41. doi:10.1162/NECO_a_00169.
- Lüttjohann, Annika, Jan-Mathijs Schoffelen, and Gilles van Luijtelaar. 2014. "Termination of Ongoing Spike-Wave Discharges Investigated by Cortico-thalamic Network Analyses." *Neurobiology of Disease* 70 (October): 127–37. doi:10.1016/j.nbd.2014.06.007.
- Lynn, Henry S. 2003. "Suppression and Confounding in Action." *The American Statistician* 57 (1): 58–61. doi:10.1198/0003130031090.
- Lytton, William W. 2008. "Computer Modelling of Epilepsy." *Nature Reviews Neuroscience* 9 (8): 626–37. doi:10.1038/nrn2416.
- MacKay, Donald M., and Warren S. McCulloch. 1952. "The Limiting Information Capacity of a Neuronal Link." *The Bulletin of Mathematical Biophysics* 14 (2): 127–35.
- Markram, Henry, Eilif Muller, Srikanth Ramaswamy, Michael W. Reimann, Marwan Abdellah, Carlos Aguado Sanchez, Anastasia Ailamaki, et al. 2015. "Reconstruction and Simulation of Neocortical Microcircuitry." *Cell* 163 (2): 456–92. doi:10.1016/j.cell.2015.09.029.
- Martinez, Joan. 2015. "Inhibitory Dynamics in the Superficial Medial Entorhinal Cortex." University of Utah.
- McCullagh, Peter, and John A Nelder. 1989. *Generalized Linear Models, 2nd Ed.* Vol. 37. CRC Press.

- McFarland, James M., Yuwei Cui, and Daniel A. Butts. 2013. "Inferring Nonlinear Neuronal Computation Based on Physiologically Plausible Inputs." Edited by Matthias Bethge. *PLoS Computational Biology* 9 (7): e1003143. doi:10.1371/journal.pcbi.1003143.
- Mensi, S., R. Naud, C. Pozzorini, M. Avermann, C. C. H. Petersen, and W. Gerstner. 2012. "Parameter Extraction and Classification of Three Cortical Neuron Types Reveals Two Distinct Adaptation Mechanisms." *Journal of Neurophysiology* 107 (6): 1756–75. doi:10.1152/jn.00408.2011.
- Merricks, Edward M., Elliot H. Smith, Guy M. McKhann, Robert R. Goodman, Lisa M. Bateman, Ronald G. Emerson, Catherine A. Schevon, and Andrew J. Trevelyan. 2015. "Single Unit Action Potentials in Humans and the Effect of Seizure Activity." *Brain* 138 (10): 2891–2906. doi:10.1093/brain/awv208.
- Moore, Christopher I., Marie Carlen, Ulf Knoblich, and Jessica A. Cardin. 2010. "Neocortical Interneurons: From Diversity, Strength." *Cell* 142 (2): 184–88. doi:10.1016/j.cell.2010.07.005.
- Moore, G P, J P Segundo, D H Perkel, and H Levitan. 1970. "Statistical Signs of Synaptic Interaction in Neurons." *Biophysical Journal* 10 (9): 876–900. doi:10.1016/S0006-3495(70)86341-X.
- Nakahara, Hiroyuki, and Shun-ichi Amari. 2002. "Information-Geometric Measure for Neural Spikes." *Neural Computation* 14 (10): 2269–2316. doi:10.1162/08997660260293238.
- Narayanan, N. S., Eyal Y. Kimchi, and Mark Laubach. 2005. "Redundancy and Synergy of Neuronal Ensembles in Motor Cortex." *Journal of Neuroscience* 25 (17): 4207–16. doi:10.1523/JNEUROSCI.4697-04.2005.
- Nowak, L, and J Bullier. 2000. "Cross-Correlations for Neuronal Spike Trains. Different Types of Temporal Correlation in Neocortex, Their Origin and Significance." In *Time and the Brain, Conceptual Advances in Brain Research*, edited by R Miller, 53–96. Hardwood Academic.
- Nowak, L. G., M. H. J. Munk, A. C. James, Pascal Girard, and J. Bullier. 1999. "Cross-Correlation Study of the Temporal Interactions between Areas V1 and V2 of the Macaque Monkey." *Journal of Neurophysiology* 81 (3): 1057–74.

- Okatan, Murat, Matthew A. Wilson, and Emery N. Brown. 2005. "Analyzing Functional Connectivity Using a Network Likelihood Model of Ensemble Neural Spiking Activity." *Neural Computation* 17 (9): 1927–61.
- Optican, Lance M., Timothy J. Gawne, Barry J. Richmond, and Joseph, P.J. 1991. "Unbiased Measures of Transmitted Information and Channel Capacity from Multivariate Neuronal Data." *Biological Cybernetics* 65: 305–10.
- Ostojic, S., N. Brunel, and V. Hakim. 2009. "How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains." *Journal of Neuroscience* 29 (33): 10234–53. doi:10.1523/JNEUROSCI.1275-09.2009.
- Ozer, Irma. 1991. "Images of Epilepsy in Literature." *Epilepsia* 32.
- Padmanabhan, Krishnan, and Nathaniel N Urban. 2010. "Intrinsic Biophysical Diversity Decorrelates Neuronal Firing While Increasing Information Content." *Nature Neuroscience* 13 (10): 1276–82. doi:10.1038/nn.2630.
- Paninski, L., Matthew R Fellows, Nicholas G Hatsopoulos, and John P Donoghue. 2004a. "Spatiotemporal Tuning of Motor Cortical Neurons for Hand Position and Velocity." *Journal of Neurophysiology* 91 (1): 515–32. doi:10.1152/jn.00587.2002.
- Paninski, Liam, Jonathan W. Pillow, and Eero P. Simoncelli. 2004b. "Maximum Likelihood Estimation of a Stochastic Integrate-and-Fire Neural Encoding Model." *Neural Computation* 16 (12): 2533–61.
- Paninski, Liam, Jonathan Pillow, and Eero Simoncelli. 2005. "Comparing Integrate-and-Fire Models Estimated Using Intracellular and Extracellular Data." *Neurocomputing* 65-66 (June): 379–85. doi:10.1016/j.neucom.2004.10.032.
- Paninski, Liam, Jonathan Pillow, and Jeremy Lewi. 2007. "Statistical Models for Neural Encoding, Decoding, and Optimal Stimulus Design." In *Progress in Brain Research*, 165:493–507. Elsevier. .
- Park, Il Memming, Miriam L R Meister, Alexander C Huk, and Jonathan W Pillow. 2014. "Encoding and Decoding in Parietal Cortex during Sensorimotor Decision-Making." *Nature Neuroscience* 17 (10): 1395–1403. doi:10.1038/nn.3800.

- Paz, Jeanne T, Thomas J Davidson, Eric S Frechette, Bruno Delord, Isabel Parada, Kathy Peng, Karl Deisseroth, and John R Huguenard. 2012. "Closed-Loop Optogenetic Control of Thalamus as a Tool for Interrupting Seizures after Cortical Injury." *Nature Neuroscience* 16 (1): 64–70. doi:10.1038/nn.3269.
- Perucca, P., F. Dubeau, and J. Gotman. 2014. "Intracranial Electroencephalographic Seizure-Onset Patterns: Effect of Underlying Pathology." *Brain* 137 (1): 183–96. doi:10.1093/brain/awt299.
- Pillow, Jonathan W., Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M. Litke, E. J. Chichilnisky, and Eero P. Simoncelli. 2008. "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population." *Nature* 454 (7207): 995–99. doi:10.1038/nature07140.
- Pillow, Jonathan W., and Eero P. Simoncelli. 2006. "Dimensionality Reduction in Neural Models: An Information-Theoretic Generalization of Spike-Triggered Average and Covariance Analysis." *Journal of Vision* 6 (4): 9–9.
- Pillow, J. W. 2005. "Prediction and Decoding of Retinal Ganglion Cell Responses with a Probabilistic Spiking Model." *Journal of Neuroscience* 25 (47): 11003–13. doi:10.1523/JNEUROSCI.3305-05.2005.
- Pnevmatikakis, Eftychios, Josh Merel, Ari Pakman, Liam Paninski, and others. 2013. "Bayesian Spike Inference from Calcium Imaging Data." In *Signals, Systems and Computers, 2013 Asilomar Conference on*, 349–53. IEEE.
- Prince, David, and Barry W. Connors. 1986. "Mechanisms of Interictal Epileptogenesis." *Advances in Neurology* 44: 275–99.
- Quiñan Quiroga, Rodrigo, and Stefano Panzeri. 2009. "Extracting Information from Neuronal Populations: Information Theory and Decoding Approaches." *Nature Reviews Neuroscience* 10 (3): 173–85. doi:10.1038/nrn2578.
- Quinn, Christopher J., Todd P. Coleman, Negar Kiyavash, and Nicholas G. Hatsopoulos. 2011. "Estimating the Directed Information to Infer Causal Relationships in Ensemble Neural Spike Train Recordings." *Journal of Computational Neuroscience* 30 (1): 17–44. doi:10.1007/s10827-010-0247-2.

- Rolls, Edmund T., and Alessandro Treves. 2011. "The Neuronal Encoding of Information in the Brain." *Progress in Neurobiology* 95 (3): 448–90. doi:10.1016/j.pneurobio.2011.08.002.
- Rolls, E. T. 2003. "An Information Theoretic Approach to the Contributions of the Firing Rates and the Correlations Between the Firing of Neurons." *Journal of Neurophysiology* 89 (5): 2810–22. doi:10.1152/jn.01070.2002.
- Rudy, Bernardo, Gordon Fishell, SooHyun Lee, and Jens Hjerling-Leffler. 2011. "Three Groups of Interneurons Account for Nearly 100% of Neocortical GABAergic Neurons." *Developmental Neurobiology* 71 (1): 45–61. doi:10.1002/dneu.20853.
- Rummel, Christian, Marc Goodfellow, Heidemarie Gast, Martinus Hauf, Frédérique Amor, Alexander Stibal, Luigi Mariani, Roland Wiest, and Kaspar Schindler. 2013. "A Systems-Level Approach to Human Epileptic Seizures." *Neuroinformatics* 11 (2): 159–73. doi:10.1007/s12021-012-9161-2.
- Sabolek, H. R., W. B. Swiercz, K. P. Lillis, S. S. Cash, G. Huberfeld, G. Zhao, L. Ste. Marie, et al. 2012. "A Candidate Mechanism Underlying the Variance of Interictal Spike Propagation." *Journal of Neuroscience* 32 (9): 3009–21. doi:10.1523/JNEUROSCI.5853-11.2012.
- Sarma, Sridevi V, Ming Cheng, Rollin Hu, Ziv Williams, Emery N Brown, and Emad N Eskandar. 2008. "Modeling Neural Spiking Activity in the Sub-Thalamic Nucleus of Parkinson's Patients and Healthy Primates." In *The International Federation of Automatic Control*, 9081–86.
- Sarma, Sridevi V, David P Nguyen, Gabriela Czanner, Sylvia Wirth, Matthew a Wilson, Wendy Suzuki, and Emery N Brown. 2011. "Computing Confidence Intervals for Point Process Models." *Neural Computation* 23 (11): 2731–45. doi:10.1162/NECO_a_00198.
- Schevon, Catherine A., Sau K. Ng, Joshua Cappell, Robert R. Goodman, Guy McKhann, Allen Waziri, Almut Branner, Alexandre Sosunov, Charles E. Schroeder, and Ronald G. Emerson. 2008. "Microphysiology of Epileptiform Activity in Human Neocortex." *Journal of Clinical Neurophysiology* 25 (6): 321–30. doi:10.1097/WNP.0b013e31818e8010.
- Schevon, Catherine A., Shennan A. Weiss, Guy McKhann, Robert R. Goodman, Rafael Yuste, Ronald G. Emerson, and Andrew J. Trevelyan. 2012. "Evidence of an

- Inhibitory Restraint of Seizure Activity in Humans.” *Nature Communications* 3 (September): 1060. doi:10.1038/ncomms2056.
- Schiff, Steven J, Tim Sauer, Rohit Kumar, and Steven L Weinstein. 2005. “Neuronal Spatiotemporal Pattern Discrimination : The Dynamical Evolution of Seizures.” *NeuroImage* 28: 1043–55. doi:10.1016/j.neuroimage.2005.06.059.
- Schindler, Kaspar A, Stephan Bialonski, Marie-therese Horstmann, Christian E Elger, Klaus Lehnertz, Kaspar A Schindler, Stephan Bialonski, and Marie-therese Horstmann. 2008. “Evolving Functional Network Properties and Synchronizability during Human Epileptic Seizures Evolving Functional Network Properties and Synchronizability during Human Epileptic Seizures.” *Chaos: An Interdisciplinary Journal of Nonlinear Science* 18: 033119. doi:10.1063/1.2966112.
- Schindler, Kaspar, Christian E. Elger, and Klaus Lehnertz. 2007. “Increasing Synchronization May Promote Seizure Termination: Evidence from Status Epilepticus.” *Clinical Neurophysiology* 118 (9): 1955–68. doi:10.1016/j.clinph.2007.06.006.
- Schneidman, Elad. 2016. “Towards the Design Principles of Neural Population Codes.” *Current Opinion in Neurobiology* 37 (April): 133–40. doi:10.1016/j.conb.2016.03.001.
- Schneidman, Elad, Michael J. Berry, Ronen Segev, and William Bialek. 2006. “Weak Pairwise Correlations Imply Strongly Correlated Network States in a Neural Population.” *Nature* 440 (7087): 1007–12. doi:10.1038/nature04701.
- Schneidman, Elad, William Bialek, and Michael J. Berry II. 2003. “Synergy, Redundancy, and Independence in Population Codes.” *Journal of Neuroscience* 23 (37): 11539–53.
- Scholvin, Jorg, Justin P. Kinney, Jacob G. Bernstein, Caroline Moore-Kochlacs, Nancy Kopell, Clifton G. Fonstad, and Edward S. Boyden. 2016. “Close-Packed Silicon Microelectrodes for Scalable Spatially Oversampled Neural Recording.” *IEEE Transactions on Biomedical Engineering* 63 (1): 120–30. doi:10.1109/TBME.2015.2406113.

- Seung, H. Sebastian, and Uygur Smbl. 2014. "Neuronal Cell Types and Connectivity: Lessons from the Retina." *Neuron* 83 (6): 1262–72.
doi:10.1016/j.neuron.2014.08.054.
- Shannon, Claude Elwood. 1948. "A Mathematical Theory of Communication." *The Bell System Technical Journal* 27: 379–423, 623–56.
- Simoncelli, Eero P, Liam Paninski, Jonathan Pillow, and Odelia Schwartz. 2004. "Characterization of Neural Responses with Stochastic Stimuli 1 Reverse Correlation." In *The New Cognitive Neurosciences*, edited by M Gazzaniga, 3rd ed.
- Smith, Elliot H., Jyun-you Liou, Tyler S. Davis, Edward M. Merricks, Spencer S. Kellis, Shennan A. Weiss, Bradley Greger, et al. 2016. "The Ictal Wavefront Is the Spatiotemporal Source of Discharges during Spontaneous Human Seizures." *Nature Communications* 7 (March): 11098. doi:10.1038/ncomms11098.
- Smith, Gregory D, and S Murray Sherman. 2002. "Detectability of Excitatory versus Inhibitory Drive in an Integrate-and-Fire-or-Burst Thalamocortical Relay Neuron Model." *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 22 (23): 10242–50.
- Spencer, Susan S., Jung Kim, Nihal DeLanerolle, and Dennis D. Spencer. 1999. "Differential Neuronal and Glial Relations with Parameters of Ictal Discharge in Mesial Temporal Lobe Epilepsy." *Epilepsia* 40 (6): 708–12.
- Staley, Kevin, Jennifer Hellier, and F. Edward Dudek. 2005. "Do Interictal Spikes Drive Epileptogenesis?" *The Neuroscientist* 11 (4): 272–76.
doi:10.1177/1073858405278239.
- Stanley, Garrett B. 2013. "Reading and Writing the Neural Code." *Nature Neuroscience* 16 (3): 259–63. doi:10.1038/nn.3330.
- Stevens, Charles F., and Anthony Zador. 1996. "When Is an Integrate-and-Fire Neuron like a Poisson Neuron?" *Advances in Neural Information Processing Systems*, 103–9.
- Stufflebeam, Steven M., Hesheng Liu, Jorge Sepulcre, Naoaki Tanaka, Randy L. Buckner, and Joseph R. Madsen. 2011. "Localization of Focal Epileptic Discharges Using Functional Connectivity Magnetic Resonance Imaging: Clinical

Article.” *Journal of Neurosurgery* 114 (6): 1693–97.
doi:10.3171/2011.1.JNS10482.

- Taniguchi, Hiroki, Miao He, Priscilla Wu, Sangyong Kim, Raehum Paik, Ken Sugino, Duda Kvitsani, et al. 2011. “A Resource of Cre Driver Lines for Genetic Targeting of GABAergic Neurons in Cerebral Cortex.” *Neuron* 71 (6): 995–1013. doi:10.1016/j.neuron.2011.07.026.
- Tibshirani, Robert J., Guenther Walther, and Trevor Hastie. 2001. “Estimating the Number of Clusters in a Data Set via the Gap Statistic.” *Journal Of The Royal Statistical Society B* 63 (2): 411–23.
- Timofeev, Igor, François Grenier, and Mircea Steriade. 2004. “Contribution of Intrinsic Neuronal Factors in the Generation of Cortically Driven Electrographic Seizures.” *Journal of Neurophysiology* 92: 1133–43. doi:10.1152/jn.00523.2003.
- Toledo-Rodriguez, M., Barak Blumenfeld, Caizhi Wu, Junyi Luo, Bernard Attali, Philip Goodman, and Henry Markram. 2004. “Correlation Maps Allow Neuronal Electrical Properties to Be Predicted from Single-Cell Gene Expression Profiles in Rat Neocortex.” *Cerebral Cortex* 14 (12): 1310–27. doi:10.1093/cercor/bhh092.
- Topolnik, Lisa, Mircea Steriade, and Igor Timofeev. 2003. “Partial Cortical Deafferentation Promotes Development of Paroxysmal Activity.” *Cerebral Cortex* 13 (8): 883–93.
- Traub, R. D., Diego Contreras, Mark O. Cunningham, Hilary Murray, Fiona E. N. LeBeau, Anita K. Roopun, Andrea Bibbig, W. Bryan Wilent, Michael J. Higley, and Whittington, Miles A. 2005a. “Single-Column Thalamocortical Network Model Exhibiting Gamma Oscillations, Sleep Spindles, and Epileptogenic Bursts.” *Journal of Neurophysiology* 93 (4): 2194–2232. doi:10.1152/jn.00983.2004.
- Traub, Roger D., Diego Contreras, and M. A. Whittington. 2005b. “Combined Experimental/simulation Studies of Cellular and Network Mechanisms of Epileptogenesis in Vitro and in Vivo.” *Journal of Clinical Neurophysiology : Official Publication of the American Electroencephalographic Society* 22 (5): 330–42.

- Tripathy, S. J., K. Padmanabhan, R. C. Gerkin, and N. N. Urban. 2013. "Intermediate Intrinsic Diversity Enhances Neural Population Coding." *Proceedings of the National Academy of Sciences* 110 (20): 8248–53. doi:10.1073/pnas.1221214110.
- Truccolo, W., Uri T Eden, Matthew R Fellows, John P Donoghue, and Emery N Brown. 2005. "A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects." *Journal of Neurophysiology* 93 (2): 1074–89. doi:10.1152/jn.00697.2004.
- Truccolo, Wilson, Jacob A Donoghue, Leigh R Hochberg, Emad N Eskandar, Joseph R Madsen, William S Anderson, Emery N Brown, Eric Halgren, and Sydney S Cash. 2011. "Single-Neuron Dynamics in Human Focal Epilepsy." *Nature Neuroscience* 14 (5): 635–41. doi:10.1038/nn.2782.
- Ventura, Valerie, Roberto Carta, Robert E Kass, Sonya N Gettner, and Carl R Olson. 2002. "Statistical Analysis of Temporal Evolution in Single-Neuron Firing Rates." *Biostatistics* 3 (1): 1–20.
- Victor, Jonathan D. 2002. "Binless Strategies for Estimation of Information from Neural Data." *Physical Review E* 66 (5). doi:10.1103/PhysRevE.66.051903.
- Vidne, Michael, Yashar Ahmadian, Jonathon Shlens, Jonathan W. Pillow, Jayant Kulkarni, Alan M. Litke, E. J. Chichilnisky, Eero Simoncelli, and Liam Paninski. 2012. "Modeling the Impact of Common Noise Inputs on the Network Activity of Retinal Ganglion Cells." *Journal of Computational Neuroscience* 33 (1): 97–121. doi:10.1007/s10827-011-0376-2.
- Vogelstein, Joshua T., Brendon O. Watson, Adam M. Packer, Rafael Yuste, Bruno Jerny, and Liam Paninski. 2009. "Spike Inference from Calcium Imaging Using Sequential Monte Carlo Methods." *Biophysical Journal* 97 (2): 636–55. doi:10.1016/j.bpj.2008.08.005.
- Wagner, Fabien B., Emad N. Eskandar, G. Rees Cosgrove, Joseph R. Madsen, Andrew S. Blum, N. Stevenson Potter, Leigh R. Hochberg, Sydney S. Cash, and Wilson Truccolo. 2015. "Microscale Spatiotemporal Dynamics during Neocortical Propagation of Human Focal Seizures." *NeuroImage* 122 (November): 114–30. doi:10.1016/j.neuroimage.2015.08.019.

- Wang, Suhua, Enhui Shi, Lizhen Zhou, and Xunli Su. 2009. "Topological Transitivity and Chaos of Group Actions on Dendrites." *International Journal of Bifurcation and Chaos* 19 (12): 4165. doi:10.1142/S0218127409025274.
- Warland, David K., Pamela Reinagel, and Markus Meister. 1997. "Decoding Visual Information from a Population of Retinal Ganglion Cells." *Journal of Neurophysiology* 78 (5): 2336–50.
- Zadeh, Amir Hassan, and Ramesh Sharda. 2015. "Hawkes Point Processes for Social Media Analytics." In *Reshaping Society through Analytics, Collaboration, and Decision Support*, edited by Lakshmi S. Iyer and Daniel J. Power, 18:51–66. Cham: Springer International Publishing.
- Zammit-Mangion, Andrew, Michael Dewar, Visakan Kadirkamanathan, and Guido Sanguinetti. 2012. "Point Process Modelling of the Afghan War Diary." *Proceedings of the National Academy of Sciences* 109 (31): 12414–19.
- Zhou, Pengcheng, Shawn D. Burton, Adam C. Snyder, Matthew A. Smith, Nathaniel N. Urban, and Robert E. Kass. 2015a. "Establishing a Statistical Link between Network Oscillations and Neural Synchrony." Edited by Olaf Sporns. *PLOS Computational Biology* 11 (10): e1004549. doi:10.1371/journal.pcbi.1004549.
- Zhou, Zhengyi, David S. Matteson, Dawn B. Woodard, Shane G. Henderson, and Athanasios C. Micheas. 2015b. "A Spatio-Temporal Point Process Model for Ambulance Demand." *Journal of the American Statistical Association* 110 (509): 6–15.

CURRICULUM VITAE

GRANT FIDDYMENT

677 Beacon St., Boston, MA 02215 neurocoding05@gmail.com Born 1987

STATEMENT

I am a computational neuroscientist interested in applying machine learning to understand the brain. Currently I build predictive models of spiking in clinical conditions (epilepsy) and experimental conditions (in vitro stimulation). I'm also interested and experienced in simulating complex systems, e.g. epileptic networks.

EDUCATION

- AUGUST 2016 PhD in COMPUTATIONAL NEUROSCIENCE, **Boston University**, Boston, MA
Thesis: "Point process modeling as a framework to dissociate intrinsic and extrinsic components in neural systems" | Advisors: Mark Kramer, Uri Eden
GPA: 3.87/4
- MAY 2011 Master of Arts in MATHEMATICS, **University of Georgia**, Athens, GA
Exams in numerical analysis & complex analysis | Advisor: Caner Kazanci
GPA: 3.9/4
- MAY 2009 Bachelor of Science in MATHEMATICS
Bachelor of Arts in ENGLISH
Magna cum laude with high honors, **University of Georgia**, Athens, GA
GPA: 3.86/4

AWARDS AND HONORS

- 2016 Finalist, ESPN Hackathon Open Division, Sloan Sports Analytics Conference 2016
2015 Participant, Summer Workshop on the Dynamic Brain, Allen Institute for Neuroscience
2015 Pre-Doctoral Training Grant: "Data-Driven Modeling of Seizure Termination"
#330118, Epilepsy Foundation
2013 CompNet Travel Award, Boston University
2011 Graduate Medical Sciences Scholarship, Boston University
2010 NSF VIGRE Internship
2009 Graduate Teaching Assistantship, University of Georgia
2006 Phi Beta Kappa, University of Georgia
2006 Hollingsworth Award, Dept. of Mathematics, University of Georgia
2005 Charter Scholarship, University of Georgia
2005 National Merit Scholar

PUBLICATIONS AND POSTERS

- Fiddyment GM, James NM, Martinet LE, Cash SS, Eden UT, Kramer MA (2016). Point process modeling reveals a unique type of enhancement during human seizures. *Journal of Neuroscience Methods* (in prep).
- Martinet LE, Fiddyment GM, Madsen JR, Eskandar EN, Eden UT, Cash SS, Kramer MA (2016). Human seizures couple across spatial scales through traveling wave dynamics. *Nature Communications* (submitted).
- Fiddyment GM, Martinet LE, Madsen J, Eskandar E, Eden UT, Cash SS, Kramer MA (2016). Two categories of ictal discharges propagate with different spatiotemporal dynamics during human seizure. *Journal of Neuroscience* (submitted).
- Fiddyment GM, Sokolowski S, Kramer MA (2015). Functional Network Observations of Diseased Brain States. In Jaeger D, Jung R (Eds.) *Encyclopedia of Computational Neuroscience*, Vol. 1, pp xx-xx. Springer New York Heidelberg Dordrecht London.
- Fiddyment, GM, Ahmed O, Martinet LE, Eden UT, Cash SS, Kramer MA (2015). Statistical and computational modeling of meso- and microscale human seizure. *Society for Neuroscience Abstracts*, 2015.

- Fiddymment, GM, Ahmed O, Martinet LE, Eden UT, Cash SS, Kramer MA (2015). Point process modeling of human seizure. Statistical Analysis of Neural Data (SAND 7) Conference Abstracts.
- Fiddymment, GM, Yazdanbakhsh A (2013). Large-scale, anatomically-constrained simulation of the visual hierarchy. ICCNS Abstracts, 2013.

WORK EXPERIENCE

SEPT 2011-AUG 2016	<p>Graduate Researcher at BOSTON UNIVERSITY <i>Data Analysis Group of M. Kramer & U. Eden</i></p> <p>Designed statistical analyses and software to analyze neural data, particularly during human seizures. Worked with a team of mathematicians and epilepsy clinicians to collect, process, analyze, and interpret data. Research was awarded doctoral training grant from Epilepsy Foundation.</p>
AUG 2009-MAY 2011	<p>Graduate Teaching Assistant at UNIVERSITY OF GEORGIA</p> <p>Instructor for Precalculus (MATH 1113, Fall) and Calculus (MATH 2200, Spring). Lectured, wrote quizzes and tests, and met with students in office hours. Teaching was accompanied by a weekly developmental seminar taught by a senior faculty member.</p>
SUMMER 2010	<p>NSF VIGRE Research Internship at UNIVERSITY OF GEORGIA <i>Traveling wave dynamics in a zebrafish seizure model with A. Sornborger</i></p> <p>Implemented simulations of leaky integrate-and-fire neurons in MATLAB. Simulated and visualized calcium dynamics to analyze an in vivo zebrafish model of seizure. Discovered bifurcation structure where increasing applied current leads to recurrent voltage patterns in a small patch of simulated tissue.</p>
JAN 2006-MAY 2009	<p>Grader at UNIVERSITY OF GEORGIA</p> <p>Graded weekly homework for a variety of math classes including Multivariable Calculus, Sequences & Series, and Intro to Higher Mathematics.</p>
MAY-DEC 2007	<p>Online Author & Editor at HOUGHTON MIFFLIN COMPANY</p> <p>As online content author, used Maple software to write interactive problems. Went on to edit problems and text for a series of middle school, high school, and college math textbooks.</p>
SUMMER 2006	<p>Co-director at MATHNASIUM</p> <p>Taught mathematics in an innovative, gym-style learning environment: tutored, gave assessments, planned curricula, managed finances, and designed recruitment materials.</p>

COMPUTER SKILLS

Advanced Knowledge: Matlab, Python, R, Linux/Unix, \LaTeX , Maple
 Basic Knowledge: Flask, SQL, Javascript, Mathematica, Perl, Java, C, C++

INTERESTS AND ACTIVITIES

Brain-computer interfaces, adaptive brain stimulation
 Numerical analysis & computer algorithms, quantum computing
[Sports analytics](#), [basketball](#)
 Ultimate Frisbee, Saxophone